

Oracle® Communications

Cloud Native Core, Service Communication Proxy Troubleshooting Guide



Release 26.1.202
G60250-02
August 2026

ORACLE®

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

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

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

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

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

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

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

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

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

Contents

Preface

Documentation Accessibility	i
Diversity and Inclusion	i
Conventions	i

1 Introduction

1.1 References	1
----------------	---

2 Troubleshooting Service Communication Proxy

2.1 Generic Checklist	1
2.2 Helm Install Failure	2
2.2.1 Incorrect Image Name in the ocsdp-custom-values Files	3
2.2.2 Docker Registry is Incorrectly Configured	3
2.2.3 Continuous Restart of Pods	3
2.3 Custom Value File Parse Failure	4
2.4 Curl HTTP2 Not Supported	4
2.5 SCP DB goes into the Deadlock State	5
2.6 Helm Test Error Scenarios	6
2.7 Using Debug Tool	7
2.7.1 Prerequisites to Use the Debug Tool	8
2.7.2 Running the Debug Tool	14
2.7.3 Debug Tool Configuration Parameters	15
2.7.4 Tools Tested in Debug Container	18
2.8 Logs	23
2.8.1 Collecting Logs	24
2.8.2 Understanding Logs	25
2.9 Verifying the Availability of Multus Container Network Interface	27
2.10 Upgrade or Rollback Failure	28
2.11 Error Messages for Mediation Rule Configuration	28
2.12 Errors Observed on Grafana and OCI Dashboards	32
2.13 TLS Connection Failure	33
2.14 Recording Rules Creation	36

2.14.1	Applying Recording Rules	43
2.14.2	Troubleshooting Stale Recording Rules and Missing Fresh Samples	45

3 Alerts

3.1	Configuring Alerts	2
3.2	System level alerts	2
3.2.1	SCPNotificationPodMemoryUsage	2
3.2.2	SCPWorkerPodMemoryUsage	3
3.2.3	SCPInstanceDown	4
3.3	Application level alerts	5
3.3.1	SCPCcaFeatureEnabledWithoutHttps	5
3.3.2	SCPIngressTrafficRateAboveMinorThreshold	6
3.3.3	SCPIngressTrafficRateAboveMajorThreshold	7
3.3.4	SCPIngressTrafficRateAboveCriticalThreshold	8
3.3.5	SCPRoutingFailedForProducer	9
3.3.6	SCPAuditErrorResponse	10
3.3.7	SCPAuditEmptyNFArrayResponse	10
3.3.8	DuplicateLocalityFoundInForeignNF	11
3.3.9	ForeignNFLocalityNotServed	12
3.3.10	UnknownLocalityFoundInForeignNF	13
3.3.11	SCPUpsstreamResponseTimeout	13
3.3.12	SCPSingleNfInstanceAvailableForNFType	14
3.3.13	SCPNoNfInstanceForNFType	15
3.3.14	SCPIngressTrafficRateExceededConfiguredLimit	16
3.3.15	SCPIngressTrafficRoutedWithoutRateLimitTreatment	16
3.3.16	SCPEgressTrafficRateExceededConfiguredLimit	17
3.3.17	SCPEgressTrafficRoutedWithoutRateLimitTreatment	18
3.3.18	SCPNotificationRejectTopologySourceLocal	19
3.3.19	SCPNotificationProcessingFailureForNF	19
3.3.20	SCPSubscriptionFailureForNFType	21
3.3.21	SCPReSubscriptionFailureForNFType	23
3.3.22	SCPNrRegistrationFailureForRegionOrSetId	23
3.3.23	SCPNrHeartbeatFailureForRegionOrSetId	25
3.3.24	SCPDOOperationFailure	25
3.3.25	SCPGeneratedErrorsResponseForNFService	26
3.3.26	SCPCircuitBreakingAppliedForNF	26
3.3.27	SCPUgradeStarted	27
3.3.28	SCPUgradeFailed	27
3.3.29	SCPUgradeSuccessful	28
3.3.30	SCPRollbackStarted	28
3.3.31	SCPRollbackFailed	29

3.3.32	SCPRollbackSuccessful	29
3.3.33	ScpWorkerPodCpuUtilizationAboveWarnThreshold	30
3.3.34	ScpWorkerPodCpuUtilizationAboveMinorThreshold	31
3.3.35	ScpWorkerPodCpuUtilizationAboveMajorThreshold	31
3.3.36	ScpWorkerPodCpuUtilizationAboveCriticalThreshold	32
3.3.37	SCPUnhealthyPeerSCPDetected	33
3.3.38	SCPDnsSrvQueryFailure	33
3.3.39	SCPProducerOverloadThrottled	34
3.3.40	SCPProducerOverloadAlternateRouted	35
3.3.41	SCPSeppNotConfigured	35
3.3.42	SCPSeppRoutingFailed	36
3.3.43	SCPGlobalEgressRLRemoteParticipantConnectivityFailure	37
3.3.44	SCPGlobalEgressRLRemoteParticipantWithDuplicateNFInstanceId	38
3.3.45	SCPMediationConnectivityFailure	39
3.3.46	SCPNotificationQueuesUtilizationAboveMinorThreshold	40
3.3.47	SCPNotificationQueuesUtilizationAboveMajorThreshold	41
3.3.48	SCPNotificationQueuesUtilizationAboveCriticalThreshold	42
3.3.49	SCPnrfProxyQueuesUtilizationAboveMinorThreshold	42
3.3.50	SCPnrfProxyQueuesUtilizationAboveMajorThreshold	43
3.3.51	SCPnrfProxyQueuesUtilizationAboveCriticalThreshold	44
3.3.52	SCPWorkerQueuesUtilizationAboveMinorThreshold	44
3.3.53	SCPWorkerQueuesUtilizationAboveMajorThreshold	45
3.3.54	SCPWorkerQueuesUtilizationAboveCriticalThreshold	46
3.3.55	SCPCacheQueuesUtilizationAboveMinorThreshold	46
3.3.56	SCPCacheQueuesUtilizationAboveMajorThreshold	47
3.3.57	SCPCacheQueuesUtilizationAboveCriticalThreshold	48
3.3.58	SCPLoadManagerQueuesUtilizationAboveMinorThreshold	48
3.3.59	SCPLoadManagerQueuesUtilizationAboveMajorThreshold	49
3.3.60	SCPLoadManagerQueuesUtilizationAboveCriticalThreshold	50
3.3.61	SCPProducerNfSetUnhealthy	50
3.3.62	SCPPeerSeppUnhealthy	51
3.3.63	SCPMicroServiceUnreachable	52
3.3.64	SCPTrafficFeedSendFailed	53
3.3.65	SCPTrafficFeedKafkaClusterUnhealthy	53
3.3.66	SCPTrafficFeedPartitionUnhealthy	54
3.3.67	SCPServiceMeshFailure	54
3.3.68	SCPHealthCheckFailedForPeerSCP	55
3.3.69	SCPHealthCheckFailed	55
3.3.70	ScpWorkerPodPendingTransUtilizationAboveWarnThreshold	56
3.3.71	ScpWorkerPodPendingTransUtilizationAboveMinorThreshold	56
3.3.72	ScpWorkerPodPendingTransUtilizationAboveMajorThreshold	57
3.3.73	ScpWorkerPodPendingTransUtilizationAboveCriticalThreshold	58

3.3.74	ScpWorkerPodResourceUtilizationAboveMinorThreshold	58
3.3.75	ScpWorkerPodResourceUtilizationAboveMajorThreshold	59
3.3.76	ScpWorkerPodResourceUtilizationAboveWarnThreshold	59
3.3.77	ScpWorkerPodResourceUtilizationAboveCriticalThreshold	60
3.3.78	SCPDNSSRVNRFMigrationTaskFailure	61
3.3.79	SCPDNSSRVNRFNonMigrationTaskFailure	61
3.3.80	SCPDNSSRVNRFDuplicateTargetDetected	62
3.3.81	SCPHighResponseTimeFromProducer	62
3.3.82	SCPCGroupVersionDetectionFailed	63
3.3.83	SCPCPUUsageFileReadFailed	63
3.3.84	SCPIgnoreUnknownService	64
3.3.85	SCPWorkerSSLCertificateOnCriticalExpiry	64
3.3.86	SCPWorkerSSLCertificateOnMajorExpiry	65
3.3.87	SCPWorkerSSLCertificateOnMinorExpiry	65
3.3.88	SCPIngressConnectionEstablishmentFailure	66
3.3.89	SCPEgressConnectionEstablishmentFailure	66
3.3.90	SCPNrfProxyOAuthQueuesUtilizationAboveMinorThreshold	67
3.3.91	SCPNrfProxyOAuthQueuesUtilizationAboveMajorThreshold	67
3.3.92	SCPNrfProxyOAuthQueuesUtilizationAboveCriticalThreshold	68
3.3.93	ScpSelfOCIThresholdAboveWarn	68
3.3.94	ScpSelfOCIThresholdAboveMinor	69
3.3.95	ScpSelfOCIThresholdAboveMajor	69
3.3.96	ScpSelfOCIThresholdAboveCritical	70
3.3.97	SCPWorkerSSLPartialSecret	70
3.3.98	SCPWorkerAndNFTimeSyncFailure	71
3.3.99	SCPCoherenceCacheHitsRateAboveThreshold	72
3.3.100	SCPNotificationRequestsRateAboveThreshold	72
3.3.101	SCPLciRxRateAboveThreshold	73
3.3.102	SCPOciRxRateAboveThreshold	73
3.3.103	SCPTrafficRateAboveRatedThreshold	74
3.3.104	SCPMaxNFProfilesLimitExceeded	74
3.3.105	SCPCustomRoamingProxyRoutingNotFound	75
3.3.106	SCPPeerCustomRoamingProxyUnhealthy	75
3.3.107	SCPCustomRoamingProxyMaxViaHeaderProcessingCountExceeded	76
3.3.108	SCPCustomRoamingProxyRoutingFailed	77
3.3.109	SCPWorkerNFSelectionHeaderMatchListDuplicates	77
3.3.110	ScpWorkerServiceRequestRoutingInfoAbsent	78
3.3.111	ScpWorkerNotificationRequestRoutingInfoAbsent	79
3.3.112	SCPWorkerNFProfileSyncFailure	79
3.3.113	SCPnfTopologyForceRefreshAuditFailure	80
3.3.114	SCPnfGroupIdDiscoveryFailure	80
3.3.115	SCPAuditDurationExceeded	81

3.3.116	SCPSelfProfileUpdateSuccess	82
3.3.117	SCPEgressRateLimitActiveKeyCountAboveCriticalThreshold	82
3.3.118	SCPEgressRateLimitActiveKeyCountAboveMajorThreshold	83
3.3.119	SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold	84
3.3.120	SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold	85
3.3.121	SCPIngressRateLimitActiveKeyCountAboveMajorThreshold	86
3.3.122	SCPIngressRateLimitActiveKeyCountAboveCriticalThreshold	87
3.3.123	SCPGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold	88
3.3.124	SCPGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold	89
3.3.125	SCPNotificationDBCertificateExpiryAboveMinorThreshold	90
3.3.126	SCPNotificationDBCertificateExpiryAboveCriticalThreshold	91
3.3.127	SCPNotificationKeystoreExpireOutOfTolerance	92
3.3.128	SCPNotificationDBCertificateExpiryAboveMajorThreshold	93
3.3.129	SCPSubscriptionDBCertificateExpiryAboveCriticalThreshold	94
3.3.130	SCPSubscriptionDBCertificateExpiryAboveMajorThreshold	95
3.3.131	SCPSubscriptionDBCertificateExpiryAboveMinorThreshold	96
3.3.132	SCPSubscriptionKeystoreExpireOutOfTolerance	97
3.3.133	SCPAuditDBCertificateExpiryAboveMinorThreshold	98
3.3.134	SCPAuditDBCertificateExpiryAboveCriticalThreshold	99
3.3.135	SCPAuditKeystoreExpireOutOfTolerance	100
3.3.136	SCPAuditDBCertificateExpiryAboveMajorThreshold	101
3.3.137	SCPConfigurationDBCertificateExpiryAboveCriticalThreshold	102
3.3.138	SCPConfigurationDBCertificateExpiryAboveMajorThreshold	103
3.3.139	SCPConfigurationDBCertificateExpiryAboveMinorThreshold	104
3.3.140	SCPConfigurationKeystoreExpireOutOfTolerance	105
3.3.141	SCPAlternateResolutionDBCertificateExpiryAboveMajorThreshold	106
3.3.142	SCPAlternateResolutionDBCertificateExpiryAboveMinorThreshold	107
3.3.143	SCPAlternateResolutionKeystoreExpireOutOfTolerance	108
3.3.144	SCPAlternateResolutionDBCertificateExpiryAboveCriticalThreshold	109
3.4	Mediation Alerts	110
3.4.1	NFMediationUserDefinedVariablesMaxSizeLimitExceeded	110

Preface

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

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 customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

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 user supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that user enter.

My Oracle Support

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

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

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

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

Acronyms

Table 1 Acronyms

Acronym	Meaning
ASM	Aspen Service Mesh
CNE	Oracle Communications Cloud Native Core, Cloud Native Environment
NRF	Oracle Communications Cloud Native Core, Network Repository Function
OCI	Oracle Cloud Infrastructure
OHC	Oracle Help Center
OSDC	Oracle Software Delivery Cloud
SCP	Oracle Communications Cloud Native Core, Service Communication Proxy
SCPC	Service Communication Proxy Control Plane
SVC	Services

What's New in This Guide

This section introduces the documentation updates for release 26.1.2xx.

Release 26.1.202 - G60250-02, August 2026

Updated the release version to 26.1.202 throughout the document.

Release 26.1.201 - G60250-01, August 2026

- Updated the release version to 26.1.201 throughout the document.
- Added the [Recording Rules Creation](#) section.
- Added the following alerts:
 - [SCPSelfProfileUpdateSuccess](#)
 - [SCPNotificationDBCertificateExpiryAboveMinorThreshold](#)
 - [SCPNotificationDBCertificateExpiryAboveCriticalThreshold](#)
 - [SCPNotificationKeystoreExpireOutOfTolerance](#)
 - [SCPNotificationDBCertificateExpiryAboveMajorThreshold](#)
 - [SCPSubscriptionDBCertificateExpiryAboveCriticalThreshold](#)
 - [SCPSubscriptionDBCertificateExpiryAboveMajorThreshold](#)
 - [SCPSubscriptionDBCertificateExpiryAboveMinorThreshold](#)
 - [SCPSubscriptionKeystoreExpireOutOfTolerance](#)
 - [SCPAuditDBCertificateExpiryAboveMinorThreshold](#)
 - [SCPAuditDBCertificateExpiryAboveCriticalThreshold](#)
 - [SCPAuditKeystoreExpireOutOfTolerance](#)
 - [SCPAuditDBCertificateExpiryAboveMajorThreshold](#)
 - [SCPConfigurationDBCertificateExpiryAboveCriticalThreshold](#)
 - [SCPConfigurationDBCertificateExpiryAboveMajorThreshold](#)
 - [SCPConfigurationDBCertificateExpiryAboveMinorThreshold](#)
 - [SCPConfigurationKeystoreExpireOutOfTolerance](#)
 - [SCPAlternateResolutionDBCertificateExpiryAboveMajorThreshold](#)
 - [SCPAlternateResolutionDBCertificateExpiryAboveMinorThreshold](#)
 - [SCPAlternateResolutionKeystoreExpireOutOfTolerance](#)
 - [SCPAlternateResolutionDBCertificateExpiryAboveCriticalThreshold](#)
 - [SCPnfgGroupIdDiscoveryFailure](#)
 - [SCPAuditDurationExceeded](#)
- Added the following alerts as part of the Rate Limiting enhancements:
 - [SCPGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold](#)
 - [SCPGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold](#)
 - [SCPEgressRateLimitActiveKeyCountAboveMajorThreshold](#)

-
- [SCPIngressRateLimitActiveKeyCountAboveMajorThreshold](#)
 - [SCPIngressRateLimitActiveKeyCountAboveCriticalThreshold](#)
 - [SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold](#)
 - [SCPEgressRateLimitActiveKeyCountAboveCriticalThreshold](#)
 - [SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold](#)
 - Updated the Description row for the [SCPnfGroupIdDiscoveryFailure](#) alert.

1

Introduction

This document provides information about Oracle Communications Cloud Native Core, Service Communication Proxy troubleshooting scenarios.

Overview

Service Communication Proxy (SCP) is a decentralized solution composed of control plane and data plane. SCP is deployed with 5G Network Functions (NFs) for providing routing control, resiliency, and observability to the core network.

SCP is deployed either as a default outbound proxy to NF instances or as a router model, where SCP is configured as HTTP2 outbound proxy at each NF in cloud native environments.

This guide provides information about resolving problems you might experience while installing, configuring, and upgrading SCP. This document also provides information about tools available to help you collect and analyze diagnostic data.

Note

The performance and capacity of the SCP system may vary based on the call model, feature or interface configuration, network conditions, and underlying CNE and hardware environment.

1.1 References

- *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, Service Communication Proxy User Guide*
- *Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, Network Repository Function User Guide*
- *Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide*
- *External reference: Drools User Guide*

2

Troubleshooting Service Communication Proxy

This section provides information to troubleshoot common errors that occur while installing and upgrading Service Communication Proxy (SCP).

Note

kubectl commands might vary based on the platform deployment. Replace kubectl with Kubernetes environment-specific command line tool to configure Kubernetes resources through kube-api server. The instructions provided in this document are as per the Oracle Communications Cloud Native Environment (OCCNE) version of kube-api server.

2.1 Generic Checklist

The following sections provide a generic checklist for troubleshooting tips.

Deployment Related Tips

Perform the following checks before the deployment:

- Are OCSCP deployment, pods, and services created, running, and available?
To check this, run the following command:

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

Inspect the output and check the following columns:

- READY, STATUS, and RESTARTS
- PORT(S)

- Is the correct image used and the correct environment variables set in the deployment?
To check this, run the following command:

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

- Check if the microservices can access each other through REST interface.
To check this, run the following command:

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

Example:

```
kubectl exec -it ocscp-scp-subscription-6bf9b7d69f-qvcnn -n ocscp curl
http://ocscp-scp-notification:8082/ocscp/scpc-notification/v2/
compositecustomobjects
```

```
kubectl exec -it ocscp-scp-notification-dd5c74869-cswkb -n ocscp curl
http://ocscp-scp-configuration:8081/ocscp/scpc-configuration/v1/
systemoptions/sse
```

Note

These commands are in their simplest form and display the logs only if scpc-notification and scpc-configuration pods are deployed.

The list of URIs for all the microservices:

- scp-worker: `http://ocscp-scp-worker:8000/hostNFMapper`
- scpc-configuration: `http://ocscp-scp-configuration:8081/ocscp/scpc-configuration/v1/systemoptions/sse` Or any other configuration URI from SWAGGER-UI
- scpc-notification: `http://ocscp-scp-notification:8082/ocscp/scpc-notification/v2/compositecustomobjects`
- scpc-subscription: `http://ocscp-scp-subscription:8080/ocscp/scpc-subscription/v1/appstate`

Application Related Tips

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

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

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

Example:

```
# kubectl -n scp-svc logs -f $(kubectl -n scp-svc get pods -o name|cut -d '/' -
f 2|grep nfr)
```

Note

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

2.2 Helm Install Failure

This section describes Helm installation failure scenarios.

2.2.1 Incorrect Image Name in the ocscp-custom-values Files

Problem

`helm install` might fail if incorrect image name is provided in the `ocscp_values.yaml` file.

Error Code or Error Message

When you run `kubectl get pods -n <ocscp_namespace>`, the status of the pods might be `ImagePullBackOff` or `ErrImagePull`.

Solution

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

1. Edit the `ocscp_values.yaml` file and provide release specific image name and tags.
2. Run the `helm install` command.
3. Run the `kubectl get pods -n <ocscp_namespace>` command to verify if the status of all the pods is **Running**.

2.2.2 Docker Registry is Incorrectly Configured

Problem

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

Error Code or Error Message

When you run `kubectl get pods -n <ocscp_namespace>`, the status of the pods might be `ImagePullBackOff` or `ErrImagePull`.

Solution

Configure docker registry on all primary and secondary nodes.

For information about docker registry configuration, see *Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide*.

2.2.3 Continuous Restart of Pods

Problem

`helm install` might fail if MySQL primary and secondary hosts may not be configured properly in `ocscp-custom-values.yaml`.

Error Code/Error Message

When you run `kubectl get pods -n <ocscp_namespace>`, the pods restart count increases continuously.

Solution

MySQL servers may not be configured properly. For more information about the MySQL configuration, see *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*.

2.3 Custom Value File Parse Failure

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

Problem

Unable to parse the ocscp_values-x.x.x.yaml file while running Helm install.

Error Code or Error Message

Error: failed to parse ocscp_values-x.x.x.yaml: error converting YAML to JSON: yaml

Symptom

When parsing the ocscp-custom-values-x.x.x.yaml file, if the above mentioned error is received, it indicates that the file is not parsed because of the following reasons:

- The tree structure may not have been followed
- There may be a tab spaces in the file

Solution

Download the ocscp_csar_23_2_0_0_0.zip folder from MOS and complete the steps as described in *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*.

2.4 Curl HTTP2 Not Supported

Problem

curl http2 is not supported on the system.

Error Code or Error Message

Unsupported protocol error is thrown or connection is established with HTTP/1.1 200 OK

Symptom

If unsupported protocol error is thrown or connection is established with http1.1, it is an indication that curl http2 support may not be present on your machine.

Solution

① Note

You must check the software platform policies before executing the following procedure.

Following is the procedure to install curl with HTTP2 support:

1. Run the following command to ensure that Git is installed:

```
$ sudo yum install git -y
```

2. Run the following commands to install nghttp2:

```
$ git clone https://github.com/tatsuhiro-t/nghttp2.git
$ cd nghttp2

$ autoreconf -i
$ automake
$ autoconf

$ ./configure
$ make
$ sudo make install

$ echo '/usr/local/lib' > /etc/ld.so.conf.d/custom-libs.conf

$ ldconfig
```

3. Run the following commands to install the latest Curl:

```
$ wget http://curl.haxx.se/download/curl-7.46.0.tar.bz2 (NOTE: Check for
latest version during Installation)
$ tar -xvjf curl-7.46.0.tar.bz2
$ cd curl-7.46.0

$ ./configure --with-nghttp2=/usr/local --with-ssl

$ make

$ sudo make install

$ sudo ldconfig
```

4. Run the following command to ensure that HTTP2 is added in features:

```
$ curl --http2-prior-knowledge -v "<http://10.75.204.35:32270/nscp-
disc/v1/nf-instances?requester-nf-type=AMF&target-nf-type=SMF>"
```

2.5 SCP DB goes into the Deadlock State

Problem

MySQL locks get struck.

Error Code/Error Message

ERROR 1213 (40001): Deadlock found when trying to get lock; try restarting transaction.

Symptom

Unable to access MySQL.

Solution

Following is the procedure to remove the deadlock:

1. Run the following command on each SQL node:

```
SELECT
CONCAT('KILL ', id, ';')
FROM INFORMATION_SCHEMA.PROCESSLIST
WHERE `User` = <DbUsername>
AND `db` = <DbName>;
```

This command retrieves the list of commands to kill each connections.

Example:

```
select
  CONCAT('KILL ', id, ';')
FROM INFORMATION_SCHEMA.PROCESSLIST
where `User` = 'scpuser'
      AND `db` = 'ocscpdb';
+-----+
| CONCAT('KILL ', id, ';') |
+-----+
| KILL 204491;              |
| KILL 200332;              |
| KILL 202845;              |
+-----+
3 rows in set (0.00 sec)
```

2. Run the kill command on each SQL node.

2.6 Helm Test Error Scenarios

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

Note

You must clean up any residual job from the SCP deployment before performing the following procedure.

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

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

2. Check for the Helm Test pod that is in the error state:

Figure 2-1 Helm test pod

NAME	READY	STATUS	RESTARTS	AGE
ocscp-scp-worker-64c7d6bc7f-vgc7v	1/1	Running	0	45m
ocscp-scp-audit-7f6f4578b6-xwh5q	0/1	Running	0	45m
ocscp-scp-configuration-7dd7bb878c-w8pgz	1/1	Running	0	45m
ocscp-scp-notification-58d489985c-7ngpz	0/1	Running	0	45m
ocscp-scp-pilot-7fd894f574-7bvjd	1/1	Running	0	45m
ocscp-scp-subscription-74977bb95c-9kfs9	0/1	Running	0	45m
ocscp-test-lmtr8	0/1	Error	0	5m32s

- Run the following command to retrieve the logs:

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

Readiness Probe Failure

Helm install might fail due to the readiness probe URL failure.

In case the following error appears, check for readiness probe URL correctness in the particular microservice helm charts under charts folder:

Figure 2-2 Readiness Probe Failure

```
{
  "logEvent": "Access URL:http://192.168.140.235:8091/audit/v1/actuator/health/auditReadiness unsuccessful, Exception Details: ",
  "Timestamp": "20-09-21 07:29:25.860+0000",
  "Application": "ocscp",
  "Engineering version": "77.88.88",
  "Marketing version": "1.8.0.0",
  "Microservice": "test",
  "Cluster": "ocscp",
  "Namespace": "scpsvc",
  "Node": "master",
  "Pod": "ocscp-test-lmtr8"}^M
{"logEvent": "Readiness check failed for URL: http://192.168.140.235:8091/audit/v1/actuator/health/auditReadiness, PodName: ocs
cp-scp-audit-7f6f4578b6-xwh5q", "Timestamp": "20-09-21 07:29:25.864+0000", "Application": "ocscp", "Engineering version": "77.88.88
", "Marketing version": "1.8.0.0", "Microservice": "test", "Cluster": "ocscp", "Namespace": "scpsvc", "Node": "master", "Pod": "ocscp-te
st-lmtr8"}^M
{"logEvent": "Access URL:http://192.168.140.244:8091/notification/v1/actuator/health/notificationReadiness unsuccessful, Except
ion Details: ", "Timestamp": "20-09-21 07:29:30.899+0000", "Application": "ocscp", "Engineering version": "77.88.88", "Marketing vers
ion": "1.8.0.0", "Microservice": "test", "Cluster": "ocscp", "Namespace": "scpsvc", "Node": "master", "Pod": "ocscp-test-lmtr8"}^M
{"logEvent": "Readiness check failed for URL: http://192.168.140.244:8091/notification/v1/actuator/health/notificationReadiness
, PodName: ocscp-scp-notification-58d489985c-7ngpz", "Timestamp": "20-09-21 07:29:30.900+0000", "Application": "ocscp", "Engineeri
ng version": "77.88.88", "Marketing version": "1.8.0.0", "Microservice": "test", "Cluster": "ocscp", "Namespace": "scpsvc", "Node": "ma
ster", "Pod": "ocscp-test-lmtr8"}^M
{"logEvent": "Access URL:http://192.168.219.172:8091/subscription/v1/actuator/health/subscriptionReadiness unsuccessful, Except
ion Details: ", "Timestamp": "20-09-21 07:29:35.906+0000", "Application": "ocscp", "Engineering version": "77.88.88", "Marketing vers
ion": "1.8.0.0", "Microservice": "test", "Cluster": "ocscp", "Namespace": "scpsvc", "Node": "master", "Pod": "ocscp-test-lmtr8"}^M
{"logEvent": "Readiness check failed for URL: http://192.168.219.172:8091/subscription/v1/actuator/health/subscriptionReadiness
, PodName: ocscp-scp-subscription-74977bb95c-9kfs9", "Timestamp": "20-09-21 07:29:35.908+0000", "Application": "ocscp", "Engineeri
ng version": "77.88.88", "Marketing version": "1.8.0.0", "Microservice": "test", "Cluster": "ocscp", "Namespace": "scpsvc", "Node": "ma
ster", "Pod": "ocscp-test-lmtr8"}^M
```

Low Resources

Helm install might fail due to low resource and following error may appear:

Figure 2-3 Low Resource

```
{
  "logEvent": "Pod check failed, current state: Pending, PodName: ocscp-scp-worker-84dd7448f7-72kwz", "Timestamp": "20-09-21 07:45
:05.414+0000", "Application": "ocscp", "Engineering version": "77.88.88", "Marketing version": "1.8.0.0", "Microservice": "test", "Cl
uster": "ocscp", "Namespace": "scpsvc", "Node": "slave1", "Pod": "ocscp-test-s6p5g"}^M
```

In this case, check the CPU and memory availability in the Kubernetes cluster.

2.7 Using Debug Tool

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

- tcpdump

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

2.7.1 Prerequisites to Use the Debug Tool

This section describes the prerequisites for using the debug tool.

① Note

- For CNE 23.2.0 and later versions, follow [Step 1](#).
- For CNE versions prior to 23.2.0, follow [Step 2](#).
- The debug tool requires security context with the following permissions:

```
securityContext:
  allowPrivilegeEscalation: true
  capabilities:
    drop:
      - ALL
    add:
      - NET_RAW
      - NET_ADMIN
  runAsUser: <user-id>
volumeMounts:
  - mountPath: /tmp/tools
    name: debug-tools-dir
```

For OpenShift environment, security context constraint must exist to allow above permissions to enable debug tool deployment.

Ensure that you have admin privileges.

1. If you are using CNE 23.2.0 and later versions, add a namespace to an empty resource by running the following command to verify if the current disallow-capabilities cluster policy has namespace in it:

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

Sample output:

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

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

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

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

Example:

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

Sample output:

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

- b. If it is required to remove the namespace added in the above mentioned step, run the following command:

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

Sample output:

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

- c. To add a namespace to an existing namespace list, run the following command to verify if the current disallow-capabilities cluster policy has namespaces in it:

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

Sample output:

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

- d. If namespaces are already added, then patch the policy by running the following command to add <namespace> to the existing list:

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

Example:

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

Sample output:

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

- e. If it is required to remove the namespace added in the above mentioned step, run the following command:

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

Example:

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

Sample output:

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

Note

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

2. If you are using CNE versions prior to 23.2.0, create PodSecurityPolicy (PSP) in the Bastion Host by following these steps:
 - a. Log in to the Bastion Host.
 - b. Run the following command to create a new PSP:

Note

readOnlyRootFilesystem, allowPrivilegeEscalation, and allowedCapabilities parameters are required by the debug container. Other parameters are mandatory for PSP creation and can be customized as per the CNE environment. It is recommended to use the default values.

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

```
apiVersion: policy/v1beta1
kind: PodSecurityPolicy
metadata:
  name: debug-tool-psp
spec:
  readOnlyRootFilesystem: false
  allowPrivilegeEscalation: true
  allowedCapabilities:
    - NET_ADMIN
    - NET_RAW
  fsGroup:
```

```

    ranges:
      - max: 65535
        min: 1
    rule: MustRunAs
  runAsUser:
    rule: MustRunAsNonRoot
  seLinux:
    rule: RunAsAny
  supplementalGroups:
    rule: RunAsAny
  volumes:
    - configMap
    - downwardAPI
    - emptyDir
    - persistentVolumeClaim
    - projected
    - secret
EOF

```

- c. Run the following command to create a role for PSP:

```

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

```

- d. Run the following command to associate the service account for your NF namespace with the role created for the tool PSP:

```

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

```

```
name: system:serviceaccounts
EOF
```

For information about parameters, see [Debug Tool Configuration Parameters](#).

3. To complete the NF specific Helm configurations, do the following:
 - a. Log in to the NF server.
 - b. Run the following command to open the `ocscp_custom_values.yaml` file:

```
$ vim <ocscp_custom_values file>
```

- c. In the global configuration section, add the following:

```
extraContainersTpl: |
  - command:
    - /bin/sleep
    - infinity
    name: tools
    resources:
      requests:
        cpu: "1"
        memory: {{ .Values.global.debugToolContainerMemoryLimit |
quote }}
        ephemeral-storage: "2Gi"
      limits:
        cpu: "1"
        memory: {{ .Values.global.debugToolContainerMemoryLimit |
quote }}
        ephemeral-storage: "4Gi"
    securityContext:
      allowPrivilegeEscalation: true
      capabilities:
        drop:
        - ALL
        add:
        - NET_RAW
        - NET_ADMIN
      runAsUser: <user-id>
    volumeMounts:
    - mountPath: /tmp/tools
      name: debug-tools-dir
```

Note

- The Debug Tool Container has a default user ID: 7000. If the you want to override this default value, it can be done using the `runAsUser` field. Otherwise, the field can be skipped.
Default value: `uid=7000(debugtool) gid=7000(debugtool) groups=7000(debugtool)`
- If you want to customize the container name, replace the `name` field in the above mentioned `values.yaml` with the following:

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

This ensures that the container name is prefixed and suffixed with the required values.

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

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

Note

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

2.7.2 Running the Debug Tool

Perform the following procedure to run the debug tool.

- To enter the debug tool container, run the following command to retrieve the pod details:

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

Example:

```
$ kubectl get pods -n ocscp
```

Sample Output:

NAME	READY	STATUS	RESTARTS
AGE			
ocscp-scp-worker-58c8469595-5d6gc	2/2	Running	0

```

3d12h
ocscp-scpc-audit-d8f5cdc96-4nw7c      2/2      Running    0
3d12h
ocscp-scpc-configuration-774f9bc65b-f9ttn  2/2      Running    0
3d12h
ocscp-scpc-notification-779965766-9ckr5    2/2      Running    0
3d12h
ocscp-scpc-subscription-7bf6d6c884-2r4vp    2/2      Running    0
3d12h

```

2. Run the following command to enter the debug tool container:

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

Example:

```
$ kubectl exec -it ocscp-scpc-notification-779965766-9ckr5 -c tools -n
ocscp bash
```

3. Run the debug tools:

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

Example:

```
bash -4.2$ tcpdump
```

4. Copy the output files from container to host:

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

Example:

```
$ kubectl cp -c tools ocscp-scpc-notification-779965766-9ckr5:/tmp/
capture.pcap -n ocscp /tmp/
```

2.7.3 Debug Tool Configuration Parameters

Following are the parameters used to configure the debug tool.

CNE Parameters

Table 2-1 CNE Parameters

Parameter	Description
apiVersion	Defines the version schema of this representation of an object.
kind	Indicates a string value representing the REST resource this object represents.
metadata	Indicates the metadata of Standard object.

Table 2-1 (Cont.) CNE Parameters

Parameter	Description
metadata.name	Indicates the metadata name that must be unique within a namespace.
spec	Defines the policy enforced.
spec.readOnlyRootFilesystem	Controls whether the containers run with a read-only root filesystem, that is, no writable layer.
spec.allowPrivilegeEscalation	Gates whether or not a user is allowed to set the security context of a container to allowPrivilegeEscalation=true.
spec.allowedCapabilities	Provides a list of capabilities that are allowed to be added to a container.
spec.fsGroup	Controls the supplemental group applied to some volumes. RunAsAny allows any fsGroup ID to be specified.
spec.runAsUser	Controls which user ID the containers are run with. RunAsAny allows any runAsUser to be specified.
spec.seLinux	RunAsAny allows any seLinuxOptions to be specified.
spec.supplementalGroups	Controls which group IDs containers add. RunAsAny allows any supplementalGroups to be specified.
spec.volumes	Provides a list of allowed volume types. The allowed values correspond to the volume sources that are defined when creating a volume.

Role Creation Parameters**Table 2-2 Role Creation**

Parameter	Description
apiVersion	Defines the versioned schema of this representation of an object.
kind	Indicates a string value representing the REST resource this object represents.
metadata	Indicates the metadata of Standard object.
metadata.name	Indicates the name of metadata that must be unique within a namespace.
metadata.namespace	Defines the space within which each name must be unique.
rules	Manages all the PolicyRules for this Role.
apiGroups	Indicates the name of the APIGroup that contains the resources.
rules.resources	Indicates the list of resources this rule applies to.
rules.verbs	Indicates the list of Verbs that applies to ALL the ResourceKinds and AttributeRestrictions contained in this rule.
rules.resourceNames	Indicates an optional white list of names that the rule applies to.

Table 2-3 Role Binding Creation

Parameter	Description
apiVersion	Defines the versioned schema of this representation of an object.
kind	Indicates the string value representing the REST resource this object represents.
metadata	Indicates the metadata of Standard object.
metadata.name	Indicates the name that must be unique within a namespace.
metadata.namespace	Defines the space within which each name must be unique.
roleRef	References a Role in the current namespace or a ClusterRole in the global namespace.
roleRef.apiGroup	Indicates the group for the resource being referenced.
roleRef.kind	Indicates the type of resource being referenced.
roleRef.name	Indicates the name of resource being referenced.
subjects	Manages references to the objects the role applies to.
subjects.kind	Indicates the type of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount".
subjects.apiGroup	Manages the API group of the referenced subject.
subjects.name	Indicates the name of the object being referenced.

Debug Tool Configuration Parameters**Table 2-4 Debug Tool Configuration Parameters**

Parameter	Description
command	Indicates the string array used for container command.
image	Indicates the docker image name.
imagePullPolicy	Indicates the Image Pull Policy.
name	Indicates the name of the container.
resources	Indicates the Compute Resources required by this container.
resources.limits	Describes the maximum amount of compute resources allowed.
resources.requests	Describes the minimum amount of compute resources required.
resources.limits.cpu	Indicates the CPU limits.
resources.limits.memory	Indicates the Memory limits.
resources.limits.ephemeral-storage	Indicates the Ephemeral Storage limits.
resources.requests.cpu	Indicates the CPU requests.
resources.requests.memory	Indicates the Memory requests.
resources.requests.ephemeral-storage	Indicates the Ephemeral Storage requests.
securityContext	Indicates the Security options the container should run with.
securityContext.allowPrivilegeEscalation	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. It directly controls whether the no_new_privs flag to be set on the container process.
securityContext.readOnlyRootFilesystem	Indicates whether this container has a read-only root filesystem. The default value is false.

Table 2-4 (Cont.) Debug Tool Configuration Parameters

Parameter	Description
securityContext.capabilities	Indicates the capabilities to add or drop when running containers. It defaults to the default set of capabilities granted by the container runtime.
securityContext.capabilities.drop	Indicates the removed capabilities.
securityContext.capabilities.add	Indicates the added capabilities.
securityContext.runAsUser	Indicates the UID to run the entrypoint of the container process.
extraContainersTpl.volumeMounts	Indicates that the parameter is used for mounting the volume.
extraContainersTpl.volumeMounts.mountPath	Indicates the path for volume mount.
extraContainersTpl.volumeMounts.name	Indicates the name of the directory for debug tool logs storage.
extraContainersVolumesTpl.name	Indicates the name of the volume for debug tool logs storage.
extraContainersVolumesTpl.emptyDir.medium	Indicates where the emptyDir volume is stored.
extraContainersVolumesTpl.emptyDir.sizeLimit	Indicates the emptyDir volume size.

2.7.4 Tools Tested in Debug Container

The following tables describe the list of debug tools that are tested.

Table 2-5 tcpdump

Options Tested	Description	Output	Capabilities
-D	Print the list of the network interfaces available on the system and on which tcpdump can capture packets.	<pre>tcpdump -D 1. eth0. 2. nflog (Linux netfilter log (NFLOG) interface) 3. nfqueue (Linux netfilter queue (NFQUEUE) interface) 4. any (Pseudo-device that captures on all interfaces) 5. lo [Loopback]</pre>	NET_ADMIN, NET_RAW
-i	Listen on interface.	<pre>tcpdump -i eth0 tcpdump: verbose output suppressed, use -v or -vv for full protocol decoding listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes 12:10:37.381199 IP cncc-core-ingress-gateway-7ffc49bb7f-2khhc.46519 > kubernetes.default.svc.cluster.local.https: Flags [P], seq 1986927241:1986927276, ack 1334332290, win 626, options [nop,nop,TS val 849591834 ecr 849561833], length 3512:10:37.381952 IP cncc-core-ingress-gateway-7ffc49bb7f-2khhc.45868 > kube-dns.kube- system.svc.cluster.local.domain: 62870+ PTR? 1.0.96.10.in- addr.arpa. (40)</pre>	NET_ADMIN, NET_RAW

Table 2-5 (Cont.) tcpdump

Options Tested	Description	Output	Capabilities
-w	Write the raw packets to file rather than parsing and printing them out.	<code>tcpdump -w capture.pcap -i eth0</code>	NET_ADMIN, NET_RAW
-r	Read packets from file, which was created with the -w option.	<code>tcpdump -r capture.pcap</code> reading from file /tmp/capture.pcap, link-type EN10MB (Ethernet)12:13:07.381019 IP cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519 > kubernetes.default.svc.cluster.local.https: Flags [P.], seq 1986927416:1986927451, ack 1334332445, win 626, options [nop,nop,TS val 849741834 ecr 849711834], length 3512:13:07.381194 IP kubernetes.default.svc.cluster.local.https > cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519: Flags [P.], seq 1:32, ack 35, win 247, options [nop,nop,TS val 849741834 ecr 849741834], length 3112:13:07.381207 IP cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519 > kubernetes.default.svc.cluster.local.https: Flags [.], ack 32, win 626, options [nop,nop,TS val 849741834 ecr 849741834], length 0	NET_ADMIN, NET_RAW

Table 2-6 ip

Options Tested	Description	Output	Capabilities
addr show	Look at the protocol addresses.	<code>ip addr show</code> 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group defaultlink/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00inet 127.0.0.1/8 scope host lovalid_lft forever preferred_lft forever2: tunl0@NONE: <NOARP> mtu 1480 qdisc noop state DOWN group defaultlink/loopback 0.0.0.0 brd 0.0.0.0: eth0@if190: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1440 qdisc noqueue state UP group defaultlink/ether aa:5a:27:8d:74:6f brd ff:ff:ff:ff:ff:ff link-netnsid 0inet 192.168.219.112/32 scope global eth0valid_lft forever preferred_lft forever	-
route show	List routes.	<code>ip route show</code> default via 169.254.1.1 dev eth0 169.254.1.1 dev eth0 scope link	-

Table 2-6 (Cont.) ip

Options Tested	Description	Output	Capabilities
addrlabel list	List address labels	<pre>ip addrlabel list prefix ::1/128 label 0 prefix ::/96 label 3 prefix ::ffff:0.0.0.0/96 label 4 prefix 2001::/32 label 6 prefix 2001:10::/28 label 7 prefix 3ffe::/16 label 12 prefix 2002::/16 label 2 prefix fec0::/10 label 11 prefix fc00::/7 label 5 prefix ::/0 label 1</pre>	-

Table 2-7 netstat

Options Tested	Description	Output	Capabilities
-a	Show both listening and non-listening sockets. For TCP, this means established connections.	<pre>netstat -a Active Internet connections (servers and established)Proto Recv-Q Send-Q Local Address Foreign Address Statetcp 0 0 0.0.0.0:tpoxy 0.0.0.0:* LISTENTcp 0 0 0.0.0.0:websm 0.0.0.0:* LISTENTcp 0 0 cncc-core-ingress:websm 10-178-254-194.ku:47292 TIME_WAITtcp 0 0 cncc-core- ingress:46519 kubernetes.default:https ESTABLISHEDtcp 0 0 cncc-core-ingress:websm 10-178-254-194.ku:47240 TIME_WAITtcp 0 0 cncc-core-ingress:websm 10-178-254-194.ku:47347 TIME_WAITudp 0 0 localhost:59351 localhost:ambit-lm ESTABLISHEDActive UNIX domain sockets (servers and established)Proto RefCnt Flags Type State I-Node Pathunix 2 [] STREAM CONNECTED 576064861</pre>	-
-l	Show only listening sockets.	<pre>netstat -l Active Internet connections (only servers)Proto Recv-Q Send-Q Local Address Foreign Address Statetcp 0 0 0.0.0.0:tpoxy 0.0.0.0:* LISTENTcp 0 0 0.0.0.0:websm 0.0.0.0:* LISTENActive UNIX domain sockets (only servers)Proto RefCnt Flags Type State I-Node Path</pre>	-
-s	Display summary statistics for each protocol.	<pre>netstat -s Ip:4070 total packets received0 forwarded0 incoming packets discarded4070 incoming packets delivered4315 requests sent outicmp:0 ICMP messages received0 input ICMP message failed.ICMP input histogram:2 ICMP messages sent0 ICMP messages failedICMP output histogram:destination unreachable: 2</pre>	-
-i	Display a table of all network interfaces.	<pre>netstat -i Kernel Interface tableiface MTU RX-OK RX-ERR RX-DRP RX-OVR TX-OK TX-ERR TX-DRP TX-OVR Flgeth0 1440 4131 0 0 0 4355 0 0 0 BMRUlo 65536 0 0 0 0 0 0 LRU</pre>	-

Table 2-8 curl

Options Tested	Description	Output	Capabilities
-o	Write output to <file> instead of stdout.	curl -o file.txt http://abc.com/file.txt	-
-x	Use the specified HTTP proxy.	curl -x proxy.com:8080 -o http://abc.com/file.txt	-
--http2	Use the specified HTTP/2	curl --http2 -v http://cncc-iam-ingress-gateway.cncc.svc.cluster.local:30085/cncc/auth/admin	-

Table 2-9 ping

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

Table 2-10 nmap

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

Table 2-10 (Cont.) nmap

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

Table 2-10 (Cont.) nmap

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

Table 2-11 dig

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

2.8 Logs

This chapter explains the process to retrieve the logs and status that can be used for effective troubleshooting. SCP provides various sources of information that may be helpful in the troubleshooting process.

Log Levels

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

Supported Log Levels

For SCP, the log level for a microservice can be set to any of the following valid values:

- **DEBUG:** A log level used for events considered to be useful during software debugging when more granular information is required.
- **INFO:** The standard log level indicating that something happened, the application entered a certain state, and so on.
- **WARN:** Indicates that something unexpected happened in the application, a problem, or a situation that might disturb one of the processes. But that doesn't mean that the application has failed. The WARN level should be used in situations that are unexpected, but the code can continue the work.
- **ERROR:** The log level that should be used when the application reaches an issue preventing one or more functionalities from properly functioning.

Configuring Log Levels

To view logging configurations and update logging levels, use the Logging Config option on the Cloud Native Configuration Console. For more information, see "Configuring Global Options" in *Oracle Communications Cloud Native Core, Service Communication Proxy User Guide*.

2.8.1 Collecting Logs

Perform this procedure to collect logs from pods or containers.

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

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

2. Run the following command to collect the logs from the specific pods or containers:

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

Example:

```
$ kubectl logs ocscp-scp-worker-xxxxx -n ocscp
```

3. Run the following command to store the log in a file:

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

Example:

```
$ kubectl logs ocscp-scp-worker-xxxxx -n ocscp > logs.txt
```

4. (Optional) Run the following command for the log stream with file redirection starting with last 100 lines of log:

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

Example:

```
$ kubectl logs ocscp-scp-worker-xxxxx -n ocscp -f --tail 100 > logs.txt
```

2.8.2 Understanding Logs

This section explains the logs required to look into, to handle different SCP debugging issues.

The log level attributes of SCP services are as follows:

- SCPC-Subscription
- SCPC-Notification
- SCP-Worker

Sample Logs

Sample log statement for SCPC-Subscription:

```
{"instant":
{"epochSecond":1614521111,"nanoOfSecond":908545000},"thread":"main","level":"I
NF0","loggerName":"com.oracle.cgbu.cne.scp.soothsayer.subscription.processor.S
ubscriptionDataConsumer","message":{"logMsg=Subscription consumer cycle
completed.Thread will now sleep for given time, cycle=232,
threadSleepTimeInMs=100},"endOfBatch":false,"loggerFqcn":"org.apache.logging.
log4j.spi.AbstractLogger","threadId":30,"threadPriority":5,"messageTimestamp":
"21-06-07 12:38:56.784+0000","application":"ocscp","microservice":"scpc-
subscription","engVersion":"1.12.0","mktgVersion":"1.12.0.0.0","vendor":"orac
le","namespace":"scpsvc","node":"primary","pod":"ocscp-scp-
subscription-7f5b7c8ccd-
z89wb","subsystem":"subscription","instanceType":"prod","processId":"1"}
```

Sample log statement for SCPC-Notification:

```
{"instant":
{"epochSecond":1623069485,"nanoOfSecond":496630558},"thread":"main","level":"I
NF0","loggerName":"com.oracle.cgbu.cne.scp.soothsayer.Process","message":{"log
Msg=Successfully processed notification, nfInstanceId=6faf1bbc-6e4a-2828-a507-
a14ef8e1bc5b,
nfType=NRF},"endOfBatch":false,"loggerFqcn":"org.apache.logging.log4j.s
pi.AbstractLogger","threadId":1,"threadPriority":5,"messageTimestamp":"21-06-07
12:38:05.496+0000","application":"ocscp","microservice":"scpc-
notification","engVersion":"1.12.0","mktgVersion":"1.12.0.0.0","vendor":"orac
le","namespace":"scpsvc","node":"primary","pod":"ocscp-scp-
notification-76597b5b7-
wfmxb","subsystem":"notification","instanceType":"prod","processId":"1"}
```

Sample log statement for SCP-Worker

```
{
  "instant": {
    "epochSecond": 1623069702,
    "nanoOfSecond": 672444454
  },
  "thread": "scp-upstream-worker-7",
  "level": "WARN",
  "loggerName": "com.oracle.cgbu.cne.scp.router.routelayer.MsgRouteChain",
  "message": "MsgRouteChain::sendRsp(): SCP generated Response, Response Code = 503, body = {\"title\": \"Service Unavailable\", \"status\": \"503\", \"detail\": \"Service Unavailable:: Service Unavailable\"}, error category = Destination webclient Connection Failure, ingress request host = ocscp-scp-worker.scpsvc.svc.cluster.local:8000, ingress request path = /nnrf-nfm/v1/subscriptions, ingress 3gpp-sbi-target-apiRoot = http://nnrf1svc.scpsvc.svc.cluster.local:8080, egress request Uri = http://nnrf2svc.scpsvc.svc.cluster.local:8080/nnrf-nfm/v1/subscriptions, egress request Destination = nnrf2svc.scpsvc.svc.cluster.local",
  "endOfBatch": false,
  "loggerFqcn": "org.apache.logging.log4j.spi.AbstractLogger",
  "threadId": 139,
  "threadPriority": 5,
  "messageId": "c49b3b48-afc1-4058-83e8-b719ee181ed8",
  "messageTimestamp": "21-06-07 12:41:42.672+0000",
  "application": "ocscp",
  "microservice": "scp-worker",
  "engVersion": "17.17.17",
  "mktgVersion": "1.12.0.0.0",
  "vendor": "oracle",
  "namespace": "scpsvc",
  "node": "primary",
  "pod": "ocscp-scp-worker-9567767dc-7bqg9",
  "subsystem": "router",
  "instanceType": "prod",
  "processId": "1"
}
```

The following table describes different log attributes:

Table 2-12 Log Attribute Details

Log Attribute	Details	Sample Value	Data Type
instant	Epoch time Note: It is a group of two values such as epochSecond and nanoOfSecond.	{\"epochSecond\":1614521244,\"nanoOfSecond\":775702000}	Object
thread	Logging Thread Name	"pool-4-thread-1"	String
level	Log Level of the log printed	"ERROR"	String
loggerName	Class or Module which printed the log	"com.oracle.cgbu.cne.scp.soothsayer.audit.process.AuditMaster"	String
message	Message related to the log providing brief details Note: WARN log level indicates that SCP connection with NRF is established.	{logMsg=NRF health check did not complete successfully. Next health check will start in given interval, timeIntervalInSec=2}	String
endOfBatch	Log4j2 Internal	false	boolean
loggerFqcn	Log4j2 Internal	org.apache.logging.log4j.spi.AbstractLogger	String
threadId	Thread Id generated internally by Log4j2	31	Integer
threadPriority	Thread Priority set internally by Log4j2	5	Integer
messageTimestamp	Timestamp when log was printed	"21-06-07 12:41:15.277+0000"	String
application	Application name	"ocscp"	String
scpFqdn	SCP FQDN	ocscp-scp-worker.ocscp-thrust5-06.svc.thrust5	String
microservice	Name of the microservice	"scpc-audit"	String

Table 2-12 (Cont.) Log Attribute Details

Log Attribute	Details	Sample Value	Data Type
engVersion	Engineering version of the software	"1.12.0"	String
mktgVersion	Marketing version of the software	"1.12.0.0.0"	String
vendor	Vendor of the software	"oracle"	String
namespace	Namespace where application is deployed	"scpsvc"	String
node	Node where the pod resides	"primary"	String
pod	Pod name of deployed application	"ocscp-scpc-audit-6c5ddb4c54-hf8kr"	String
subsystem	Subsystem inside microservice group	"audit"	String
instanceType	Instance type	"prod"	String
processId	Process ID internally assigned	1	Integer

2.9 Verifying the Availability of Multus Container Network Interface

Perform the following procedure to verify whether the Multus Container Network Interface (CNI) feature is enabled after SCP installation is complete.

Note

Ensure that Multus Container Network Interface configuration is completed as described in *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*.

1. To verify whether the pod contains signaling network, go to Kubernetes cluster and run the following command:

```
kubectl describe pod <pod-name> -n <namespace>
```

2. Check whether the pod output contains the "net1" interface.
Sample pod output:

```
k8s.v1.cni.cncf.io/network-status:
[
  {
    "name": "",
    "ips": [
      "192.168.219.111"
    ],
    "default": true,
    "dns": {}
  }, {
    "name": "scpsvc/macvlan-siga",
    "interface": "net1",
    "ips": [
      "192.68.3.78"
    ],
    "mac": "b6:41:f9:a6:c8:8e",
```

```

      "dns": {}
    }
  }
  k8s.v1.cni.cncf.io/networks: [{ "name": "macvlan-siga"}]

```

2.10 Upgrade or Rollback Failure

When SCP upgrade or rollback fails, perform the following procedure:

1. Check the preupgrade or postupgrade or rollback hook logs in Kibana as applicable. You can filter upgrade or rollback logs using the following filters:

- For upgrade: lifeCycleEvent=9001 or 9011
- For rollback: lifeCycleEvent=9002

Sample Log:

```

{"time_stamp":"2021-11-22
10:28:11.820+0000","thread":"main","level":"INFO","logger":"com.oracle.cgbu
.cne.scp.soothsayer.hooks.releases.ReleaseHelmHook_1_14_1","message":{"logM
sg=Starting Pre-Upgrade hook Execution, lifeCycleEvent=9001 | Upgrade,
sourceRelease=101400,
targetRelease=101500},"loc":"com.oracle.cgbu.cne.scp.soothsayer.common.uti
ls.EventSpecificLogger.submit(EventSpecificLogger.java:94)"}

```

2. Check the pod logs in Kibana to analyze the cause of failure.
3. After detecting the cause of failure, do the following:
 - For upgrade failure:
 - If the cause of upgrade failure is database or network connectivity issue, contact [My Oracle Support](#). When the issue is resolved, rerun the upgrade command.
 - If the cause of failure is not related to database or network connectivity issue and is observed during the preupgrade phase, do not perform rollback because SCP deployment remains in the source or older release.
 - When the preupgrade fails and the error log contains Can't write; duplicate key in table in the preupgrade hooks logs, obtain the dump of both application and backup database of SCP as described in the "SCP Data Corruption" section in *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*, and then contact [My Oracle Support](#).
 - If the upgrade failure occurs during the postupgrade phase, for example, post upgrade hook failure due to target release pod not moving to ready state, then perform a rollback.
 - For rollback failure: If the cause of rollback failure is database or network connectivity issue, contact [My Oracle Support](#). When the issue is resolved, rerun the rollback command.
4. If the issue persists, contact [My Oracle Support](#).

2.11 Error Messages for Mediation Rule Configuration

This section allows you to troubleshoot and resolve problems in Mediation Rules using Drools Rule Language (DRL).

The following types of errors are handled:

- Custom Errors
- Errors specific to Drool library

The custom errors generated by application are as follows:


- Value not valid for type: This error is shown in cases where a field in the request body contains a value that is not expected.
- Fields are required and missing for rules with state: This error is shown in cases where a required field file is missing in the request.
- Rule name already exists in the database: When cloning a rule and the new name exists in the database, this error appears.
- Rule cannot be deleted on status: An APPLIED rule cannot be deleted. In the case of a delete request with APPLIED status, this exception will appear.
- Request body ruleName does not match with param ruleName: If the value of the rule name in the URL does not match the value in the "name" field in the request body, this exception will appear.
- Rule was not found: This error is shown when a get or delete request has a rule name that doesn't exist in the database.
- The State field is not valid at the creation of the rule; it must be SAVE state: The rule is not going to be created if the value of state in the request body is different from SAVE state.
- State transition is not valid for status: For more information, refer "Mediation Rule API Migration" section in the *Oracle Communications Cloud Native Core, Service Communication Proxy User Guide*.
- The APPLIED Status is not valid at the creation of the rule, it must be in the DRAFT status.

 Note

Once a rule is generated, ruleName, Format, and Status fields cannot be changed. Some fields are dependent on Status fields to be able to modify. For more information, refer "Mediation Rule API Migration" section in the *Oracle Communications Cloud Native Core, Service Communication Proxy User Guide*.

For DRL errors, the Drools provide standardized messages in the following format:

- **1st Block:** Indicates the error code.
- **2nd Block:** Indicates the DRL source line and column where the problem occurred.
- **3rd Block:** Indicates the description of the problem.
- **4th Block:** Indicates the Components in the DRL source where the issue occurred, such as rules, functions, and queries.
- **5th Block:** Indicates the DRL source pattern where the issue occurred.

Figure 2-4 Error Message Format

The following standardized error messages are supported by Drools:

- **101: no viable alternative**

Indicates that the parser reached a decision point but was unable to find an alternative.

Example rule with incorrect spelling

```
//incorrect spelling
//incorrect syntax
package com.oracle.cgbu.ocmediation.nfmediation;
import com.oracle.cgbu.ocmediation.factdetails.Request;

dialect "mvel"

rule "ruleTest"
when
  Â req : Request(headers.has("header") == true) //special character
then
  req.headers.add("newHeader", "132465")
end
```

Note

The parser reported an invalid token due to a special character.

- **102: mismatched input**

Indicates that the parser expects a specific input that is not present at the current input point.

Example rule with an incomplete rule statement

```
//incomplete rule statement
// incorrect syntax
package com.oracle.cgbu.ocmediation.nfmediation;
import com.oracle.cgbu.ocmediation.factdetails.Request;

dialect "mvel"

//Must have a rule name

when
  req : Request(headers.has("header1") == true)
```

```
then
  req.headers.add("TEST", "132465")
end
```

Note

The parser encountered an incomplete construct due to a missing rule name.

- **103: failed predicate**

Indicates that a validating semantic predicate is evaluated as false. These semantic predicates are commonly used in DRL files to identify component terms such as declare, rule, exists, not, and others.

Example rule with an invalid keyword

```
package com.oracle.cgbu.ocmediation.nfmediation;
import com.oracle.cgbu.ocmediation.factdetails.Request;

dialect "mvel"

fdsfsdfs

rule "ruleTest"
when
  req : Request(headers.has("header") == true)
then
  req.headers.add("newHeader", "132465")
end
```

Note

This text line is not a DRL keyword construct, hence the parser cannot validate the rest of the DRL file.

- **105: did not match anything**

Indicates the parser reached a grammar sub-rule that must match an alternative at least once but did not match anything.

Example rule with invalid text in an empty condition

```
package com.oracle.cgbu.ocmediation.nfmediation;
import com.oracle.cgbu.ocmediation.factdetails.Request;

dialect "mvel"

rule "ruleTest"
when
  None // Must remove `None` if condition is empty
then
  req.headers.add("TEST", "132465")
end
```

Note

The condition should be empty, but the word None is used. This error is resolved by removing None, which is not a valid DRL keyword.

For more information on the list of errors specific to the Drool library, see *Drools User Guide*.

2.12 Errors Observed on Grafana and OCI Dashboards

This section provides information to troubleshoot errors observed on Grafana and OCI dashboards. These errors occur when the number of records fetched by a query, which is an expression that uses application metrics and dimensions, exceeds the configured limit, due to which charts or widgets are invisible on dashboards.

The following sample error messages are displayed on the:

- Grafana dashboard: execution: execution: query processing would load too many samples into memory in query execution.
- OCI dashboard: Query cannot result in more than 2000 streams

To resolve this issue, perform the following tasks to view the data on these dashboards:

1. On the Grafana or OCI dashboard, minimize the query interval to check if the data appears.
 - a. To debug an issue, check the error logs and select the interval on the dashboard based on the timestamp of the observed error logs to minimize the search results.

This task fetches only records within that query interval and reduces the number of records on the dashboard.

2. You can add more filters using the metric dimensions (as per the traffic being run) to the query to minimize the search results. For more information about metric dimensions, see *Oracle Communications Cloud Native Core, Service Communication Proxy User Guide*. To query a particular metric, these dimensions act as filter keys as defined in [Table 2-13](#) and [Table 2-14](#).
3. Ensure that the metric used in the query is pegged by the SCP application.

The following tables describe examples of queries on Grafana and OCI dashboards:

Table 2-13 Examples of Query on Grafana Dashboard

Metric Used	Default Query on Dashboard	Query after Applying Additional Filters
ocscp_metric_http_rx_req_total	sum(irate(ocscp_metric_http_rx_req_total{namespace="scpsvc"}[2m])) by (pod)	sum(irate(ocscp_metric_http_rx_req_total{namespace="scpsvc",ocscp_nf_service_type="nausf-auth"}[2m])) by (pod)
ocscp_metric_http_tx_req_total	sum(irate(ocscp_metric_http_tx_req_total{namespace="ocscp"}[5m])) by (ocscp_nf_type,ocscp_producer_service_instance_id)	sum(irate(ocscp_metric_http_tx_req_total{namespace="ocscp",ocscp_nf_service_type="nausf-auth"}[5m])) by (ocscp_nf_type,ocscp_producer_service_instance_id)

Table 2-14 Examples of Query on OCI Dashboard

Metric Used	Default Query on Dashboard	Query after Applying Additional Filters
ocscp_metric_http_rx_req_total	ocscp_metric_http_rx_req_total[10m] {k8Namespace="ocscp-2"}.rate().groupBy(podname).sum()	- ocscp_metric_http_rx_req_total[10m] {k8Namespace="ocscp-2",ocscp_nf_service_type="nausf-auth"}.rate().groupBy(podname).sum() - ocscp_metric_http_rx_req_total[10m] {k8Namespace="ocscp-2",ocscp_nf_service_type="nausf-auth", ocscp_consumer_info="smf#smf8svc.ocscp-2.svc.cluster.local#NA"}.rate().groupBy(podname).sum()
ocscp_metric_http_tx_res_total	ocscp_metric_http_tx_res_total[10m] {k8Namespace="ocscp-2"}.rate().groupBy(ocscp_consumer_info).sum()	- ocscp_metric_http_tx_res_total[10m] {k8Namespace="ocscp-2",ocscp_nf_service_type="nausf-auth"}.rate().groupBy(ocscp_consumer_info).sum() - ocscp_metric_http_tx_res_total[10m] {k8Namespace="ocscp-2",ocscp_nf_service_type="nausf-auth", ocscp_consumer_info="smf#smf8svc.ocscp-2.svc.cluster.local#NA"}.rate().groupBy(ocscp_consumer_info).sum()

2.13 TLS Connection Failure

This section describes the TLS related issues and their resolution steps. It is recommended to attempt the resolution steps provided in this guide before contacting Oracle Support:

Problem: Ingress Connection Failure Due to TLS 1.3 Version Mismatch Between Consumer NF and SCP Server

Scenario: The Consumer NF has requested a connection using TLS 1.3, but the SCP server is currently configured to support only TLS 1.2. This version mismatch prevents successful communication between the client and server.

Server Error Message

The client supported protocol versions [TLSv1.3] are not accepted by server preferences [TLSv1.2]

Client Error Message

Received fatal alert: protocol_version

Resolution Steps:

- Update the ocscp-custom-values.yaml file on the SCP server to set the sbiProxySslConfigurations.server.tlsVersion parameters to include TLS 1.3 and TLS 1.2. This configuration applies to fresh installations. For fixing issues or changing the version, you can use the REST API for configuration.
- Verify TLS 1.3 for secure communication between the consumer NF and SCP to ensure that the issue has been resolved.

Problem: Egress Connection Failure Due to Cipher Mismatch: SCP Client and Producer Server for TLS v1.3.

Scenario: The SCP client is configured to request a connection using TLSv1.3 with specific ciphers that are not supported by the Producer server. As a result, the connection fails due to the cipher mismatch, preventing secure communication between the client and server.

Client Error Message

No appropriate protocol(protocol is disabled or cipher suites are inappropriate)

Server Error Message

Received fatal alert: handshake failure

Resolution Steps:

- Ensure that the following cipher suites are configured for the SCP client to use with TLS 1.3:
 - TLS_AES_128_GCM_SHA256
 - TLS_AES_256_GCM_SHA384
 - TLS_CHACHA20_POLY1305_SHA256
- Ensure that both the client and server have at least one common TLS 1.3 cipher configured.
- Verify TLS 1.3 for secure communication between the SCP client and the Producer server to ensure that the issue has been resolved.

Problem: Egress Connection Failure for TLS v1.3 Due to Expired Certificates.

Scenario: The SCP client is attempting to establish an egress connection using TLSv1.3, but the connection fails due to expired certificates. Specifically, the SCP client is presenting TLSv1.3 certificates that have passed their validity period, which causes the Producer server to reject the connection.

Client Error Message

Service Unavailable for producer due to Certificate Expired

Server Error Message

Received fatal alert: handshake failure

Resolution Steps:

- Verify the validity of the current certificate (client_ocscp.cer).
- If the certificate has expired, renew it or extend its validity.
- Attempt to establish a connection between the SCP client and the Producer server to confirm that the issue has been resolved.
- Verify the TLSv1.3 for secure communication.

Problem: Pods are not available after populating the `clientDisabledExtension` or `serverDisabledExtension` parameter.

Resolution Steps:

- Check the value of the `clientDisabledExtension` or `serverDisabledExtension` parameters. The following values cannot be added to these parameters:
 - `supported_versions`
 - `key_share`
 - `supported_groups`
 - `signature_algorithms`
 - `certificate_authorities`

If any of the above values are present, remove them or revert to the default configuration for the pod to start properly.

Note

The `certificate_authorities` value can be added to the `clientDisabledExtension` parameter but cannot be added to the `serverDisabledExtension` parameter.

Problem: Pods are not available after populating the `clientAllowedSignatureSchemes` parameter.

Solution:

- Check the value of the `clientAllowedSignatureSchemes` parameter.
- The following values should be present for this parameter:
 - `ecdsa_secp521r1_sha512`
 - `ecdsa_secp384r1_sha384`
 - `ecdsa_secp256r1_sha256`
 - `ed448`
 - `ed25519`
 - `rsa_pss_rsae_sha512`
 - `rsa_pss_rsae_sha384`
 - `rsa_pss_rsae_sha256`
 - `rsa_pss_pss_sha512`
 - `rsa_pss_pss_sha384`
 - `rsa_pss_pss_sha256`
 - `rsa_pkcs1_sha512`
 - `rsa_pkcs1_sha384`
 - `rsa_pkcs1_sha256`

If any of the above values are not present, add them or revert to the default configuration to ensure the pod starts properly.

2.14 Recording Rules Creation

Create OC-SCP recording rules from dashboard aggregations and migrate Grafana panels to query the precomputed series.

Create recording rules to precompute expensive dashboard aggregations, so Grafana does not repeatedly query high-cardinality raw counters. For OC-SCP KPIs, derive rules from existing dashboard `sum by (...)` dimensions and update Grafana to query the recorded series directly.

The creation of recording rules has the following benefits:

- Reduces dashboard query cost and Prometheus CPU use during refreshes.
- Keeps Grafana expressions aligned with the labels used by existing panels.
- Avoids repeating aggregations that Prometheus has already recorded.
- Improves query stability during high-load periods by using wider rate windows.

Recording Rule Design Principle

Create rules based on labels used in dashboard `sum by (...)` expressions. Preserve `kubernetes_namespace` whenever Grafana needs namespace filtering.

Example source expression:

```
sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="$namespace",
ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (kubernetes_pod_name)
```

Corresponding rule:

```
- record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_kubernetes_p
od_name
  expr: sum by (kubernetes_namespace, kubernetes_pod_name) (
    ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"}
  )
```

Updated Grafana expression:

```
rate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_kuberne
tes_pod_name{kubernetes_namespace="$namespace"}[5m])
```

Note

Do not repeat `sum by (kubernetes_pod_name)` in Grafana: the recording rule already produces one series for each namespace and pod.

Sample Source Metrics

The dashboard uses the following base HTTP counter metrics:

- `ocscp_metric_http_rx_req_total`

- `ocscp_metric_http_tx_req_total`
- `ocscp_metric_http_rx_res_total`
- `ocscp_metric_http_tx_res_total`

Common filters are `kubernetes_namespace="$namespace"`, `ocscp_consumer_host!="OCSCP-INITIATED-MSG"`, `ocscp_response_code_class="5"`, `ocscp_nf_service_type="nnrf-oauth2"`, and `ocscp_response_code="200"`.

Sample Recording Rules

Use a dedicated OC-SCP HTTP rule group. Start with an interval of 1m; increase it to 5m if Prometheus is overloaded or misses rule evaluations.

```
apiVersion: v1
data:
  recording_rules.yml: |
    groups:
      - name: SCPHttpRecordingRules
        interval: 1m
        rules:
          - record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace
      expr: sum by (kubernetes_namespace)
(ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
          - record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_kubernetes_p
od_name
      expr: sum by (kubernetes_namespace, kubernetes_pod_name)
(ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
          - record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_kubernetes_p
od_name_container
      expr: sum by (kubernetes_namespace, kubernetes_pod_name, container)
(ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
          - record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_typ
e
      expr: sum by (kubernetes_namespace, ocscp_nf_type)
(ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
          - record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_typ
e_ocscp_req_msg_size
      expr: sum by (kubernetes_namespace, ocscp_nf_type,
ocscp_req_msg_size) (ocscp_metric_http_rx_req_total{ocscp_consumer_host!
="OCSCP-INITIATED-MSG"})
          - record:
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_autho
rity
      expr: sum by (kubernetes_namespace, ocscp_authority)
(ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
          - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace
      expr: sum by (kubernetes_namespace) (ocscp_metric_http_tx_req_total)
          - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_kubernetes_p
od_name_container
```

```

        expr: sum by (kubernetes_namespace, kubernetes_pod_name, container)
(ocscp_metric_http_tx_req_total)
    - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type
e
        expr: sum by (kubernetes_namespace, ocscp_nf_type)
(ocscp_metric_http_tx_req_total)
    - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type
ocscp_nf_end_point
        expr: sum by (kubernetes_namespace, ocscp_nf_type,
ocscp_nf_end_point) (ocscp_metric_http_tx_req_total)
    - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_produc
er_host
        expr: sum by (kubernetes_namespace, ocscp_producer_host)
(ocscp_metric_http_tx_req_total)
    - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_end
_point
        expr: sum by (kubernetes_namespace, ocscp_nf_end_point)
(ocscp_metric_http_tx_req_total)
    - record:
ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_ser
vice_type
        expr: sum by (kubernetes_namespace, ocscp_nf_service_type)
(ocscp_metric_http_tx_req_total)
    - record:
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace
        expr: sum by (kubernetes_namespace) (ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_kubernetes_p
od_name
        expr: sum by (kubernetes_namespace, kubernetes_pod_name)
(ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_respon
se_code_class
        expr: sum by (kubernetes_namespace, ocscp_response_code_class)
(ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_type
e
        expr: sum by (kubernetes_namespace, ocscp_nf_type)
(ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_respon
se_code_class_ocscp_producer_host
        expr: sum by (kubernetes_namespace, ocscp_response_code_class,
ocscp_producer_host) (ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_rx_res_5xx_total:aggregated_by_kubernetes_namespace_ocscp_re
sponse_code
        expr: sum by (kubernetes_namespace, ocscp_response_code)
(ocscp_metric_http_rx_res_total{ocscp_response_code_class="5"})
    - record:

```

```

ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_respon
se_code_ocscp_nf_end_point
    expr: sum by (kubernetes_namespace, ocscp_response_code,
ocscp_nf_end_point) (ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_ser
vice_type_ocscp_response_code_kubernetes_pod_name
    expr: sum by (kubernetes_namespace, ocscp_nf_service_type,
ocscp_response_code, kubernetes_pod_name) (ocscp_metric_http_rx_res_total)
    - record:
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace
    expr: sum by (kubernetes_namespace)
(ocscp_metric_http_tx_res_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
    - record:
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_respon
se_code_class
    expr: sum by (kubernetes_namespace, ocscp_response_code_class)
(ocscp_metric_http_tx_res_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
    - record:
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_typ
e_ocscp_nf_service_type
    expr: sum by (kubernetes_namespace, ocscp_nf_type,
ocscp_nf_service_type) (ocscp_metric_http_tx_res_total{ocscp_consumer_host!
="OCSCP-INITIATED-MSG"})
    - record:
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_respon
se_code_ocscp_nf_type
    expr: sum by (kubernetes_namespace, ocscp_response_code,
ocscp_nf_type) (ocscp_metric_http_tx_res_total{ocscp_consumer_host!="OCSCP-
INITIATED-MSG"})
    - record:
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_respon
se_code
    expr: sum by (kubernetes_namespace, ocscp_response_code)
(ocscp_metric_http_tx_res_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})

```

Sample Grafana Migration Patterns

Replace raw aggregations with the matching recorded series:

Table 2-15 Sample Grafana Migration Patterns

Pattern	Before	After
Total request rate	sum(irate(ocscp_metric_http_rx_req_to tal{...}[2m]))	irate(ocscp_metric_http_rx_req_total: aggregated_by_kubernetes_namespace{ku bernetes_namespace="\$namespace"}[2m])
Grouped by NF type	sum(irate(...[2m])) by (ocscp_nf_type)	irate(ocscp_metric_http_rx_req_total: aggregated_by_kubernetes_namespace_oc scp_nf_type{kubernetes_namespace="\$na mespace"}[2m])
Grouped by pod	sum(irate(...[2m])) by (kubernetes_pod_name)	irate(ocscp_metric_http_rx_req_total: aggregated_by_kubernetes_namespace_ku bernetes_pod_name{kubernetes_namespac e="\$namespace"}[2m])

Table 2-15 (Cont.) Sample Grafana Migration Patterns

Pattern	Before	After
Response code class	sum(irate(ocscp_metric_http_rx_res_total{...}[2m])) by (ocscp_response_code_class)	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class{kubernetes_namespace="\$namespace"}[2m])
5xx counter	sum(ocscp_metric_http_rx_res_total{...,ocscp_response_code_class="5"}) by (ocscp_response_code)	ocscp_metric_http_rx_res_5xx_total:aggregated_by_kubernetes_namespace_ocscp_response_code{kubernetes_namespace="\$namespace"}

Example success-rate calculation:

```
(sum(irate(ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class{kubernetes_namespace="$namespace",
ocscp_response_code_class="2"}[2m])) /
sum(irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="$namespace"}[2m]))) * 100
```

When to Keep an Outer sum

Keep an outer sum only when a panel intentionally collapses labels retained by a rule, such as summing a pod-and-container recording rule across pods and containers.

```
sum(irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_kubernetes_pod_name_container{kubernetes_namespace="$namespace",
container=~"scp-worker"}[2m]))
```

Do not retain sum by (...) when the recording rule already has exactly the labels required by the panel.

Recording Rule to Grafana Expression Mapping

The following table describes rule name mapping with the corresponding recorded-series Grafana query. Each replaces the raw metric aggregation with `irate(...[2m])`, retaining the original panel filters.

Table 2-16 Recording Rule to Grafana Expression Mapping

Recording Rule	Impacted Grafana Expression (Old)	New Grafana Expression
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace	sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m]))	irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type	sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_nf_type)	irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type{kubernetes_namespace="\$namespace"}[2m])

Table 2-16 (Cont.) Recording Rule to Grafana Expression Mapping

Recording Rule	Impacted Grafana Expression (Old)	New Grafana Expression
ocscp_metric_http_rx_req_to tal:aggregated_by_kubernete s_namespace_kubernetes_pod_name	sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (kubernetes_pod_name)	irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_kubernetes_pod_name{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_nf_type_ocscp_req_msg_size	sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_nf_type, ocscp_req_msg_size)	irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type_ocscp_req_msg_size{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_authority	sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_authority)	irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace_ocscp_authority{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_req_to tal:aggregated_by_kubernete s_namespace	sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m]))	irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_nf_type	sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_nf_type)	irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_nf_type_ocscp_nf_end_point	sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_nf_type, ocscp_nf_end_point)	irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_type_ocscp_nf_end_point{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_producer_host	sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_producer_host)	irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_producer_host{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_nf_end_point	sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_nf_end_point)	irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_end_point{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_req_to tal:aggregated_by_kubernete s_namespace_ocscp_nf_service_type	sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_nf_service_type)	irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace_ocscp_nf_service_type{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_res_to tal:aggregated_by_kubernete s_namespace	sum(irate(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace"}[2m]))	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="\$namespace"}[2m])

Table 2-16 (Cont.) Recording Rule to Grafana Expression Mapping

Recording Rule	Impacted Grafana Expression (Old)	New Grafana Expression
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class	sum(irate(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_response_code_class)	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_type	sum(irate(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_nf_type)	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_type{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class_ocscp_producer_host	sum(irate(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_response_code_class, ocscp_producer_host)	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class_ocscp_producer_host{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_rx_res_5xx_total:aggregated_by_kubernetes_namespace_ocscp_response_code	sum(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace", ocscp_response_code_class="5"}) by (ocscp_response_code)	ocscp_metric_http_rx_res_5xx_total:aggregated_by_kubernetes_namespace_ocscp_response_code{kubernetes_namespace="\$namespace"}
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_ocscp_nf_end_point	sum(irate(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace"}[2m])) by (ocscp_response_code, ocscp_nf_end_point)	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_ocscp_nf_end_point{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace	sum(irate(ocscp_metric_http_tx_res_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m]))	irate(ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_type_ocscp_nf_service_type	sum(irate(ocscp_metric_http_tx_res_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_nf_type, ocscp_nf_service_type)	irate(ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_type_ocscp_nf_service_type{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class	sum(irate(ocscp_metric_http_tx_res_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_response_code_class)	irate(ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_ocscp_nf_type	sum(irate(ocscp_metric_http_tx_res_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_response_code, ocscp_nf_type)	irate(ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_ocscp_nf_type{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code	sum(irate(ocscp_metric_http_tx_res_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (ocscp_response_code)	irate(ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code{kubernetes_namespace="\$namespace"}[2m])

Table 2-16 (Cont.) Recording Rule to Grafana Expression Mapping

Recording Rule	Impacted Grafana Expression (Old)	New Grafana Expression
ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_service_type_ocscp_response_code_kubernetes_pod_name	sum(irate(ocscp_metric_http_rx_res_total{kubernetes_namespace="\$namespace", ocscp_nf_service_type="nnrf-oauth2", ocscp_response_code="200"}[2m])) by (kubernetes_pod_name)	irate(ocscp_metric_http_rx_res_total:aggregated_by_kubernetes_namespace_ocscp_nf_service_type_ocscp_response_code_kubernetes_pod_name{kubernetes_namespace="\$namespace", ocscp_nf_service_type="nnrf-oauth2", ocscp_response_code="200"}[2m])
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace + ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace	sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) + sum(irate(ocscp_metric_http_tx_req_total{kubernetes_namespace="\$namespace"}[2m]))	irate(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="\$namespace"}[2m]) + irate(ocscp_metric_http_tx_req_total:aggregated_by_kubernetes_namespace{kubernetes_namespace="\$namespace"}[2m])
ocscp_metric_http_tx_res_total:aggregated_by_kubernetes_namespace_ocscp_response_code_class / ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace	((sum(irate(ocscp_metric_http_tx_res_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG", ocscp_response_code_class="2"}[2m])) / (sum(irate(ocscp_metric_http_rx_req_total{kubernetes_namespace="\$namespace", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])))) * 100	The pasted source ends mid-expression: ((sum(irate(ocscp_metric_http_tx_res_total:aggregated_by_kube

2.14.1 Applying Recording Rules

Perform the following procedure to apply and validate OC-SCP Prometheus recording rules in the Prometheus server ConfigMap.

Prerequisite:

Obtain access to the target Kubernetes cluster and namespace, and identify the Prometheus server ConfigMap used by vz-oso-prom-svr.

1. Create a merge-patch file named `ocscp_recordingrules.yaml`.

Match the existing ConfigMap data structure. If the ConfigMap also carries alert rules, define recording rules before alerting rules.

```
apiVersion: v1
data:
  recording_rules.yml: |
    groups:
      - name: SCPHttpRecordingRules
        interval: 1m
        rules:
          - record:
```

```
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace
  expr: sum by (kubernetes_namespace)
(ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-MSG"})
```

2. Identify the Prometheus server ConfigMap and back it up.

```
kubectl get configmap -n <namespace> | grep oso-prom-svr
kubectl get configmap <configmap-name> -n <namespace> -o yaml > configmap-backup.yaml
```

3. Apply the recording-rule merge patch.

```
kubectl patch configmap <cm-name> -n <namespace> --type=merge --patch-file
ocscp_recordingrules.yaml
```

4. Verify that the ConfigMap contains the SCPHttpRecordingRules group.

```
kubectl get configmap <cm-name> -n <namespace> -o yaml
kubectl get configmap <cm-name> -n <namespace> -o yaml | grep -A20
SCPHttpRecordingRules
```

5. Confirm that Prometheus loaded the recording rules.

In Prometheus, select **Status**, and then **Rules**. Confirm that SCPHttpRecordingRules is visible and has no evaluation errors.

6. Confirm that the recorded series is present and fresh.

Query the recorded metric and then query its sample age.

```
ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace

time() -
timestamp(ocscp_metric_http_rx_req_total:aggregated_by_kubernetes_namespace
)
```

7. Order rule groups so recording rules precede dependent alert rules.

This order enables alerts to use already evaluated recorded metrics. It does not resolve stale rules caused by Prometheus overload or rule-evaluation failures.

```
groups:
- name: SCPHttpRecordingRules
  interval: 1m
  rules:
  - record: ...

- name: SCPAlerts
  interval: 1m
  rules:
  - alert: ...
```

8. Complete the post-application validation.

Verify the rule group in Prometheus by navigating to **Status** and **Rules**, then validate the recorded metric output, freshness, and source aggregation.

```
sum by (kubernetes_namespace) (  
  ocscp_metric_http_rx_req_total{ocscp_consumer_host!="OCSCP-INITIATED-  
MSG"}  
)
```

The Prometheus ConfigMap contains the OC-SCP recording-rule group, Prometheus has loaded it without evaluation errors, and the recorded series returns fresh samples that agree with the source aggregation.

2.14.2 Troubleshooting Stale Recording Rules and Missing Fresh Samples

Use Prometheus freshness and rule-evaluation metrics to identify stale OC-SCP recording rules, then reduce or correct the condition causing Prometheus overload.

A recorded metric can exist while `rate()` or `irate()` returns no data. In this case, determine whether the recorded series is stale or lacks fresh samples.

Check the age of the most recent recorded sample:

```
time() - timestamp(<recording_rule_metric>)
```

Inspect rule-evaluation failures, missed group iterations, execution duration, and configured interval:

```
increase(prometheus_rule_evaluation_failures_total[30m])  
increase(prometheus_rule_group_iterations_missed_total[30m])  
prometheus_rule_group_last_duration_seconds  
prometheus_rule_group_interval_seconds
```

If the source expression returns data but the recorded metric is stale, Prometheus rule evaluation is unhealthy or overloaded.

Recommendations

- Keep recording rules in a dedicated rule group before alerting-rule groups.
- Start with a 1m evaluation interval. If Prometheus is overloaded, increase the interval to 5m and change the related Grafana expression to use the same 5m window.
- Correct the underlying cause of Prometheus overload. For example, investigate and remediate an alert storm so that excessive alert processing does not prevent recording-rule evaluation.

3

Alerts

This section provides information about the supported alerts and how to configure the alerts.

Note

The performance and capacity of the SCP system may vary based on the call model, feature or interface configuration, network conditions, and underlying CNE and hardware environment.

You can configure alerts in Prometheus and ScpAlertrules.yaml file.

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

Table 3-1 Alerts Levels or Severity Types

Alerts 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 may affect the service of SCP.
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 may affect the service of SCP.
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 may affect the service of SCP.
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 SCP.

Caution

User, computer and applications, and character encoding settings may cause an issue when copy-pasting commands or any content from PDF. The PDF reader version also affects the copy-pasting functionality. It is recommended to verify the pasted content when the hyphens or any special characters are part of the copied content.

Note

- `kubectl` commands might vary based on the platform deployment. Replace `kubectl` with Kubernetes environment-specific command line tool to configure Kubernetes resources through kube-api server. The instructions provided in this document are as per the Oracle Communications Cloud Native Environment (OCCNE) version of kube-api server.
- The alert file can be customized as required by the deployment environment. For example, namespace can be added as a filtered criteria to the alert expression to filter alerts only for a specific namespace.

3.1 Configuring Alerts

This section describes how to configure alert rules for SCP in both Prometheus and OCI environments. It provides guidance on setting up measurement-based alert rules, where the alerting system evaluates metrics reported by SCP microservices against specified rule conditions to generate alerts as required.

Prometheus Alert Configuration

In a Prometheus environment, SCP alert rules are configured based on metrics reported by SCP components. The alerting workflow monitors these metrics and issues notifications when the defined conditions are met.

For more information about configuring SCP alerts in Prometheus, see the "Alert Configuration" section in *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*.

OCI Alert Configuration

In OCI environments, SCP alert rules are defined using Metric Query Language (MQL). Configuring these alerts requires creating a dedicated SCP alert file for use with the OCI observability platform.

For more information about configuring SCP alerts in OCI, see the "Configuring SCP Alerts for OCI" section in *Oracle Communications Cloud Native Core, Service Communication Proxy Installation, Upgrade, and Fault Recovery Guide*.

3.2 System level alerts

This section lists the system level alerts.

3.2.1 SCPNotificationPodMemoryUsage

Table 3-2 SCPNotificationPodMemoryUsage

Field	Description
Description	Notify Notification service Pod memory usage if it is above threshold Threshold value is 85% of allocated (4GB) memory: 3.4 GB

Table 3-2 (Cont.) SCPNotificationPodMemoryUsage

Field	Description
Summary	Memory usage is above 85% for podname: {{labels.kubernetes_pod_name}}, instance: {{labels.instance}}, namespace: {{labels.kubernetes_namespace}} with current value {{ \$value }},scpfqdn: {{labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Major
Conditions	sum(container_memory_usage_bytes{image!="",pod=~".*scpc-notification.+"}) by (kubernetes_pod_name,kubernetes_namespace, instance) > 3650722201
OID	1.3.6.1.4.1.323.5.3.35.1.2.3001
Metric Used	ocscp_nrf_notifications_requests_nf_total
Recommended Actions	Cause: When high notification rate or very large NF profile size is present in notifications. Diagnostic Information: Monitor the notification metric: ocscp_nrf_notifications_requests_nf_total. Notification usage reduces after some time when it crosses 2.5 GB or 3 GB. Recovery: This alert is cleared automatically when the scpc-notification pod memory usage reduces below the defined threshold. Reduce the notification rate. These notifications are generated by NRF and can be controlled through NRF. For any assistance, contact My Oracle Support.

3.2.2 SCPWorkerPodMemoryUsage

Table 3-3 SCPWorkerPodMemoryUsage

Field	Description
Description	Notify Worker per Pod memory usage is above threshold Threshold value is 85% of allocated (16GB) memory: 13.6 GB
Summary	Memory usage is above 85% for podname: {{labels.kubernetes_pod_name}}, instance: {{labels.instance}}, namespace: {{labels.kubernetes_namespace}} with current value {{ \$value }},scpfqdn: {{labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	major
Conditions	sum(container_memory_usage_bytes{image!="",pod=~".*scp-worker.+"}) by (kubernetes_pod_name,kubernetes_namespace, instance) > 14602888806
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.7004
Metric Used	- ocscp_metric_http_rx_req_total - ocscp_metric_http_tx_req_total - ocscp_metric_http_rx_res_total - ocscp_metric_http_tx_res_total

Table 3-3 (Cont.) SCPWorkerPodMemoryUsage

Field	Description
Recommended Actions	Cause: When there is high traffic rate, alternate routing, more number of routing rules and rules size, and due to network or producer NF latency. Diagnostic Information: Monitor traffic rate, alerts, and latency on the KPI Dashboard. Check the traffic rates of the following metrics if they are too high: • ocscp_metric_http_rx_req_total • ocscp_metric_http_tx_req_total • ocscp_metric_http_rx_res_total • ocscp_metric_http_tx_res_total Check the upstream response time by using the following command and ensure whether upstream is taking too long to respond: ocscp_metric_upstream_service_time_total. Check the following platform metric for current memory usage by the scp-worker pod: container_memory_usage_bytes. Recovery: This alert is cleared automatically when the scp-worker pod memory usage reduces below the defined threshold. Reduce the traffic rate and improve the latency. For any assistance, contact My Oracle Support.

3.2.3 SCPIInstanceDown

Table 3-4 SCPIInstanceDown

Field	Description
Description	Notify that if any pod in ocscp release is down. Provides information like pod name, instance id and app name.
Summary	Pod with podname: {{\$labels.kubernetes_pod_name}} is Down , timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Critical
Conditions	kube_pod_status_ready{pod =~ '*scp-worker.*'.*scpc-notification.*'.*scpc-subscription.*'.*scpc-configuration.*'.*scpc-audit.*'.*scpc-nrfproxy.*'.*scpc-load-manager.*'.*scpc-cache.*'.*scpc-alternate-resolution.*'.*scpc-mediation.*',condition =~ 'true'} !=1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.7006

Table 3-4 (Cont.) SCPIInstanceDown

Field	Description
Recommended Actions	Cause: When the following issues occur: - The Control plane pods, such as configuration, subscription, notification, audit, and alternate-resolution are down due to connection failure with DB. - Pod restarts due to kubernetes liveness or readiness probe failures. - Application restarts or starts failure. Diagnostic Information: - Check if DB services are active by running the following command: <pre>kubectl describe pod <podname> -n <namespace></pre> - Check kubernetes events for probe failures in the platform logs. - Check if any exception is reported in the SCP application logs. Recovery: This alert is cleared automatically when the inactive pod is active. Recover DB services if down. Collect the application logs and contact My Oracle Support for any assistance.

3.3 Application level alerts

This section lists the application level alerts.

3.3.1 SCPCcaFeatureEnabledWithoutHttps

Table 3-5 SCPCcaFeatureEnabledWithoutHttps

Field	Description
Severity	Info
Condition	ocscp_worker_cca_validation_feature_enabled_without_https > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.9022
Description	An alert is raised when the CCA validation feature is enabled without enabling HTTPS.
Recommended Actions	Cause: CCA validation feature is enabled without enabling HTTPS. Diagnostic Information: Deploy HTTPS SCP deployment. Recovery: The alert is cleared automatically if either the CCA feature is disabled or deployment is changed to HTTPS. For any assistance, contact My Oracle Support.

3.3.2 SCPIngressTrafficRateAboveMinorThreshold

Table 3-6 SCPIngressTrafficRateAboveMinorThreshold

Field	Details
Description	'Ingress Traffic Rate for ocscp_authority: "{{labels.ocscp_authority}}" is above minor threshold (i.e. 5600 mps)'
Summary	'namespace: {{labels.namespace}}, podname: {{labels.pod}}, scpfqdn: {{labels.ocscp_fqdn}}, scpauthority: {{labels.ocscp_authority}}, timestamp: {{ with query "time()" }} {{ . first value humanizeTimestamp }}{{ end }}: Current Ingress Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 70 Percent of Max Ingress MPS(8000)'
Severity	minor
Condition	This alert notifies that the traffic rate is above 1200 MPS (user configure minor threshold value) with Locality and current value of traffic rate. sum(rate(ocscp_metric_http_rx_req_total{container="scp-worker"}[2m])) by (namespace,pod,ocscp_fqdn,ocscp_authority) >= 5600 < 6400
OID	1.3.6.1.4.1.323.5.3.35.1.2.7001
Metric Used	ocscp_metric_http_rx_req_total
Recommended Action	Cause: When a consumer NF sends more traffic than expected. Diagnostic Information: - Monitor ingress traffic to the pod using the KPI dashboard. - Refer to the rate of ocscp_metric_total_http_rx_req on the Grafana dashboard. Recovery: - The alert clears automatically when ingress traffic falls below the minor threshold or exceeds the major threshold. - If the alert does not clear, check the Consumer NF for uneven traffic distribution per connection or other abnormalities. - For any assistance, contact My Oracle Support.

Note

The alert expression is configured for the SCP profile (12 vCPUs and 16 Gi of memory). For a smaller resource profile (for example, 8 vCPUs and 12 Gi of memory), the per-worker pod MPS should be evaluated based on the dimensioning sheet, and the alert file must be updated accordingly.

3.3.3 SCPIngressTrafficRateAboveMajorThreshold

Table 3-7 SCPIngressTrafficRateAboveMajorThreshold

Field	Details
Description	'Ingress Traffic Rate for ocscp_authority: <code>{{ \$labels.ocscp_authority }}</code> is above major threshold (i.e. 6400 mps)'
Summary	'namespace: <code>{{ \$labels.namespace }}</code> , podname: <code>{{ \$labels.pod }}</code> , scpfqdn: <code>{{ \$labels.ocscp_fqdn }}</code> , scpauthority: <code>{{ \$labels.ocscp_authority }}</code> , timestamp: <code>{{ with query "time()" }} {{ . first value humanizeTimestamp }} {{ end }}</code> : Current Ingress Traffic Rate is <code>{{ \$value printf "%.2f" }}</code> mps which is above 80 Percent of Max Ingress MPS(8000)'
Severity	major
Condition	This alert notifies that the traffic rate is above 1400 MPS (user configure major threshold value) with Locality and current value of traffic rate. expr: <code>sum(rate(ocscp_metric_http_rx_req_total{container="scp-worker"}[2m])) by (namespace,pod,ocscp_fqdn,ocscp_authority) >= 6400 < 7600</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7001
Metric Used	ocscp_metric_http_rx_req_total
Recommended Action	Cause: When a consumer NF sends more traffic than expected. Diagnostic Information: - Monitor ingress traffic to the pod using the KPI dashboard. - Refer to the rate of ocscp_metric_total_http_rx_req on the Grafana dashboard. Recovery: - The alert clears automatically when ingress traffic falls below the major threshold or exceeds the critical threshold. - If the alert does not clear, check the Consumer NF for uneven traffic distribution per connection or other abnormalities. - If the alert persists for an extended duration, lower ingress traffic from the Consumer NF to the pod. - For any assistance, contact My Oracle Support.

Note

The alert expression is configured for the SCP profile (12 vCPUs and 16 Gi of memory). For a smaller resource profile (for example, 8 vCPUs and 12 Gi of memory), the per-worker pod MPS should be evaluated based on the dimensioning sheet, and the alert file must be updated accordingly.

3.3.4 SCPIngressTrafficRateAboveCriticalThreshold

Table 3-8 SCPIngressTrafficRateAboveCriticalThreshold

Field	Details
Description	'Ingress Traffic Rate for ocscp_authority: {{ \$labels.ocscp_authority }} is above critical threshold (i.e. 7600 mps)'
Summary	'namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfqdn: {{ \$labels.ocscp_fqdn }}, scpauthority: {{ \$labels.ocscp_authority }}, timestamp: {{ with query "time()" }} {{ . first value humanizeTimestamp }}{{ end }}: Current Ingress Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 95 Percent of Max Ingress MPS(8000)'
Severity	Critical
Condition	Notify that Traffic rate is above 1600mps (user configure critical threshold value) with Locality and current value of traffic rate. sum(rate(ocscp_metric_http_rx_req_total{container="scp-worker"}[2m])) by (namespace,pod,ocscp_fqdn,ocscp_authority) >= 7600
OID	1.3.6.1.4.1.323.5.3.35.1.2.7001
Metric Used	ocscp_metric_http_rx_req_total
Recommended Action	Cause: A Consumer NF sends more traffic than expected. Diagnostic Information: - Monitor ingress traffic to the pod using the KPI dashboard. - Refer to the rate of ocscp_metric_total_http_rx_req on the Grafana dashboard. Recovery: - The alert clears automatically when ingress traffic falls below the critical threshold. - If the alert does not clear, check the Consumer NF for uneven traffic distribution per connection or other abnormalities. - If the alert persists for an extended duration, lower ingress traffic from the Consumer NF to the pod. - For any assistance, contact My Oracle Support.

Note

The alert expression is configured for the SCP profile (12 vCPUs and 16 Gi of memory). For a smaller resource profile (for example, 8 vCPUs and 12 Gi of memory), the per-worker pod MPS should be evaluated based on the dimensioning sheet, and the alert file must be updated accordingly.

3.3.5 SCPRoutingFailedForProducer

Table 3-9 SCPRoutingFailedForProducer

Field	Details
Description	'Routing failed for nfclustername: {{ \$labels.ocscp_producer_host }}'
Summary	'Routing failed for producer with nfclustername: {{ \$labels.ocscp_producer_host }}, nfservicetype = {{ \$labels.ocscp_nf_service_type }}, scpauthority: {{ \$labels.ocscp_authority }}, nftype = {{ \$labels.ocscp_nf_type }}, nfendpoint = {{ \$labels.ocscp_nf_end_point }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scp_fqdn: ' {{ \$labels.scp_fqdn }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }} and value = {{ \$value }} '
Severity	info
Condition	This alert notifies that routing failed for producer. It provides information about NFService Type, NFType, Locality, producer fqdn, and value. increase(ocscp_metric_routing_attempt_fail_total{app_kubernetes_io_name="scp-worker"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.7005
Metric Used	ocscp_metric_routing_attempt_fail_total
Recommended Action	Cause: When routing fails to select a producer NF because routing rules do not exist for the NF service or producer. Diagnostic Information: - Verify that routing rules are configured for the NF for which routing failed. - Check notification logs for errors that occur when processing notifications for the NF for which routing failed. Use the following command: <pre>kubectl logs <podname> -n <namespace></pre> - Verify that the NF is reachable by using one of the following methods: - Run the ping command from cluster nodes by using the Service IP address. For example, run ping <IPAddress>. - If the Service Fqdn is used, run the ping command from within the pod. - If the pod does not support the ping command, use the SCP pod debug container. - As an alternative to ping, collect tcpdump data to verify connectivity. For example, run tcpdump -w capture.pcap -i <pod interface>. - Refer to the Debug Container Guide for usage and sample commands. Recovery: - The alert clears automatically when routing succeeds for the producer NF or when no further traffic arrives during the next Prometheus scrape interval. - Verify that the NF is registered. Register the NF to create routing rules if the rules do not exist. - For any assistance, contact My Oracle Support.

3.3.6 SCPAuditErrorResponse

Table 3-10 SCPAuditErrorResponse

Field	Description
Severity	Info
Conditions	ocscp_audit_error_response > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.4001
Description	Alert is raised when Audit module receives a 3xx, 4xx, or 5xx error from NRF. This alert is labeled with specific nftype, nrfRegionOrSetId, and auditmethod. Note: Alert is cleared on the next audit cycle.
Recommended Actions	Cause: When the configured NRF sends error responses, down, or not reachable. Diagnostic Information: - Check if NRF is up and reachable. To check the NRF status, see <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. - Check if the NF is reachable or not using one of the following steps: - Run the ping command from primary/secondary nodes using IP of NRF. Example of a ping command: <code>ping <IpAddress></code>. - Run the ping command from inside the pod if FQDN of NRF is used. - Instead of using ping, you can collect tcpdump for ensuring connectivity. Example of a tcpdump: <code>tcpdump -w capture.pcap -i <pod interface></code>. - Monitor audit and worker service logs: <code>kubectl logs <pod name> -n <namespace></code> - Check jaeger traces for scp-worker. Recovery: The alert is cleared automatically during the next audit cycle and when no more errors are received. Collect audit and worker service logs and contact My Oracle Support for any assistance.

3.3.7 SCPAuditEmptyNFArrayResponse

Table 3-11 SCPAuditEmptyNFArrayResponse

Field	Description
Description	Alert is generated when Audit module receives a 2xx response with empty NFInstance array from NRF. Alert is labeled with specific nftype, nrfRegionOrSetId, and auditmethod. Alert is cleared if Audit receives a success response with non-empty NFInstance array or on next audit cycle if topology source is changed to LOCAL.
Summary	SCP Audit received Empty NF Array Response for nftype {{ \$labels.nftype }}, nrfRegionOrSetId: {{ \$labels.nrfRegionOrSetId }}, auditmethod: {{ \$labels.auditmethod }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scpfdqn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }}

Table 3-11 (Cont.) SCPAuditEmptyNFArrayResponse

Field	Description
Severity	Critical
Conditions	ocscp_audit_2xx_empty_nf_array > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.4002
Recommended Actions	Cause: When NRF does not have any NF registered or due to any error condition on NRF. Diagnostic Information: Check if NRF contains any registered NF and validate as required. For more information, refer to NRF documents. Recovery: This alert is cleared automatically if Audit receives a success response with non-empty NFInstance array or during the next audit cycle when the topology source is changed to LOCAL. Register a NF with NRF or change the topology source to LOCAL. For any assistance, contact My Oracle Support.

3.3.8 DuplicateLocalityFoundInForeignNF

Table 3-12 DuplicateLocalityFoundInForeignNF

Field	Details
Description	'Duplicate Locality found for Foreign NF'
Summary	'Duplicate locality for remote Region. namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scpFqdn: ' {{ \$labels.scpFqdn }} ', timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }} and value = {{ \$value }} '
Severity	Major
Condition	This alert is raised when a foreign SCP is registered with a duplicate locality from the current region.
OID	1.3.6.1.4.1.323.5.3.35.1.2.3002
Metric Used	ocscp_notification_duplicate_foreign_location

Table 3-12 (Cont.) DuplicateLocalityFoundInForeignNF

Field	Details
Recommended Action	Cause: When a foreign SCP profile is accepted, but its MateScpInfo localities overlap with the local SCP's serving or mate localities. Diagnostic Information: Check if ocscp_notification_duplicate_foreign_location = 1. Group by NF type: sum by (nfType) (ocscp_notification_duplicate_foreign_location) Determine a specific instance: ocscp_notification_duplicate_foreign_location{nfInstanceId="<id>"} Look for WARN: "Common Serving or Mate localities found between local SCP and foreign SCP" Look for pegging from pegInterScpMetrics in ForeignSCPRuleProcessorImp Verify the row for nfInstanceId has duplicateRemoteLocality=true (and the NF is foreign) Recovery: Ensure that the foreign SCP's serving and MateScpInfo localities do not overlap with the local SCP's serving or mate localities. Update the foreign SCP's profile and re-register it. After the profile is corrected and processed, the gauge automatically drops to 0 on the next REGISTERED event with no overlap (or clears on DEREGISTERED). For any assistance, contact My Oracle Support.

3.3.9 ForeignNFLocalityNotServed

Table 3-13 ForeignNFLocalityNotServed

Field	Description
Severity	Critical
Conditions	ocscp_notification_foreign_nf_locality_unserved > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.3003
Description	Alert is raised when a Foreign Producer NF's locality is not served by any SCP.

Table 3-13 (Cont.) ForeignNFLocalityNotServed

Field	Description
Recommended Actions	Cause: When SCP discovers an unknown Producer NF's without any locality served by an SCP. Diagnostic Information: Check logs for received NF notification by running the following command: <code>kubectl get pods -n <namespace></code>. Note: Use the complete name of notification pod in the following command: <code>kubectl logs <pod> -n <namespace></code>. Check the following metric to get the NFInstanceId information for which this alert is raised: <code>ocscp_notification_foreign_nf_locality_unserved(nfInstanceId)</code>. Recovery: This alert is cleared automatically if the unknown NF is deregistered or registers update received with locality served by SCP. Re-register NF with correct locality information. For any assistance, contact My Oracle Support.

3.3.10 UnknownLocalityFoundInForeignNF

Table 3-14 UnknownLocalityFoundInForeignNF

Field	Description
Severity	critical
Conditions	<code>ocscp_notification_foreign_nf_locality_absent > 0</code>
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.3004
Description	Alert will be raised when a Foreign Producer NF's locality is unknown.
Recommended Actions	Cause: When SCP discovers an unknown Producer NF's without locality information. Diagnostic Information: Check logs for the received NF notification by running the following command: <code>kubectl get pods -n <namespace></code>. Use the complete name of notification pod in the following command: <code>kubectl logs <pod> -n <namespace> -f --tail=0</code> Check the following metric to get the NFInstanceId information for which this alert is raised: <code>ocscp_notification_foreign_nf_locality_absent(nfInstanceId)</code>. Recovery: This alert is cleared automatically if unknown NF is deregistered or registers update received with locality known to SCP. Re-register NF with correct locality information. For any assistance, contact My Oracle Support.

3.3.11 SCPUpsstreamResponseTimeout

Table 3-15 SCPUpsstreamResponseTimeout

Field	Description
Description	Alert is raised when Upstream connection to a producer NF fails

Table 3-15 (Cont.) SCPUstreamResponseTimeout

Field	Description
Summary	SCP SUstream Response Timeout for nfservicetype: <code>{{labels.ocscp_nf_service_type}}</code> , nftype <code>{{labels.ocscp_nf_type}}</code> , responsecode <code>{{labels.ocscp_response_code}}</code> , nfclustername <code>{{labels.ocscp_producer_host}}</code> , namespace: <code>{{labels.kubernetes_namespace}}</code> , podname: <code>{{labels.kubernetes_pod_name}}</code> , timestamp: <code>{{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}</code>
Severity	info
Conditions	<code>increase(ocscp_metric_upstream_response_timeout_total[2m]) > 0</code>
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.7011
Recommended Actions	Cause: When a producer NF is down, not reachable, or latency is high. Diagnostic Information: Check whether the producer NF is up and network connectivity to the producer NF is established by using one of the following steps: - Run the ping command from primary/secondary nodes by using IP of producer NF. Example of a ping command: <code>ping <IPAddress></code>. - Run the ping command from inside the pod, if FQDN of producer NF is used. - Instead of using the ping command, you can collect tcpdump for ensuring the connectivity. tcpdump must be run on the debug container for scp-worker microservice. Example of a tcpdump: <code>tcpdump -w capture.pcap -i <pod interface></code>. Check the upstream response time by using the following metric and determine if upstream is taking too long to respond: <code>ocscp_metric_upstream_service_time</code> (producer FQDN) Recovery: This alert is cleared automatically in the next scrape interval if the system does not observe any error. For any assistance, contact My Oracle Support.

3.3.12 SCPSingleNfInstanceAvailableForNFType

Table 3-16 SCPSingleNfInstanceAvailableForNFType

Field	Description
Severity	Major
Conditions	<code>ocscp_no_nf_instance == 0</code>
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.3005
Description	Alert is raised when there is a single NFInstance available with SCP for an NFType.

Table 3-16 (Cont.) SCPSingleNFInstanceAvailableForNFType

Field	Description
Recommended Actions	Cause: When the preventiveAuditOnLastNFInstanceDeletion attribute is set to true, SCP has single NFInstance available for an NFType. Diagnostic Information: Check all SCP NRFs for specific NFType in the alert if only one NFInstance is available. For information about registered NFs, see <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. Check the number of NFs of a particular type by using API or CNC console of SCP. For information about procedures to check the NFs available with SCP, see "Configuring Service Communication Proxy using the CNC Console" in <i>Oracle Communications Cloud Native Core, Service Communication Proxy User Guide</i>. Recovery: This alert is cleared automatically in the next scrape interval if more than one NFInstance is available for a specified NFType in the alert. For any assistance, contact My Oracle Support.

3.3.13 SCPNoNFInstanceForNFType

Table 3-17 SCPNoNFInstanceForNFType

Field	Description
Severity	Critical
Conditions	ocscp_no_nf_instance == 1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.3006
Description	Alert is raised when there is no NFInstance available with SCP for an NFType
Recommended Actions	Cause: When the preventiveAuditOnLastNFInstanceDeletion flag is set to true, SCP has no NFInstance available for a NFType. Diagnostic Information: Check all SCP NRFs for specific NFType in the alert if no NFInstance is available. For information about registered NFs, see <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. Check the number of NFs of a particular type by using API or CNC console of SCP. For information about procedures to check the NFs available with SCP, see "Configuring Service Communication Proxy using the CNC Console" in <i>Oracle Communications Cloud Native Core, Service Communication Proxy User Guide</i>. Recovery: This alert is cleared automatically in the next scrape interval if at least one NFInstance is available for a specified NFType in the alert. For any assistance, contact My Oracle Support.

3.3.14 SCPIngressTrafficRateExceededConfiguredLimit

Table 3-18 SCPIngressTrafficRateExceededConfiguredLimit

Alert Parameters	Value
Description	Ingress traffic rate exceeds configured rate limit for consumer fqdn: <code>{{labels.ocscp_consumer_fqdn}}</code>
Summary	'Ingress traffic rate exceeds configured rate limit for consumer fqdn: ocscpconsumerfqdn = <code>{{labels.ocscp_consumer_host}}</code> , consumernfinstanceid = <code>{{labels.ocscp_consumer_nf_instance_id}}</code> , consumernftype = <code>{{labels.ocscp_consumer_nf_type}}</code> , configuredingressrate = <code>{{labels.ocscp_configured_ingress_rate}}</code> , namespace: <code>{{labels.kubernetes_namespace}}</code> , podname: <code>{{labels.kubernetes_pod_name}}</code> , scp_fqdn: ' <code>{{labels.scp_fqdn}}</code> ', timestamp: <code>{{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }}</code> and value = <code>{{ \$value }}</code> '
Severity	Critical
Condition	This alert is raised when the ingress traffic rate exceeds the configured rate for consumer FQDN. <code>increase(ocscp_metric_ingress_rate_limiting_throttle_req_total[2m]) > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7012
Metric Used	ocscp_metric_ingress_rate_limiting_throttle_req_total
Recommended Actions	Cause: When the ingress traffic rate exceeds the configured rate limit for the consumer FQDN. Diagnostic Information: - Check the ingress traffic rate from the consumer FQDN. - To check the ingress rate, refer to the following metrics: <code>cscp_metric_http_rx_req_total</code> - Check the ingress rate limit configuration as described in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. Recovery: This alert is cleared when no more requests get suppressed due to ingress rate limiting in the next scrape interval. For any assistance, contact My Oracle Support.

3.3.15 SCPIngressTrafficRoutedWithoutRateLimitTreatment

Table 3-19 SCPIngressTrafficRoutedWithoutRateLimitTreatment

Alert Parameters	Value
Description	Ingress traffic routed without rate limit treatment
Summary	'Ingress traffic routed without rate limit treatment: consumernftype = <code>{{labels.ocscp_consumer_nf_type}}</code> , consumernfinstanceid = <code>{{labels.ocscp_consumer_nf_instance_id}}</code> , consumerfqdn = <code>{{labels.ocscp_consumer_host}}</code> , cause = <code>{{labels.ocscp_cause}}</code> , namespace: <code>{{labels.kubernetes_namespace}}</code> , podname: <code>{{labels.kubernetes_pod_name}}</code> , scp_fqdn: ' <code>{{labels.scp_fqdn}}</code> ', timestamp: <code>{{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }}</code> and value = <code>{{ \$value }}</code> '
Severity	Major

Table 3-19 (Cont.) SCPIngressTrafficRoutedWithoutRateLimitTreatment

Alert Parameters	Value
Condition	This alert is raised when the ingress traffic routes without rate limiting treatment. increase(ocscp_metric_ingress_rate_limiting_not_applied_req_total[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.7013
Metric Used	ocscp_metric_ingress_rate_limiting_not_applied_req_total
Recommended Actions	Cause: When the ingress traffic routes without rate limiting treatment. Diagnostic Information: - Check the ingress rate limiting configurations for the untreated FQDNs that can be obtained from the following metric: ocscp_metric_ingress_rate_limiting_not_applied_req_total(ocscp_consumer_fqdn) - Check the ingress rate limit configuration as described in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. Recovery: This alert is cleared when no more requests get routed without ingress rate limiting treatment in the next scrape interval. For any assistance, contact My Oracle Support.

3.3.16 SCPEgressTrafficRateExceededConfiguredLimit

Table 3-20 SCPEgressTrafficRateExceededConfiguredLimit

Field	Details
Description	'Egress traffic rate exceeds configured rate limit'
Summary	'Egress traffic rate exceeds configured rate limit: producerntype = {{ \$labels.ocscp_nf_type }}, producerntfservicetype = {{ \$labels.ocscp_nf_service_type }}, producerfqdn = {{ \$labels.ocscp_producer_host }}, configuredegressrate = {{ \$labels.ocscp_configured_egress_rate }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scp_fqdn: ' {{ \$labels.scp_fqdn }}', timestamp: {{ with query "time()" }} {{ . first value humanizeTimestamp }} {{ end }} and value = {{ \$value }} '
Severity	critical
Condition	This alert is raised when the egress traffic rate exceeds the configured rate. increase(ocscp_metric_egress_rate_limiting_throttle_req_total{app_kubernetes_io_name="scp-worker"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.7014
Metric Used	ocscp_metric_egress_rate_limiting_throttle_req_total

Table 3-20 (Cont.) SCPEgressTrafficRateExceededConfiguredLimit

Field	Details
Recommended Action	Cause: When the egress traffic rate exceeds the configured rate. Diagnostic Information: - Check Egress Traffic Rate by using <code>ocscp_metric_total_http_tx_req</code>. - Check the egress rate limit configuration. - Refer to the SCP User Guide for steps to check the rate limit configuration. Recovery: - The alert clears automatically when no further requests are throttled because of egress rate limiting during the next scrape interval. - For any assistance, contact My Oracle Support.

3.3.17 SCPEgressTrafficRoutedWithoutRateLimitTreatment

Table 3-21 SCPEgressTrafficRoutedWithoutRateLimitTreatment

Field	Description
Description	Alert is raised when egress traffic routes without rate limiting.
Summary	Egress traffic routed without rate limiting: <code>producernfstype = {{ \$labels.ocscp_nf_type }}</code> , <code>producernfservicetype = {{ \$labels.ocscp_nf_service_type }}</code> , <code>producernfinstanceid = {{ \$labels.ocscp_nf_instance_id }}</code> , <code>producerfqdn = {{ \$labels.ocscp_producer_host }}</code> , <code>cause = {{ \$labels.ocscp_cause }}</code> , <code>namespace: {{ \$labels.kubernetes_namespace }}</code> , <code>podname: {{ \$labels.kubernetes_pod_name }}</code> , <code>scpfqdn: {{ \$labels.ocscp_fqdn }}</code> , <code>timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}</code> and <code>value = {{ \$value }}</code>
Severity	Major
Conditions	<code>increase(ocscp_metric_egress_rate_limiting_not_applied_req_total{app_kubernetes_io_name="scp-worker"}[2m]) > 0</code>
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.7015
Recommended Actions	Cause: When the egress traffic routes without rate limiting treatment. Diagnostic Information: Check the egress rate limiting configurations for the untreated producer FQDN. Obtain the producer FQDN by using the following metric: <code>ocscp_metric_egress_rate_limiting_not_applied_req_total(ocscp_producer_fqdn)</code> Check the egress rate limit configuration as described in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. Recovery: This alert is cleared when no more requests get routed without egress rate limiting treatment in the next scrape interval. For any assistance, contact My Oracle Support.

3.3.18 SCPNotificationRejectTopologySourceLocal

Table 3-22 SCPNotificationRejectTopologySourceLocal

Field	Description
Severity	Info
Conditions	increase(ocscp_notifications_rejected_topologysource_local_total[15m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.3007
Description	Alert is raised when SCP rejects a notification from NRF due to topology source set to LOCAL for NF Type.
Recommended Actions	Cause: When NF Topology Source Info is set to LOCAL. Diagnostic Information: Check the topology source information of an NF Type. For information about the topology source APIs, see <i>Oracle Communications Cloud Native Core, Service Communication Proxy User Guide</i>. Recovery: This alert is cleared automatically after 15 minutes when NF Topology Source Info is set to NRF from LOCAL. For any assistance, contact My Oracle Support.

3.3.19 SCPNotificationProcessingFailureForNF

Table 3-23 SCPNotificationProcessingFailureForNF

Field	Description
Description	Alerts is raised when Notification processing has failed on SCP.
Summary	SCP Notification Processing failure for nfInstanceId: {{ \$labels.nfInstanceId }}, nfType: {{ \$labels.nfType }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scpfqdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Major
Conditions	increase(ocscp_notification_nf_profile_processing_failure_total[15m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.3008

Table 3-23 (Cont.) SCPNotificationProcessingFailureForNF

Field	Description
Recommended Actions	Cause: When Notification processing has failed on SCP. Diagnostic Information: Check notification pod logs for any errors by running the following command: <pre>kubectl logs <notification pod name> -n <scp namespace></pre> To get the list of pods, run the following command: <pre>kubectl get pod -n <scp namespace></pre> Sample logs: <pre>{ "instant": { "epochSecond":1620272241,"nanoOfSecond":609935406}, "thread":"runQueueThreadPool1","level":"ERROR","loggerName":"com.oracle.cgbu.cne.scp.soothsayer.Process","message":"{logMsg=Notified profile IP endpoints already present in stored profile, logMsgCode=DUPLICATE_IPENDPOINT_OR_FQDN_FOUND_IN_STORED_PROFILE, rootCause=Notified Profile contains an ipEndPoint {\"ipv4Address\\\":\\\"10.75.203.74\\\", \"transport\\\":\\\"TCP\\\", \"port\\\":32673} which is already present in stored Profile with nfInstanceId 93ED74AA-A29C-4450-9D7A-9278CAF6266D for serviceInstanceId audmcl08nv08-udmueauthn-589d6d5bcc-4dslh}\", \"endOfBatch\":false, \"loggerFqcn\":\"org.apache.logging.log4j.spi.AbstractLogger\", \"threadId\":34, \"threadPriority\":5, \"messageTimestamp\":\"21-05-06 03:37:21.609+0000\", \"application\":\"ocscp-soothsayer\", \"microservice\":\"ocscp-scpc-notification\", \"engVersion\":\"26.1.202\", \"mktgVersion\":\"26.1.202.0.0\", \"vendor\":\"oracle\", \"namespace\":\"scpsvc\", \"node\":\"worker1\", \"pod\":\"ocscp-scpc-notification-547f699c96-m7nc8\", \"subsystem\":\"notification\", \"instanceType\":\"prod\", \"processId\":\"1\"} { \"instant\": { \"epochSecond\":1620272241,\"nanoOfSecond\":734873890}, \"thread\":\"runQueueThreadPool1\",\"level\":\"WARN\", \"loggerName\":\"com.oracle.cgbu.cne.scp.soothsayer.scheduler.RunQueueConsumer\", \"message\":\"{logMsg=Profile Processing failed, nfInstanceId=93ED74AA-A29C-4450-9D7A-9278CAF6266D}\", \"endOfBatch\":false, \"loggerFqcn\":\"org.apache.logging.log4j.spi.AbstractLogger\", \"threadId\":34, \"threadPriority\":5, \"messageTimestamp\":\"21-05-06 03:37:21.734+0000\", \"application\":\"ocscp-soothsayer\", \"microservice\":\"ocscp-scpc-</pre>

Table 3-23 (Cont.) SCPNotificationProcessingFailureForNF

Field	Description
	notification", "engVersion": "26.1.202", "mktgVersion": "26.1.202.0.0", "vendor": "oracle", "namespace": "scpsvc", "node": "worker1", "pod": "ocscp-scpc-notification-547f699c96-m7nc8", "subsystem": "notification", "instanceType": "prod", "processId": "1"}" For information about the topology source APIs, see <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. Recovery: This alert is cleared automatically after 15 minutes. For any assistance, contact My Oracle Support.

3.3.20 SCPSubscriptionFailureForNFType

Table 3-24 SCPSubscriptionFailureForNFType

Field	Description
Severity	Critical
Conditions	increase(ocscp_subscription_nf_failure_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.2001
Description	Alerts is raised when SCP subscription to NRF has failed. This alert is labeled with specific nfType, nrfRegionOrSetId, and auditMethod.

Table 3-24 (Cont.) SCPSubscriptionFailureForNFType

Field	Description
Recommended Actions	Cause: When the subscription fails for an NF Type with NRF. Diagnostic Information: Check whether NRF is up. To check NRF status, see <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. Check whether the NRF is reachable or not by using one of the following steps: - Run the ping command from primary or secondary nodes using IP of NRF. Example of a ping command: ping <IPAddress>. - Run the ping command from inside the pod if FQDN of NRF is used. - Instead of using ping, you can collect tcpdump for ensuring the connectivity. Example of a tcpdump: tcpdump -w capture.pcap -i <pod interface>. If NRF is up, check scp-worker logs to find any error response from NRF. If there are error responses, monitor NRF logs: kubectl logs <pod name> -n <namespace>. Sample logs: <pre>{ "instant": {"epochSecond":1620275506,"nanoOfSecond":134773910}, "thread":"pool-8- thread-1","level":"ERROR","loggerName":"com.oracle.cgbu.cne.scp.soothsayer.subscription.processor.SubscriptionDataConsumer","message":"{logMsg=Exception occurred while handling action for subscriber data, action=TO_BE_RENEWED, stackTrace=com.oracle.cgbu.cne.scp.soothsayer.subscription.operations.NrfSubscriptionClient.triggerPatchRequest(NrfSubscriptionClient.java:177)\ncom.oracle.cgbu.cne.scp.soothsayer.subscription.processor.SubscriptionDataConsumer.handleAction(SubscriptionsDataConsumer.java:322)\ncom.oracle.cgbu.cne.scp.soothsayer.subscription.processor.SubscriptionDataConsumer.consumeSQ(SubscriptionsDataConsumer.java:140)\ncom.oracle.cgbu.cne.scp.soothsayer.subscription.processor.SubscriptionDataConsumer.run(SubscriptionsDataConsumer.java:84)\njava.base/java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1130)\njava.base/java.util.concurrent.ThreadPoolExecutor\$Worker.run(ThreadPoolExecutor.java:630)\njava.base/java.lang.Thread.run(Thread.java:832)}", "endOfBatch":false, "loggerFqn":"org.apache.logging.log4j.spi.AbstractLogger", "threadId":33, "threadPriority":5, "messageTimestamp":"21-05-06 04:31:46.134+0000", "application":"ocscp-soothsayer", "microservice":"ocscp-scpc-subscription", "engVersion":"1.15.0", "mktgVersion":"1.15.0.0.0", "vendor":"oracle", "namespace":"scpsvc", "node":"primary", "pod":"ocscp-scpc-</pre>

Table 3-24 (Cont.) SCPSubscriptionFailureForNFType

Field	Description
	subscription-55cfb57cc6-2qp2g", "subsystem": "subscription", "instanceType": "prod", "processId": "1"} Recovery: This alert is cleared automatically when NRF is up and running or errors are corrected for received error responses. For any assistance, contact My Oracle Support.

3.3.21 SCPReSubscriptionFailureForNFType

Table 3-25 SCPReSubscriptionFailureForNFType

Field	Description
Severity	Critical
Conditions	increase(ocscp_patch_subscription_nf_failure_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.2002
Description	Alerts is raised when SCP re-subscription to NRF has failed. This alert is labeled with specific nftype, nrfRegionOrSetId, and auditmethod.
Recommended Actions	Cause: When the re-subscription fails for an NF Type with NRF. Diagnostic Information: Check whether NRF is up. To check the NRF status, see the <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. Check whether the NRF is reachable or not by using one of the following steps: - Run the ping command from primary/secondary nodes by using IP of NRF. Example of a ping command: ping <IPAddress>. - Run the ping command from inside the pod, if FQDN of NRF is used. - Instead of using ping, you can collect tcpdump for ensuring the connectivity. Example of a tcpdump: tcpdump -w capture.pcap -i <pod interface>. Check scp-worker logs to find any error response from NRF. If there are error responses, monitor NRF logs: kubectl logs <pod name> -n <namespace>. Recovery: This alert is cleared automatically when NRF is up and running or errors are corrected for received error responses. For any assistance, contact My Oracle Support.

3.3.22 SCPNrfRegistrationFailureForRegionOrSetId

Table 3-26 SCPNrfRegistrationFailureForRegionOrSetId

Field	Description
Description	'SCP Registration Failure for Region Or SetId: {{\$labels.nrfRegionOrSetId}}'

Table 3-26 (Cont.) SCPNrfRegistrationFailureForRegionOrSetId

Field	Description
Summary	'SCP Registration Failure for Region Or SetId: {{{labels.nrfRegionOrSetId}}}, namespace: {{{labels.kubernetes_namespace}}}, podname: {{{labels.kubernetes_pod_name}}}, scpFqdn: ' {{{labels.scpFqdn}}} ', timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Conditions	This alert is raised when SCP registration fails. This alert is labeled with specific nftype, nrfRegionOrSetId, and auditmethod. increase(ocscp_subscription_nrf_registration_failure_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.2003
Metric Used	ocscp_subscription_nrf_registration_failure_total
Recommended Actions	Cause: When the registration fails for an NF Type with NRF. Diagnostic Information: Check whether NRF is up. To check the NRF status, see the <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. Check whether the NRF is reachable or not by using one of the following steps: - Run the ping command from primary or secondary nodes by using IP of NRF. Example of a ping command: ping <IPAddress>. - Run the ping command from inside the pod if FQDN of NRF is used. - Instead of using ping, you can collect tcpdump for ensuring the connectivity. Example of a tcpdump: tcpdump -w capture.pcap -i <pod interface>. Check scp-worker logs to find any error response from NRF. If there are error responses, monitor NRF logs: kubectl logs <pod name> -n <namespace>. Sample logs: <pre>{ "instant": { "epochSecond":1620638888,"nanoOfSecond":78728229,"thread":"reg istration-0","level":"ERROR","loggerName":"com.oracle.cgbu.cne.s cp.soothsayer.subscription.processor.NrfRegistrationProcessor", "message":"{logMsg=Registration will be retried after configured interval, configuredIntervalInSec=6}","endOfBatch":false,"loggerFqcn":"org .apache.logging.log4j.spi.AbstractLogger","threadId":33,"threadP riority":5,"messageTimestamp":"21-05-10 09:28:08.078+0000","application":"ocscp- soothsayer","microservice":"ocscp-scp- subscription","engVersion":"26.1.202","mktgVersion":"26.1.202.0. 0","vendor":"oracle","namespace":"scpsvc","node":"worker1","pod" : "ocscp-scp- subscription-66c68b9db6-6g582","subsystem":"subscription","insta nceType":"prod","processId":"1"} }</pre> Recovery: This alert is cleared automatically when NRF is up and running or errors are corrected for received error responses. For any assistance, contact My Oracle Support.

3.3.23 SCPNrfHeartbeatFailureForRegionOrSetId

Table 3-27 SCPNrfHeartbeatFailureForRegionOrSetId

Field	Description
Description	Alerts is raised when SCP Heartbeat fails.
Summary	SCP Heartbeat to NRF Failure for Region Or SetId: {{ \$labels.nrfRegionOrSetId }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scpfdqn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Major
Conditions	increase(ocscp_subscription_nrf_heartbeat_failures_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.2004
Recommended Actions	Cause: When the Heartbeat fails for an NF Type with NRF. Diagnostic Information: Check whether NRF is up. To check the NRF status, see the <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i>. Check whether the NRF is reachable or not by using one of the following steps: - Run the ping command from primary or secondary nodes by using IP of NRF. Example of a ping command: ping <IPAddress>. - Run the ping command from inside the pod if FQDN of NRF is used. - Instead of using ping, you can collect tcpdump for ensuring the connectivity. Example of a tcpdump: tcpdump -w capture.pcap -i <pod interface>. Check scp-worker logs to find any error response from NRF. If there are error responses, monitor NRF logs: kubectl logs <pod name> -n <namespace>. Recovery: This alert is cleared automatically when NRF is up and running or errors are corrected for received error responses. For any assistance, contact My Oracle Support.

3.3.24 SCPDBOperationFailure

Table 3-28 SCPDBOperationFailure

Field	Description
Severity	Warning
Conditions	increase(ocscp_db_operation_failure_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.2005
Description	Alert is raised for any DB operation failures.

Table 3-28 (Cont.) SCPDBOperationFailure

Field	Description
Recommended Actions	Cause: When the SCP DB operation fails. Diagnostic Information: Check whether the DB service is up. Check the status/age of the mysql pod by using the following command: <code>kubectl get pods -n <namespace></code>. Where, <namespace> is the namespace in which mysql pod is deployed. Recovery: This alert is cleared automatically when the DB service is up and running. For any assistance, contact My Oracle Support.

3.3.25 SCPGeneratedErrorsResponseForNFService

Table 3-29 SCPGeneratedErrorsResponseForNFService

Field	Description
Severity	Info
Conditions	<code>increase(ocscp_metric_scp_generated_response_total[2m]) > 0</code>
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.7016
Description	Alert is raised for NF type for which SCP generated response is triggered.
Recommended Actions	Cause: When the error response is generated for NF Service Type by SCP. Diagnostic Information: Monitor scp-worker logs to determine the reason for error responses generated by SCP. Check for error reason in the logs: <code>kubectl logs <pod name> -n <namespace></code>. Recovery: This alert is cleared automatically when the cause for error response at SCP worker is corrected and configured. For any assistance, contact My Oracle Support.

3.3.26 SCPCircuitBreakingAppliedForNF

Table 3-30 SCPCircuitBreakingAppliedForNF

Field	Description
Severity	Info
Conditions	<code>ocscp_circuit_breaking_applied > 0</code>
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.7017
Description	Alert is raised for NF when circuit breaking is applied.

Table 3-30 (Cont.) SCPCircuitBreakingAppliedForNF

Field	Description
Recommended Actions	Cause: When Circuit Breaking applies for producer NF FQDN based on the configured http2MaxRequests value. Diagnostic Information: Monitor scp-worker logs for number of error responses when outstanding requests exceed the configured http2MaxRequests value: <code>kubectl logs <pod name> -n <namespace></code>. Check the latency to upstream producer from SCP. Use the following metric to check the same: <code>ocscp_metric_upstream_service_time_total(ocscp_producer_host or ocscp_nf_end_point)</code>. Recovery: This alert is cleared automatically when the configuration for http2MaxRequests for circuit breaking is configured beyond the traffic at worker or lower the traffic than the value configured for circuit breaking. For any assistance, contact My Oracle Support.

3.3.27 SCPUpgradeStarted

Table 3-31 SCPUpgradeStarted

Field	Description
Severity	Info
Conditions	When the SCP upgrade process for an SCP microservice starts.
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.6001
Description	Alert is raised when the SCP upgrade process for an SCP microservice starts.
Recommended Actions	Cause: When SCP upgrade is performed for a particular microservice. Diagnostic Information: Not applicable. Recovery: This alert is cleared automatically in 5 minutes when the customAlertExpiryEnabled parameter is set to false in the <code>ocscp_values.yaml</code> file. Otherwise, it is cleared after a specific duration as specified in the customAlertExpiryDuration parameter when the customAlertExpiryEnabled value is true. For any assistance, contact My Oracle Support.

3.3.28 SCPUpgradeFailed

Table 3-32 SCPUpgradeFailed

Field	Description
Severity	Critical
Conditions	When any SCP microservice upgrade fails during the upgrade process.
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.6002
Description	Alert is raised when any SCP microservice upgrade fails.

Table 3-32 (Cont.) SCPUpgradeFailed

Field	Description
Recommended Actions	Cause: When any SCP microservice upgrade fails during the upgrade process. Diagnostic Information: Monitor new hook-jobs that might have failed after multiple attempts. Also, monitor any failed log. Run the following command to check the pod of hook-job: <code>kubectl get pods -n <namespace></code>. Run the following command to check the logs: <code>kubectl logs <pod name> -n <namespace></code>. Recovery: This alert is cleared automatically in 5 minutes when the <code>customAlertExpiryEnabled</code> parameter is set to <code>false</code> in the <code>ocscp_values.yaml</code> file. Otherwise, it is cleared after a specific duration as specified in the <code>customAlertExpiryDuration</code> parameter when the <code>customAlertExpiryEnabled</code> value is <code>true</code>. For any assistance, contact My Oracle Support.

3.3.29 SCPUpgradeSuccessful

Table 3-33 SCPUpgradeSuccessful

Field	Description
Severity	Info
Conditions	When any SCP microservice upgrade is completed.
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.6003
Description	Alert is raised when any SCP microservice upgrade is completed.
Recommended Actions	Cause: When any SCP microservice upgrade is completed. Diagnostic Information: Not applicable. Run the following command to check the pod of hook-job: <code>kubectl get pods -n <namespace></code>. Run the following command to check the logs: <code>kubectl logs <pod name> -n <namespace></code>. Recovery: This alert is cleared automatically in 5 minutes when the <code>customAlertExpiryEnabled</code> parameter is set to <code>false</code> in the <code>ocscp_values.yaml</code> file. Otherwise, it is cleared after a specific duration as specified in the <code>customAlertExpiryDuration</code> parameter when the <code>customAlertExpiryEnabled</code> value is <code>true</code>. For any assistance, contact My Oracle Support.

3.3.30 SCPRollbackStarted

Table 3-34 SCPRollbackStarted

Field	Description
Severity	Info
Conditions	When the rollback process for an SCP microservice starts.
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.6004

Table 3-34 (Cont.) SCPRollbackStarted

Field	Description
Description	Alert is raised when the rollback process for an SCP microservice starts.
Recommended Actions	Cause: When the rollback process for an SCP microservice starts. Diagnostic Information: Not applicable. Recovery: This alert is cleared automatically in 5 minutes when the <code>customAlertExpiryEnabled</code> parameter is set to <code>false</code> in the <code>ocscp_values.yaml</code> file. Otherwise, it is cleared after a specific duration as specified in the <code>customAlertExpiryDuration</code> parameter when the <code>customAlertExpiryEnabled</code> value is <code>true</code>. For any assistance, contact My Oracle Support.

3.3.31 SCPRollbackFailed

Table 3-35 SCPRollbackFailed

Field	Description
Severity	Critical
Conditions	When any SCP microservice rollback fails during the rollback process.
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.6005
Description	Alert is raised when any SCP microservice rollback fails.
Recommended Actions	Cause: When any SCP microservice rollback fails during the rollback process. Diagnostic Information: Monitor new hook-jobs that might have failed after multiple attempts. Also, monitor any failed log. Run the following command to check the pod of hook-job: <pre>kubectl get pods -n <namespace></pre> Run the following command to check the logs: <pre>kubectl logs <pod name> -n <namespace></pre> Recovery: This alert is cleared automatically in 5 minutes when the <code>customAlertExpiryEnabled</code> parameter is set to <code>false</code> in the <code>ocscp_values.yaml</code> file. Otherwise, it is cleared after a specific duration as specified in the <code>customAlertExpiryDuration</code> parameter when the <code>customAlertExpiryEnabled</code> value is <code>true</code>. For any assistance, contact My Oracle Support.

3.3.32 SCPRollbackSuccessful

Table 3-36 SCPRollbackSuccessful

Field	Description
Severity	Info

Table 3-36 (Cont.) SCPRollbackSuccessful

Field	Description
Conditions	When any SCP microservice rollback is completed.
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.6006
Description	Alert is raised when any SCP microservice rollback is completed.
Recommended Actions	Cause: When any SCP microservice rollback is completed. Diagnostic Information: Not applicable. Recovery: This alert is cleared automatically in 5 minutes when the customAlertExpiryEnabled parameter is set to false in the ocscp_values.yaml file. Otherwise, it is cleared after a specific duration as specified in the customAlertExpiryDuration parameter when the customAlertExpiryEnabled value is true. For any assistance, contact My Oracle Support.

3.3.33 ScpWorkerPodCpuUtilizationAboveWarnThreshold

Table 3-37 ScpWorkerPodCpuUtilizationAboveWarnThreshold

Field	Details
Description	CPU utilization of SCP worker at warn level
Summary	CPU utilization of SCP worker at warn level. namespace: <code>{{ \$labels.kubernetes_namespace }}</code> , podname: <code>{{ \$labels.kubernetes_pod_name }}</code> , timestamp: <code>{{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}</code> and value = <code>{{ \$value }}</code>
Severity	Warning
Condition	This alert is raised when CPU utilization of SCP-Worker reaches the WARN level. <code>increase(ocscp_worker_pod_overload_control_cpu_utilization_total{ocscp_threshold_level="WARN"}[2m]) > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7018
Metric Used	ocscp_worker_pod_overload_control_cpu_utilization_total
Recommended Action	Cause: When CPU utilization of scp-worker reaches the WARN level. Diagnostic Information: 1. Get the configured threshold level values using Pod Overload Control Policy REST APIs. For more information about Pod Overload Control Policy, see <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. 2. Check the CPU threshold level status from the scp-worker pod logs under the warn level. Recovery: • Reduce the incoming service request rate. • This alert is automatically cleared when CPU utilization is reduced to below WARN threshold level. For any assistance, contact My Oracle Support.

3.3.34 ScpWorkerPodCpuUtilizationAboveMinorThreshold

Table 3-38 ScpWorkerPodCpuUtilizationAboveMinorThreshold

Field	Details
Description	CPU utilization of SCP worker at minor level
Summary	CPU utilization of SCP worker at minor level. namespace: <code>{{labels.kubernetes_namespace}}</code> , podname: <code>{{labels.kubernetes_pod_name}}</code> , timestamp: <code>{{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}</code> and value = <code>{{ \$value }}</code>
Severity	Minor
Condition	This alert is raised when CPU utilization of scp-worker reaches the MINOR level. <code>increase(ocscp_worker_pod_overload_control_cpu_utilization_total{ocs_cp_threshold_level="MINOR"}[2m]) > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7019
Metric Used	ocscp_worker_pod_overload_control_cpu_utilization_total
Recommended Action	Cause: When CPU utilization of SCP-Worker reaches the MINOR level. Diagnostic Information: 1. Get the configured threshold level values using Pod Overload Control Policy REST APIs. For more information about Pod Overload Control Policy, see <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. 2. Check the CPU threshold level status from the scp-worker pod logs under the warn level. Recovery: • Reduce the incoming service request rate. • This alert is automatically cleared when CPU utilization is reduced to below MINOR threshold level. For any assistance, contact My Oracle Support.

3.3.35 ScpWorkerPodCpuUtilizationAboveMajorThreshold

Table 3-39 ScpWorkerPodCpuUtilizationAboveMajorThreshold

Field	Details
Description	CPU utilization of SCP worker at major level
Summary	Worker CPU utilization lead to major level. namespace: <code>{{labels.kubernetes_namespace}}</code> , podname: <code>{{labels.kubernetes_pod_name}}</code> , timestamp: <code>{{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}</code>
Severity	Major
Condition	This alert is raised when CPU utilization of scp-worker reaches the MAJOR level. <code>increase(ocscp_worker_pod_overload_control_cpu_utilization_total{ocs_cp_threshold_level="MAJOR"}[2m]) > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7020
Metric Used	ocscp_worker_pod_overload_control_cpu_utilization_total

Table 3-39 (Cont.) ScpWorkerPodCpuUtilizationAboveMajorThreshold

Field	Details
Recommended Action	Cause: When CPU utilization of SCP-Worker reaches the MAJOR level. Diagnostic Information: 1. Get the configured threshold level values using Pod Overload Control Policy REST APIs. For more information about Pod Overload Control Policy, see <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. 2. Check the CPU threshold level status from the scp-worker pod logs under the warn level. Recovery: • Reduce the incoming service request rate. • This alert is automatically cleared when CPU utilization is reduced to below MAJOR threshold level. For any assistance, contact My Oracle Support.

3.3.36 ScpWorkerPodCpuUtilizationAboveCriticalThreshold

Table 3-40 ScpWorkerPodCpuUtilizationAboveCriticalThreshold

Field	Details
Description	CPU utilization of SCP worker at critical level
Summary	Worker CPU utilization lead to critical level. namespace: <code>{{labels.kubernetes_namespace}}</code> , podname: <code>{{labels.kubernetes_pod_name}}</code> , timestamp: <code>{{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}</code>
Severity	Critical
Condition	This alert is raised when CPU utilization of scp-worker reaches the CRITICAL level. <code>increase(ocscp_worker_pod_overload_control_cpu_utilization_total{ocs_cp_threshold_level="CRITICAL"}[2m]) > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7021
Metric Used	ocscp_worker_pod_overload_control_cpu_utilization_total
Recommended Action	Cause: When CPU utilization of scp-worker reaches the CRITICAL level. Diagnostic Information: 1. Get the configured threshold level values using Pod Overload Control Policy REST APIs. For more information about Pod Overload Control Policy, see <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. 2. Check the CPU threshold level status from the scp-worker pod logs under the warn level. Recovery: • Reduce the incoming service request rate. • This alert is automatically cleared when CPU utilization is reduced to below CRITICAL threshold level. For any assistance, contact My Oracle Support.

3.3.37 SCPUnhealthyPeerSCPDetected

Table 3-41 SCPUnhealthyPeerSCPDetected

Field	Details
Description	Next hop SCP is marked unhealthy
Summary	'Next hop SCP is marked unhealthy. peerscp host: {{labels.peerScpName}}, scpFqdn: {{labels.scpFqdn}}, namespace: {{labels.kubernetes_namespace}}, podname: {{labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }} and value = {{ \$value }}'
Severity	Info
Condition	This alert is raised when peer SCP is marked as unhealthy. ocscp_peer_scp_unhealthy > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.7022
Metric Used	ocscp_peer_scp_unhealthy
Recommended Action	Cause: The peer SCP is marked as unhealthy because of consecutive failure responses. Diagnostic Information: 1. Check transport failures and routing errors on the peer SCP. 2. Run the ping command from primary or worker nodes using IP of Service. Sample command: ping <IPAddress>. 3. If FQDN of service is used, run the ping command from inside the pod. 4. If the pod does not support the ping command, get the debug container of SCP pod. 5. If you do not want to use the ping command, collect tcpdump to establish the connection. Sample command: tcpdump -w capture.pcap -i <pod interface> Recovery: This alert is automatically cleared after the degradation time is over. Degradation time = Number of consecutive degradations multiplied by configured base ejection. For any assistance, contact My Oracle Support.

3.3.38 SCPDnsSrvQueryFailure

Table 3-42 SCPDnsSrvQueryFailure

Field	Details
Description	DNS SRV Query failed with cause {{labels.cause}}
Summary	'DNS SRV Query failed with cause {{labels.cause}}, namespace: {{labels.kubernetes_namespace}}, podname: {{labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical

Table 3-42 (Cont.) SCPDnsSrvQueryFailure

Field	Details
Condition	This alert is raised when the DNS server lookup for SRV fails due to network, servfail, or timed-out errors. ocscp_alternate_resolution_dnssrv_rx_error_res == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.8001
Metric Used	ocscp_alternate_resolution_dnssrv_rx_error_res
Recommended Action	Cause: When the dnsSRVAlternateRouting flag is set to true, if the DNS SRV lookup fails due to network, servfail, or timed-out errors. Diagnostic Information: Check the DNS SRV server status and re-establish the status to normal. Recovery: This alert is automatically cleared when SCP performs a successful DNS SRV query. For any assistance, contact My Oracle Support.

3.3.39 SCPPProducerOverloadThrottled

Table 3-43 SCPPProducerOverloadThrottled

Field	Details
Description	Producer is in Throttled Overload state
Summary	Producer is in Throttled Overload state.ocscp_peer_fqdn: {{labels.ocscp_peer_fqdn}}, ocscp_peer_nf_instance_id: {{labels.ocscp_peer_nf_instance_id}}, ocscp_peer_service_instance_id: {{labels.ocscp_peer_service_instance_id}}, scpfdn: {{labels.ocscp_fqdn}}, podname: {{labels.kubernetes_pod_name}}, namespace: {{labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Info
Condition	ocscp_load_manager_peer_load_throttled_threshold == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7023
Metric Used	ocscp_load_manager_peer_load_throttled_threshold
Recommended Action	Cause: When the load of producer NF is higher than the throttled threshold configured for the service. Diagnostic Information: 1. Check and configure the throttled threshold for each service using Routing Options REST APIs configurations as described in "Configuring Routing Options" in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. 2. Check the load for each service using the NF Profile REST APIs as described in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. Recovery: This alert clears automatically when the NF profile is deregistered or changed with load less than the throttled abatement threshold. For any assistance, contact My Oracle Support.

3.3.40 SCPPProducerOverloadAlternateRouted

Table 3-44 SCPPProducerOverloadAlternateRouted

Field	Details
Description	Producer is in Alternate Route Overload state
Summary	Producer is in Alternate Route Overload state.ocscp_peer_fqdn: {{labels.ocscp_peer_fqdn}}, ocscp_peer_nf_instance_id: {{labels.ocscp_peer_nf_instance_id}}, ocscp_peer_service_instance_id: {{labels.ocscp_peer_service_instance_id}}, scpfdn: {{labels.ocscp_fqdn}}, podname: {{labels.kubernetes_pod_name}}, namespace: {{labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }}
Severity	Info
Condition	ocscp_load_manager_peer_load_alternateRoute_threshold == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7024
Metric Used	ocscp_load_manager_peer_load_alternateRoute_threshold
Recommended Action	Cause: When the load of producer NF is higher than alternate routing threshold configured for the service. Diagnostic Information: 1. Check and configure the alternate routing threshold for each service using Routing Options REST APIs configurations as described in "Configuring Routing Options" in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. 2. Check the load for each service using the NF Profile REST APIs as described in <i>Oracle Communications Cloud Native Core, Service Communication Proxy REST Specification Guide</i>. Recovery: This alert clears automatically when the NF profile is deregistered or changed with load less than the alternate routing abatement threshold. For any assistance, contact My Oracle Support.

3.3.41 SCPSeppNotConfigured

Table 3-45 SCPSeppNotConfigured

Field	Details
Description	SEPP is not configured for PLMN
Summary	'SEPP is not configured for PLMN'Summary: 'SEPP is not configured for PLMN. plmnid: {{labels.plmn_id}}, scpfdn: {{labels.scp_fqdn}}, namespace: {{labels.kubernetes_namespace}}, podname: {{labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when Security Edge Protection Proxy (SEPP) is not configured. ocscp_metric_sepp_not_configured_current == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7025

Table 3-45 (Cont.) SCPSeppNotConfigured

Field	Details
Metric Used	cscp_metric_sepp_not_configured_current
Recommended Action	Cause: When SEPP routing related rules are not configured at SCP for selected PLMN in the inter-PLMN routing. Diagnostic Information: 1. Check whether SEPP profile is registered or SEPP related configuration is made at SCP. 2. Verify the routing rules created for selected inter-PLMN in the PLMN_SEPP_MAPPING table. Recovery: This alert clears automatically when the SEPP related routing rules are created at SCP for selected PLMN in the inter-PLMN routing. For any assistance, contact My Oracle Support.

3.3.42 SCPSeppRoutingFailed

Table 3-46 SCPSeppRoutingFailed

Field	Details
Description	Routing towards SEPP failed
Summary	Routing towards SEPP failed. sepp_fqdn: {{\$labels.ocscp_sepp_fqdn}}, scpfqdn: {{\$labels.scp_fqdn}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Minor
Condition	This alert is raised when routing towards SEPP fails. ocscp_metric_sepp_routing_attempt_fail_current == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7026
Metric Used	ocscp_metric_sepp_routing_attempt_fail_current

Table 3-46 (Cont.) SCPSeppRoutingFailed

Field	Details
Recommended Action	Cause: Inter-PLMN routing failed for the selected SEPP instances. Diagnostic Information: 1. Check whether selected SEPP is up and healthy. 2. Check if selected SEPP is reachable. 3. Use the ping command from primary or secondary nodes using IP of SEPP. Sample ping command is ping <IpAddress> 4. If FQDN of SEPP is used, try ping command from inside the pod. 5. Alternative to ping, collect tcpdump for ensuring proper connectivity. Sample command: tcpdump -w capture.pcap -i <pod interface> 6. Check scp-worker logs for any error response from selected SEPP. If there are error responses, monitor selected SEPP logs by running this command: kubectl logs <pod name> -n <namespace> Recovery: This alert clears automatically when routing is successful for selected SEPP. For any assistance, contact My Oracle Support.

3.3.43 SCPGlobalEgressRLRemoteParticipantConnectivityFailure

Table 3-47 SCPGlobalEgressRLRemoteParticipantConnectivityFailure

Field	Details
Description	'SCP Global Egress RL Remote Participant Connectivity Failure for participant
Summary	'SCP Global Egress RL Remote Participant Connectivity Failure for participant: {{ \$labels.scp_remote_coh_cluster_name }}, scp_fqdn: {{ \$labels.scp_fqdn }}, scp_local_coh_cluster_name: {{ \$labels.scp_local_coh_cluster_name }}, scp_remote_coh_cluster_fqdn: {{ \$labels.scp_remote_coh_cluster_fqdn }}, scp_remote_coh_cluster_port: {{ \$labels.scp_remote_coh_cluster_port }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when the remote participant SCP connection is not established or goes down. ocscp_global_egress_rl_remote_participant_connectivity_failure == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9001
Metric Used	ocscp_global_egress_rl_bucketkey_not_rate_controlled_total

Table 3-47 (Cont.) SCPGlobalEgressRLRemoteParticipantConnectivityFailure

Field	Details
Recommended Action	Cause: When the remote participant SCP connection is not established or down. Diagnostic Information: 1. Check whether the FQDN or IP port are properly configured with remote SCP's scp-cache microservice. 2. Check whether clusterName and NFInstanceId are configured for the remote SCP. 3. Check whether communication path is active between the two SCP's scp-cache microservice. 4. Run the following command to monitor scp-cache logs: <code>kubectl logs <pod name> -n <namespace></code> Recovery: This alert clears automatically if the connection is established with the remote participant SCP. For any assistance, contact My Oracle Support.

3.3.44 SCPGlobalEgressRLRemoteParticipantWithDuplicateNFInstanceId

Table 3-48 SCPGlobalEgressRLRemoteParticipantWithDuplicateNFInstanceId

Field	Details
Description	This alert is raised when a duplicate remote coherence participant is found.
Summary	SCP Global Egress RL Remote Participant Configured With Duplicate NFInstanceId for participant: {{ \$labels.scp_remote_coh_cluster_name }}, scp_fqdn: {{ \$labels.ocscp_fqdn }}, scp_nf_instance_id: {{ \$labels.scp_nf_instance_id }}, scp_local_coh_cluster_name: {{ \$labels.scp_local_coh_cluster_name }}, scp_remote_coh_cluster_fqdn: {{ \$labels.scp_remote_coh_cluster_fqdn }}, scp_remote_coh_cluster_port: {{ \$labels.scp_remote_coh_cluster_port }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }}
Severity	Major
Condition	ocscp_global_egress_rl_remote_participant_is_duplicate == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.13001
Metric Used	ocscp_global_egress_rl_remote_participant_is_duplicate

Table 3-48 (Cont.) SCPGlobalEgressRLRemoteParticipantWithDuplicateNFInstanceId

Field	Details
Recommended Action	Cause: Duplicate configuration of remote coherence participants with local SCP. Diagnostic Information: 1. Ensure Global Rate Limit feature is enabled. 2. Check whether the clusterName and NFInstanceId of local SCPs and remote SCPs are not duplicates. 3. Check whether the clusterName and NFInstanceId of the first remote SCP and the second remote SCP are not duplicates. 4. Configure unique values for clusterName and NFInstanceId for local SCPs as well as remote SCPs. Recovery: This alert is cleared automatically if no duplicate configurations between local and remote SCPs are found. For any assistance, contact My Oracle Support.

3.3.45 SCPMediationConnectivityFailure

Table 3-49 SCPMediationConnectivityFailure

Field	Details
Description	'SCP Mediation Connectivity Failed, scp_fqdn
Summary	'SCP Mediation Connectivity Failed, scp_fqdn: {{\$labels.scp_fqdn}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when Mediation connection is not established or request to Mediation is not successful. ocscp_mediation_http_not_reachable == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9002
Metric Used	ocscp_mediation_http_not_reachable

Table 3-49 (Cont.) SCPMediationConnectivityFailure

Field	Details
Recommended Action	Cause: The remote Mediation connection is not established or request to Mediation is not successful. Diagnostic Information: 1. Check the errors of the ocscp_mediation_http_not_reachable metric on the Grafana dashboard. 2. Run the following command to check the status of the mediation pod: <pre>kubectl get pods -n <namespace></pre> Recovery: 1. If the mediation pod is not in the ready state, run the following command to check the scp-mediation logs: <pre>kubectl logs <pod name> -n <namespace></pre> 2. If the mediation pod is absent, then redeploy or upgrade SCP with mediationService set to true. This alert clears automatically when the connection is established with Mediation when Mediation is invoked from any of the trigger points. For any assistance, contact My Oracle Support.

3.3.46 SCPNotificationQueuesUtilizationAboveMinorThreshold

Table 3-50 SCPNotificationQueuesUtilizationAboveMinorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Notification Queues Utilization Above Minor Threshold'
Summary	'SCP Notification Queues Utilization Above Minor Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when the queues in the notification service are utilized above 65% of the maximum size (user configure minor threshold value). ocscp_notification_queue_alert{severity="MINOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.3009
Metric Used	ocscp_notification_queue_alert

Table 3-50 (Cont.) SCPNotificationQueuesUtilizationAboveMinorThreshold

Field	Details
Recommended Action	Cause: The Notification module is getting more traffic than expected. Diagnostic Information: 1. Monitor Notification traffic to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_notification_queue_utilization Recovery: This alert clears automatically when notification traffic goes below minor threshold or exceeds major threshold. For any assistance, contact My Oracle Support.

3.3.47 SCPNotificationQueuesUtilizationAboveMajorThreshold

Table 3-51 SCPNotificationQueuesUtilizationAboveMajorThreshold

Field	Details
Description	'instancetype: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Notification Queues Utilization Above Major Threshold'
Summary	'SCP Notification Queues Utilization Above Major Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when the queues in the notification service is utilized above 75% of the maximum size (user configure major threshold value). ocscp_notification_queue_alert{severity="MAJOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.3009
Metric Used	ocscp_notification_queue_alert
Recommended Action	Cause: The Notification module is getting more traffic than expected. Diagnostic Information: 1. Monitor Notification traffic to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_notification_queue_utilization Recovery: This alert clears automatically when notification traffic goes below major threshold or above critical major threshold. For any assistance, contact My Oracle Support.

3.3.48 SCPNotificationQueuesUtilizationAboveCriticalThreshold

Table 3-52 SCPNotificationQueuesUtilizationAboveCriticalThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Notification Queues Utilization Above Critical Threshold'
Summary	'SCP Notification Queues Utilization Above Critical Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the queues in notification service are utilized above 85% of the maximum size (user configure critical threshold value). ocscp_notification_queue_alert{severity="CRITICAL"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.3009
Metric Used	ocscp_notification_queue_alert
Recommended Action	Cause: The Notification module is getting more traffic than expected. Diagnostic Information: 1. Monitor Notification traffic to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_notification_queue_utilization Recovery: This alert clears automatically when notification traffic goes below critical threshold. For any assistance, contact My Oracle Support.

3.3.49 SCPNrfProxyQueuesUtilizationAboveMinorThreshold

Table 3-53 SCPNrfProxyQueuesUtilizationAboveMinorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Nrfproxy Queues Utilization Above Minor Threshold'
Summary	'SCP Nrfproxy Queues Utilization Above Minor Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when the task queues in scp-nrfproxy service are utilized above 65% of the maximum size (user configure minor threshold value). ocscp_nrfproxy_queue_alert{severity="MINOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9010
Metric Used	ocscp_nrfproxy_queue_alert

Table 3-53 (Cont.) SCPNrfProxyQueuesUtilizationAboveMinorThreshold

Field	Details
Recommended Action	Cause: NrfProxy task queues are getting filled and the traffic is more than expected. Diagnostic Information: 1. Monitor traffic towards nrfProxy to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_nrfproxy_queue_utilization Recovery: This alert clears automatically when the traffic goes below minor threshold or above major threshold. For any assistance, contact My Oracle Support.

3.3.50 SCPNrfProxyQueuesUtilizationAboveMajorThreshold

Table 3-54 SCPNrfProxyQueuesUtilizationAboveMajorThreshold

Field	Details
Description	'instancetype: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Nrfproxy Queues Utilization Above Major Threshold'
Summary	'SCP Nrfproxy Queues Utilization Above Major Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when the task queues in scp-nrfproxy service are utilized above 75% of the maximum size (user configure major threshold value). ocscp_nrfproxy_queue_alert{severity="MAJOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9010
Metric Used	ocscp_nrfproxy_queue_alert
Recommended Action	Cause: NrfProxy task queues are getting filled and the traffic is more than expected. Diagnostic Information: 1. Monitor traffic towards nrfProxy to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_nrfproxy_queue_utilization Recovery: This alert clears automatically when the traffic goes below major threshold or above critical threshold. For any assistance, contact My Oracle Support.

3.3.51 SCPNrfProxyQueuesUtilizationAboveCriticalThreshold

Table 3-55 SCPNrfProxyQueuesUtilizationAboveCriticalThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Nrfproxy Queues Utilization Above Critical Threshold'
Summary	'SCP Nrfproxy Queues Utilization Above Critical Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the task queues in scp-nrfproxy service are utilized above 85% of the maximum size (user configure critical threshold value). ocscp_nrfproxy_queue_alert{severity="CRITICAL"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9010
Metric Used	ocscp_nrfproxy_queue_alert
Recommended Action	Cause: NrfProxy task queues are getting filled and the traffic is more than expected. Diagnostic Information: 1. Monitor traffic towards nrfProxy to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_nrfproxy_queue_utilization Recovery: This alert clears automatically when the traffic goes below critical threshold. For any assistance, contact My Oracle Support.

3.3.52 SCPWorkerQueuesUtilizationAboveMinorThreshold

Table 3-56 SCPWorkerQueuesUtilizationAboveMinorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Worker Queues Utilization Above Minor Threshold'
Summary	'SCP Worker Queues Utilization Above Minor Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when task queues in scp-worker service are utilized above 65% of the maximum size (user configure minor threshold value). ocscp_worker_queue_alert{severity="MINOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9007
Metric Used	ocscp_worker_queue_alert

Table 3-56 (Cont.) SCPWorkerQueuesUtilizationAboveMinorThreshold

Field	Details
Recommended Action	Cause: Worker task queues are getting filled and the traffic is more than expected. Diagnostic Information: 1. Monitor traffic towards nrfProxy to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_worker_queue_utilization Recovery: This alert clears automatically when the traffic goes below minor threshold or above major threshold. For any assistance, contact My Oracle Support.

3.3.53 SCPWorkerQueuesUtilizationAboveMajorThreshold

Table 3-57 SCPWorkerQueuesUtilizationAboveMajorThreshold

Field	Details
Description	'instancetype: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Worker Queues Utilization Above Major Threshold'
Summary	'SCP Worker Queues Utilization Above Major Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when task queues in scp-worker service are utilized above 75% of the maximum size (user configure major threshold value). ocscp_worker_queue_alert{severity="MAJOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9007
Metric Used	ocscp_worker_queue_alert
Recommended Action	Cause: Worker task queues are getting filled and the traffic is more than expected. Diagnostic Information: 1. Monitor traffic towards nrfProxy to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_worker_queue_utilization Recovery: This alert clears automatically when the traffic goes below major threshold or goes above critical threshold. For any assistance, contact My Oracle Support.

3.3.54 SCPWorkerQueuesUtilizationAboveCriticalThreshold

Table 3-58 SCPWorkerQueuesUtilizationAboveCriticalThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Worker Queues Utilization Above Critical Threshold'
Summary	'SCP Worker Queues Utilization Above Critical Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when task queues in scp-worker service are utilized above 85% of the maximum size (user configure critical threshold value). ocscp_worker_queue_alert{severity="CRITICAL"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9007
Metric Used	ocscp_worker_queue_alert
Recommended Action	Cause: Worker task queues are getting filled and the traffic is more than expected. Diagnostic Information: 1. Monitor traffic towards nrfProxy to pod using the KPI dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_worker_queue_utilization Recovery: This alert clears automatically when the traffic goes below critical threshold. For any assistance, contact My Oracle Support.

3.3.55 SCPCacheQueuesUtilizationAboveMinorThreshold

Table 3-59 SCPCacheQueuesUtilizationAboveMinorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Cache Queues Utilization Above Minor Threshold'
Summary	'SCP Cache Queues Utilization Above Minor Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when the task queues in the scp-cache service are utilized above 65% of their maximum size (the user-configured minor threshold value). ocscp_cache_queue_alert{severity="MINOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.13000
Metric Used	ocscp_cache_queue_alert

Table 3-59 (Cont.) SCPCacheQueuesUtilizationAboveMinorThreshold

Field	Details
Recommended Action	Cause: When the cache task queues are getting filled, and traffic is higher than expected. Diagnostic Information: 1. Monitor traffic towards the cache pod using the KPI Dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocsdp_cache_queue_utilization. Recovery: The alert is cleared automatically when minor threshold or goes above major threshold. For any assistance, contact My Oracle Support.

3.3.56 SCPCacheQueuesUtilizationAboveMajorThreshold

Table 3-60 SCPCacheQueuesUtilizationAboveMajorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Cache Queues Utilization Above Major Threshold'
Summary	SCP Cache Queues Utilization Above Major Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}
Severity	Major
Condition	This alert is raised when the task queues in the scp-cache service are utilized above 75% of their maximum size (the user-configured major threshold value). ocsdp_cache_queue_alert{severity="MAJOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.13000
Metric Used	ocsdp_cache_queue_alert
Recommended Action	Cause: When the cache task queues are getting filled, and traffic is higher than expected. Diagnostic Information: 1. Monitor traffic towards the cache pod using the KPI Dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocsdp_cache_queue_utilization. Recovery: The alert is cleared automatically when traffic falls below a major threshold or goes above a critical threshold. For any assistance, contact My Oracle Support.

3.3.57 SCPCacheQueuesUtilizationAboveCriticalThreshold

Table 3-61 SCPCacheQueuesUtilizationAboveCriticalThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Cache Queues Utilization Above Critical Threshold'
Summary	'SCP Cache Queues Utilization Above Critical Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the task queues in the scp-cache service are utilized above 85% of their maximum size (the user-configured critical threshold value). ocscp_cache_queue_alert{severity="CRITICAL"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.13000
Metric Used	ocscp_cache_queue_alert
Recommended Action	Cause: When the cache task queues are getting filled, and traffic is higher than expected. Diagnostic Information: 1. Monitor traffic towards the cache pod using the KPI Dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_cache_queue_utilization. Recovery: The alert is cleared automatically when traffic falls below a critical threshold. For any assistance, contact My Oracle Support.

3.3.58 SCPLoadManagerQueuesUtilizationAboveMinorThreshold

Table 3-62 SCPLoadManagerQueuesUtilizationAboveMinorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Load Manage Queues Utilization Above Minor Threshold'
Summary	SCP Load Manager Queues Utilization Above Minor Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when the task queues in the scp-load-manager service are utilized above 65% of their maximum size (the user-configured major threshold value). ocscp_load_manager_queue_alert{severity="MINOR"}
OID	1.3.6.1.4.1.323.5.3.35.1.2.11000
Metric Used	ocscp_load_manager_queue_alert

Table 3-62 (Cont.) SCPLoadManagerQueuesUtilizationAboveMinorThreshold

Field	Details
Recommended Action	Cause: When the cache task queues are getting filled, and traffic is higher than expected. Diagnostic Information: 1. Monitor traffic towards the cache pod using the KPI Dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: <code>ocscp_load_manager_queue_utilization</code>. Recovery: The alert is cleared automatically when traffic falls below a minor threshold. For any assistance, contact My Oracle Support.

3.3.59 SCPLoadManagerQueuesUtilizationAboveMajorThreshold

Table 3-63 SCPLoadManagerQueuesUtilizationAboveMajorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Load Manage Queues Utilization Above Major Threshold'
Summary	'SCP Load Manager Queues Utilization Above Major Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when the task queues in the scp-load-manager service are utilized above 75% of their maximum size (the user-configured major threshold value). <code>ocscp_load_manager_queue_alert{severity="MAJOR"} == 1</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.11000
Metric Used	<code>ocscp_load_manager_queue_alert</code>
Recommended Action	Cause: When the cache task queues are getting filled, and traffic is higher than expected. Diagnostic Information: 1. Monitor traffic towards the cache pod using the KPI Dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: <code>ocscp_load_manager_queue_utilization</code>. Recovery: The alert is cleared automatically when traffic falls below a major threshold. For any assistance, contact My Oracle Support.

3.3.60 SCPLoadManagerQueuesUtilizationAboveCriticalThreshold

Table 3-64 SCPLoadManagerQueuesUtilizationAboveCriticalThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}: SCP Load Manager Queues Utilization Above Critical Threshold'
Summary	'SCP Load Manager Queues Utilization Above Critical Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the task queues in the scp-load-manager service are utilized above 85% of their maximum size (the user-configured critical threshold value). ocscp_load_manager_queue_alert{severity="CRITICAL"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.11000
Metric Used	ocscp_load_manager_queue_alert
Recommended Action	Cause: When the cache task queues are getting filled, and traffic is higher than expected. Diagnostic Information: 1. Monitor traffic towards the cache pod using the KPI Dashboard. 2. Refer to rate of the following metric on the Grafana dashboard: ocscp_load_manager_queue_utilization. Recovery: The alert is cleared automatically when traffic falls below a critical threshold. For any assistance, contact My Oracle Support.

3.3.61 SCPProducerNfSetUnhealthy

Table 3-65 SCPProducerNfSetUnhealthy

Field	Details
Description	All producer NFs in NF set are marked unhealthy
Summary	Summary: 'All producer NFs in NF set are marked unhealthy. nfSet: {{\$labels.ocscp_nf_setid}}, scpFqdn: {{\$labels.scp_fqdn}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when all producer NFs in an NF Set are marked unhealthy. ocscp_metric_nf_set_unhealthy == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7027
Metric Used	ocscp_metric_nf_set_unhealthy

Table 3-65 (Cont.) SCPProducerNfSetUnhealthy

Field	Details
Recommended Action	Cause: All the producer NFs are marked unhealthy because of consecutive failure responses. Diagnostic Information: 1. Check transport failures and routing errors on producer NFs. 2. Run the ping command from primary or secondary nodes using IP of Service. Sample command: ping <IPAddress>. 3. If FQDN of service is used, run the ping command from inside the pod. 4. If the pod does not support the ping command, get the debug container of SCP pod. 5. If you do not want to use the ping command, collect tcpdump to establish the connection. Sample command: tcpdump -w capture.pcap -i <pod interface> Recovery: This alert is automatically cleared after the degradation time is over. Degradation time = Number of consecutive degradations multiplied by configured base ejection. For any assistance, contact My Oracle Support.

3.3.62 SCPPeerSeppUnhealthy

Table 3-66 SCPPeerSeppUnhealthy

Field	Details
Description	Peer Sepp is marked unhealthy
Summary	Peer Sepp is marked unhealthy. seppFqdn: {{ \$labels.ocscp_sepp_fqdn }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }}
Severity	Info
Condition	ocscp_sepp_unhealthy == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7028
Metric Used	ocscp_sepp_unhealthy

Table 3-66 (Cont.) SCPPeerSeppUnhealthy

Field	Details
Recommended Action	Cause: The peer SEPP is marked unhealthy because of consecutive failure responses. Diagnostic Information: 1. Check transport failures and routing errors on peer SCP. 2. Run the ping command from primary or secondary nodes using IP of Service. Sample command: ping <IPAddress>. 3. If FQDN of service is used, run the ping command from inside the pod. 4. If the pod does not support the ping command, get the debug container of SCP pod. 5. If you do not want to use the ping command, collect tcpdump to establish the connection. Sample command: tcpdump -w capture.pcap -i <pod interface> Recovery: This alert is automatically cleared after the degradation time is over. Degradation time = Number of consecutive degradations multiplied by configured base ejection. For any assistance, contact My Oracle Support.

3.3.63 SCPMicroServiceUnreachable

Table 3-67 SCPMicroServiceUnreachable

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}: SCP communication between the micro-services indicated by source and destination has failed'
Summary	Summary: 'SCP communication between the micro-services indicated by source and destination has failed: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, source:{{\$labels.source}}, destination: {{\$labels.destination}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the communication between SCP microservices indicated by source and destination has failed. ocscp_metric_svc_unreachable==1
OID	1.3.6.1.4.1.323.5.3.35.1.2.7029
Metric Used	ocscp_metric_svc_unreachable
Recommended Action	Cause: Communication between SCP microservices has failed. Diagnostic Information: Verify whether endpoints of all the services are in Running and Ready state. If not, restart the services. Recovery: This alert clears automatically when the required services are in Running and Ready state. For any assistance, contact My Oracle Support.

3.3.64 SCPTrafficFeedSendFailed

Table 3-68 SCPTrafficFeedSendFailed

Field	Details
Description	Sending messages to Traffic Feed failed. Cause : {{labels.ocscp_cause}}
Summary	Sending messages to Traffic Feed failed, cause: {{labels.ocscp_cause}}, scp_fqdn: {{labels.ocscp_fqdn}}, namespace: {{labels.kubernetes_namespace}}, podname: {{labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Minor
Condition	increase(ocscp_metric_trafficfeed_failed_total{app_kubernetes_io_name="scp-worker"}[1h]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9003
Metric Used	ocscp_metric_trafficfeed_failed_total
Recommended Action	Cause: Sending of message to traffic feed failed. Diagnostic Information: • Check failure reason. • Check if the traffic feed OCNADD configuration is correct. • Check if the OCNADD server is reachable and available. Recovery: This alert clears automatically after 24 hrs if sending messages to traffic feed stops failing. For any assistance, contact My Oracle Support.

3.3.65 SCPTrafficFeedKafkaClusterUnhealthy

Table 3-69 SCPTrafficFeedKafkaClusterUnhealthy

Field	Details
Description	'Kafka cluster is marked unhealthy, Cause : {{labels.ocscp_cause}}'
Summary	'Kafka cluster is marked unhealthy, cause: {{labels.ocscp_cause}}, scp_fqdn: {{labels.ocscp_fqdn}}, namespace: {{labels.kubernetes_namespace}}, podname: {{labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the Kafka cluster is unhealthy. ocscp_metric_trafficfeed_cluster_unhealthy == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9026
Metric Used	ocscp_metric_trafficfeed_cluster_unhealthy
Recommended Action	Cause: The Kafka cluster is unhealthy. Diagnostic Information: • Check and diagnose OCNADD Kafka cluster. Recovery: This alert clears when the Kafka cluster recovers from the failure condition. For any assistance, contact My Oracle Support.

3.3.66 SCPTrafficFeedPartitionUnhealthy

Table 3-70 SCPTrafficFeedPartitionUnhealthy

Field	Details
Description	'Kafka partition {{labels.kafka_partition_id}} is marked unhealthy, Cause : {{labels.ocscp_cause}}'
Summary	'Kafka cluster is marked unhealthy, cause: {{labels.ocscp_cause}}, partition_id: {{labels.kafka_partition_id}}, topic: {{labels.topic}}, scp_fqdn: {{labels.ocscp_fqdn}}, namespace: {{labels.namespace}}, podname: {{labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when the Kafka partition is unhealthy. ocscp_metric_trafficfeed_partition_unhealthy == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9025
Metric Used	ocscp_metric_trafficfeed_partition_unhealthy
Recommended Action	Cause: The Kafka partition is unhealthy. Diagnostic Information: Check and diagnose OCNADD Kafka cluster. Recovery: This alert clears when the Kafka partition recovers from the failure condition. For any assistance, contact My Oracle Support .

3.3.67 SCPServiceMeshFailure

Table 3-71 SCPServiceMeshFailure

Field	Details
Description	'SCP servicemesh failure encountered'
Summary	'SCP servicemesh failure encountered for nfservicetype: {{labels.ocscp_nf_service_type}}, nftype: {{labels.ocscp_nf_type}}, nfinstanceid: {{labels.ocscp_nf_instance_id}}, serviceinstanceid: {{labels.ocscp_service_instance_id}}, producerfqdn: {{labels.ocscp_producer_host}}, responsecode: {{labels.ocscp_response_code}}, serverheader: {{labels.ocscp_server_header}}, namespace: {{labels.namespace}}, podname: {{labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}SeverityInfo
Severity	Info
Condition	This alert is raised when service mesh failure occurs. increase(ocscp_metric_sidecarproxy_failures_total[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.7030
Metric Used	ocscp_metric_sidecarproxy_failures_total

Table 3-71 (Cont.) SCPServiceMeshFailure

Field	Details
Recommended Action	Cause: Service mesh failure observed at SCP. Diagnostic Information: - Service mesh is unable to reach peer NF due to connection failures or host is not known to service mesh or some other errors at service mesh. - Verify the service mesh logs to detect error details. - Check sidecar status of peer NF. Recovery: This alert clears automatically after 2 minutes if there is no service mesh failure observed by SCP with the same dimensions. For any assistance, contact My Oracle Support.

3.3.68 SCPHealthCheckFailedForPeerSCP

Table 3-72 SCPHealthCheckFailedForPeerSCP

Field	Details
Description	'SCP HealthCheck failed for peer SCP'
Summary	'SCP HealthCheck failed for peer SCP. namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when peer SCP or inter-SCP becomes unhealthy due to health check status and outlier detection. ocscp_interscp_health_check_status_failed == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9023
Metric Used	ocscp_interscp_health_check_status_failed
Recommended Action	Cause: When peer SCP is unhealthy to receive any SBI message requests due to health check and outlier detection. Diagnostic Information: - Monitor if the overall average load of peer SCP is greater than the configured threshold value. - Check outlier detection. Recovery: This alert clears automatically if SCP-C decides SCP-P is healthy or available based on the current and previous status of outlier detection and health check. For any assistance, contact My Oracle Support.

3.3.69 SCPHealthCheckFailed

Table 3-73 SCPHealthCheckFailed

Field	Details
Description	'SCP HealthCheck failed'
Summary	'SCP HealthCheck failed. namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'

Table 3-73 (Cont.) SCPHealthCheckFailed

Field	Details
Severity	Info
Condition	This alert is raised when SCP is unhealthy because the overall average load of SCP is greater than the configured threshold. ocscp_health_check_status_failed == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.9024
Metric Used	ocscp_health_check_status_failed
Recommended Action	Cause: When SCP is unhealthy to receive any SBI message requests due to the overall average load. Diagnostic Information: Monitor if the overall average load of SCP is greater than the configured threshold value. Recovery: This alert clears automatically when the overall average load of SCP is less than the configured threshold value. For any assistance, contact My Oracle Support.

3.3.70 ScpWorkerPodPendingTransUtilizationAboveWarnThreshold

Table 3-74 ScpWorkerPodPendingTransUtilizationAboveWarnThreshold

Field	Details
Description	Worker Pending Transaction lead to Warn level
Summary	'Worker Pending Transaction lead to warn level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Warn
Condition	This alert is raised when pending transaction utilization of SCP-Worker reaches WARN level. ocscp_worker_pod_overload_control_pendingTrans_utilization_warn > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9014
Metric Used	ocscp_worker_pod_overload_control_pendingTrans_utilization_warn
Recommended Action	Cause: When pending transactions utilization of SCP-Worker reaches WARN level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of warn pending transaction utilization can be checked from database. - Check WARN level logs of SCP-Worker for pending transaction status. Recovery: This alert clears automatically when pending transaction utilization is below WARN threshold level. For any assistance, contact My Oracle Support.

3.3.71 ScpWorkerPodPendingTransUtilizationAboveMinorThreshold

Table 3-75 ScpWorkerPodPendingTransUtilizationAboveMinorThreshold

Field	Details
Description	Worker Pending Transaction lead to minor level

Table 3-75 (Cont.) ScpWorkerPodPendingTransUtilizationAboveMinorThreshold

Field	Details
Summary	'Worker Pending Transaction lead to minor level.namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when pending transaction utilization of SCP-Worker reaches MINOR level. ocscp_worker_pod_overload_control_pendingTrans_utilization_minor > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9014
Metric Used	ocscp_worker_pod_overload_control_pendingTrans_utilization_minor
Recommended Action	Cause: When pending transactions utilization of SCP-Worker reaches MINOR level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of minor pending transaction utilization can be checked from database. - Check MINOR level logs of SCP-Worker for pending transaction status. Recovery: This alert clears automatically when pending transaction utilization is below MINOR threshold level. For any assistance, contact My Oracle Support .

3.3.72 ScpWorkerPodPendingTransUtilizationAboveMajorThreshold

Table 3-76 ScpWorkerPodPendingTransUtilizationAboveMajorThreshold

Field	Details
Description	Worker Pending Transaction lead to major level
Summary	'Worker Pending Transaction lead to major level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when pending transaction utilization of SCP-Worker reaches MAJOR level. ocscp_worker_pod_overload_control_pendingTrans_utilization_major > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9014
Metric Used	ocscp_worker_pod_overload_control_pendingTrans_utilization_major
Recommended Action	Cause: When pending transactions utilization of SCP-Worker reaches MAJOR level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of major pending transaction utilization can be checked from database. - Check MAJOR level logs of SCP-Worker for pending transaction status. Recovery: This alert clears automatically when pending transaction utilization is below MAJOR threshold level. For any assistance, contact My Oracle Support .

3.3.73 ScpWorkerPodPendingTransUtilizationAboveCriticalThreshold

Table 3-77 ScpWorkerPodPendingTransUtilizationAboveCriticalThreshold

Field	Details
Description	Worker Pending Transaction lead to critical level
Summary	'Worker Pending Transaction lead to critical level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when pending transaction utilization of SCP-Worker reaches CRITICAL level. ocscp_worker_pod_overload_control_pendingTrans_utilization_critical > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9014
Metric Used	ocscp_worker_pod_overload_control_pendingTrans_utilization_critical
Recommended Action	Cause: When pending transactions utilization of SCP-Worker reaches CRITICAL level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of critical pending transaction utilization can be checked from database. - Check CRITICAL level logs of SCP-Worker for pending transaction status. Recovery: This alert clears automatically when pending transaction utilization is below CRITICAL threshold level. For any assistance, contact My Oracle Support.

3.3.74 ScpWorkerPodResourceUtilizationAboveMinorThreshold

Table 3-78 ScpWorkerPodResourceUtilizationAboveMinorThreshold

Field	Details
Description	Worker overload control lead to minor level
Summary	'Worker overload control lead to minor level.namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when overload control resource utilization of SCP-Worker reaches MINOR level. ocscp_worker_pod_overload_control_resource_utilization_minor > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9018
Metric Used	ocscp_worker_pod_overload_control_resource_utilization_minor

Table 3-78 (Cont.) ScpWorkerPodResourceUtilizationAboveMinorThreshold

Field	Details
Recommended Action	Cause: When overload control resource utilization of SCP-Worker reaches MINOR level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of minor overload control resource utilization can be checked from database. - Check MINOR level logs of SCP-Worker for overload control resource utilization status. Recovery: This alert clears automatically when overload control resource utilization is below MINOR threshold level. For any assistance, contact My Oracle Support.

3.3.75 ScpWorkerPodResourceUtilizationAboveMajorThreshold

Table 3-79 ScpWorkerPodResourceUtilizationAboveMajorThreshold

Field	Details
Description	Worker overload control lead to major level
Summary	'Worker overload control lead to major level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when overload control resource utilization of SCP-Worker reaches MAJOR level. ocscp_worker_pod_overload_control_resource_utilization_major > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9018
Metric Used	ocscp_worker_pod_overload_control_resource_utilization_major
Recommended Action	Cause: When overload control resource utilization of SCP-Worker reaches MAJOR level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of major overload control resource utilization can be checked from database. - Check MAJOR level logs of SCP-Worker for overload control resource utilization status. Recovery: This alert clears automatically when overload control resource utilization is below MAJOR threshold level. For any assistance, contact My Oracle Support.

3.3.76 ScpWorkerPodResourceUtilizationAboveWarnThreshold

Table 3-80 ScpWorkerPodResourceUtilizationAboveWarnThreshold

Field	Details
Description	'Worker overload control lead to Warn level'
Summary	'Worker overload control lead to warn level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'

Table 3-80 (Cont.) ScpWorkerPodResourceUtilizationAboveWarnThreshold

Field	Details
Severity	Warning
Condition	This alert is raised when overload control resource utilization of SCP-Worker reaches WARN level. ocscp_worker_pod_overload_control_resource_utilization_warn > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9018
Metric Used	ocscp_worker_pod_overload_control_resource_utilization_warn
Recommended Action	Cause: When overload control resource utilization of SCP-Worker reaches WARN level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of warn overload control resource utilization can be checked from database. - Check WARN level logs of SCP-Worker for overload control resource utilization status. Recovery: This alert clears automatically when overload control resource utilization is below WARN threshold level. For any assistance, contact My Oracle Support .

3.3.77 ScpWorkerPodResourceUtilizationAboveCriticalThreshold

Table 3-81 ScpWorkerPodResourceUtilizationAboveCriticalThreshold

Field	Details
Description	Worker overload control lead to critical level
Summary	"Worker overload control lead to critical level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}"
Severity	Critical
Condition	This alert is raised when overload control resource utilization of SCP-Worker reaches CRITICAL level. ocscp_worker_pod_overload_control_resource_utilization_critical > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.9018
Metric Used	ocscp_worker_pod_overload_control_resource_utilization_critical
Recommended Action	Cause: When overload control resource utilization of SCP-Worker reaches CRITICAL level. Diagnostic Information: - Monitor pending transactions usage while processing traffic. The threshold value of critical overload control resource utilization can be checked from database. - Check CRITICAL level logs of SCP-Worker for overload control resource utilization status. Recovery: This alert clears automatically when overload control resource utilization is below CRITICAL threshold level. For any assistance, contact My Oracle Support .

3.3.78 SCPDNSSRVNRFMigrationTaskFailure

Table 3-82 SCPDNSSRVNRFMigrationTaskFailure

Field	Description
Severity	critical
Condition	ocscp_configuration_dnssrv_nrf_migration_task_failure == 1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15001
Description	An alert is raised to notify that migration from static to DNS has failed.
Recommended Actions	Cause: - If migration to DNS task was waiting for DNS SRV data and wait time elapsed. - If any migration task fails due to no acknowledgement from other microservices. - Wait time elapsed in task completion response for any migration task. Diagnostic Information: Monitor that all the DNS SRV configurations are proper and that all SCP pods are up and running in the proper state. Recovery: - DNS SRV data, wait time elapsed: Once data is received from DNS, the alert will be cleared. - No acknowledgement: Keep retrying for success acknowledgement and clear or remove the alert on receiving success acknowledgement. - Wait time elapsed for completion response: Resend the event and wait. Repeat this until a completion response is received, and on receiving a completion response, clear the alert. In all the raised conditions, the alert will also be cleared on receiving the new migration task. For any assistance, contact My Oracle Support.

3.3.79 SCPDNSSRVNRFNonMigrationTaskFailure

Table 3-83 SCPDNSSRVNRFNonMigrationTaskFailure

Field	Description
Severity	critical
Condition	ocscp_configuration_dnssrv_nrf_non_migration_task_failure == 1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15003
Description	An alert is raised to notify that the non-migrated task has failed.
Recommended Actions	Cause: - If a non-migrated task fails due to no acknowledgement from other microservices. - Wait time elapsed in task completion response. Diagnostic Information: Monitor that all the DNS SRV configurations are proper and that all SCP pods are up and running in the proper state. Recovery: - No acknowledgement: Keep retrying for success acknowledgement and clear or remove the alert on receiving success acknowledgement. - Wait time elapsed: When a new task is submitted for the same SPN, this alert is immediately cleared; if subsequent task processing fails, again, the alert will be raised. For any assistance, contact My Oracle Support.

3.3.80 SCPDNSSRVNRFDuplicateTargetDetected

Table 3-84 SCPDNSSRVNRFDuplicateTargetDetected

Field	Description
Severity	critical
Condition	ocscp_configuration_dnssrv_nrf_duplicate_target_detected == 1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15002
Description	An alert is raised to notify that a duplicate target NRF has been detected in the DNS SRV records.
Recommended Actions	Cause: This alert is raised when a duplicate target FQDN is received from the DNS SRV for different NRF SRV FQDN(s). In this case, the target FQDNs received against the first NRF SRV FQDN in the scpc-configuration service from the scpc-alternate-resolution service shall be processed, but the target FQDNs received against the subsequent NRF SRV FQDN will be ignored, and this alert shall be raised. Diagnostic Information: Monitor that all the DNS SRV configurations are proper and that all SCP pods are up and running in the proper state. Recovery: If for the same NRF SRV FQDN receives non-overlapping target FQDNs, then the alert will be cleared. For any assistance, contact My Oracle Support.

3.3.81 SCPHighResponseTimeFromProducer

Table 3-85 SCPHighResponseTimeFromProducer

Field	Description
Description	It notifies when the traffic exceeds 200 messages per second and the response delay from the Producer NF takes more than 50 seconds.
Summary	More than 200 responses received by the SCP have a response time exceeding 50,000 milliseconds. Instance name: {{ \$labels.instance }}, Namespace: {{ \$labels.namespace }}, Pod name: {{ \$labels.pod }}, Timestamp: {{ with query "time()" }} {{ . first value humanizeTimestamp }}{{ end }}
Severity	Info
Condition	(sum(rate(ocscp_metric_upstream_service_time_total{ocscp_upstream_service_time="50000ms"}[2m])) by (kubernetes_namespace) + sum(rate(ocscp_metric_upstream_service_time_total{ocscp_upstream_service_time=">50000ms"}[2m])) by (kubernetes_namespace)) > 200
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15004
Metric Used	ocscp_metric_upstream_service_time_total

Table 3-85 (Cont.) SCPHighResponseTimeFromProducer

Field	Description
Recommended Actions	Cause: More than 200 messages per second have an upstream response time above 50 seconds. Diagnostic Information: Monitor metric metricocscp_metric_upstream_service_time_total with ocscp_upstream_service_time="50000ms" and ocscp_upstream_service_time=">50000ms". Recovery: The alert is cleared automatically when the number of responses with delays exceeding 50 seconds falls below 200 messages per second. If the alert does not clear, check for any producer NFs or specific service request types that are taking more than 50 seconds to respond and take corrective actions if necessary. Note that immediate action may not be required, as this alert is informational. However, a high number of requests with long response delays could lead to performance degradation at the SCP. For any assistance, contact My Oracle Support.

3.3.82 SCPCGroupVersionDetectionFailed

Table 3-86 SCPCGroupVersionDetectionFailed

Field	Description
Severity	critical
Condition	ocscp_worker_cgroup_version_detection_failed == 1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15005
Description	Notify that cgroup version detection has failed.
Recommended Actions	Cause: SCP is unable to detect the cgroup version from the underlying kernel with the command "stat -fc %T /sys/fs/cgroup/." The possible expected value is either tmpfs or cgroup2fs. Diagnostic Information: - Option 1: Check worker error level logs for failure of cgroup version detection. - Option 2: - Login to the worker pod and run the command "stat -fc %T /sys/fs/cgroup." - Verify whether it outputs either tmpfs or cgroup2fs. Recovery: - Make sure the cgroup has either tmpfs or cgroup2fs. - Do same version upgrade on SCP. For any assistance, contact My Oracle Support.

3.3.83 SCPCPUUsageFileReadFailed

Table 3-87 SCPCPUUsageFileReadFailed

Field	Description
Severity	critical
Condition	ocscp_worker_cpu_usage_file_read_failed == 1
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15006

Table 3-87 (Cont.) SCPCPUUsageFileReadFailed

Field	Description
Description	Notify that the CPU usage file read operation failed within the detected cgroup version.
Recommended Actions	Cause: SCP encountered a failure in performing a read operation for the CPU usage file within the detected cgroup version. The file path is determined based on the detected cgroup version. Diagnostic Information: - Option 1: Check worker warn level logs for failure in performing a read operation for the CPU usage file. - Option 2: - Log in to the worker pod and run the command "stat -fc %T /sys/fs/cgroup" to confirm whether it outputs either 'tmpfs' or 'cgroup2fs'. - Depending on the detected cgroup version, inspect the file located at /sys/fs/cgroup/cpu/cpuacct.usage for 'tmpfs' or /sys/fs/cgroup/cpu.stat for 'cgroup2fs'. - Ensure that the file exists at the specified path and possesses the required permissions for read operations. Recovery: - Verify that the files are appropriate according to the cgroup version and possess the necessary permissions for read operations. - Perform the same version upgrade on the SCP. For any assistance, contact My Oracle Support.

3.3.84 SCPIgnoreUnknownService

Table 3-88 SCPIgnoreUnknownService

Field	Description
Severity	Info
Condition	increase(ocscp_ignore_unknown_service_total[24h]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15000
Description	An alert is raised to notify that SCP ignored an unknown service in the NF profile.
Recommended Actions	Cause: SCP has received the NF profile with an unknown service and processed the profile by ignoring this unknown service. Diagnostic Information: Check the received NF profile for the unknown services. Recovery: If the unknown services are not present in the NF profile in the next scrapping interval, then the alert will be cleared. For any assistance, contact My Oracle Support.

3.3.85 SCPWorkerSSLCertificateOnCriticalExpiry

Table 3-89 SCPWorkerSSLCertificateOnCriticalExpiry

Field	Description
Severity	Critical
Condition	ocscp_metric_ssl_certificate_expire_total == 1

Table 3-89 (Cont.) SCPWorkerSSLCertificateOnCriticalExpiry

Field	Description
OID used for SNMP Traps	11.3.6.1.4.1.323.5.3.35.1.2.15010
Description	An alert is raised whenever the SCP SSL certificate is about to expire, based on the configured threshold values.
Recommended Actions	Cause: The SSL certificate expiration is approaching the configured critical expiry time. Diagnostic Information: - Whenever this alert is raised, it indicates that the SSL certificates configured for SCP are about to expire. - Verify the certificate expiry date. Recovery: The SCP SSL secret needs to be updated with renewed SSL certificates. For any assistance, contact My Oracle Support.

3.3.86 SCPWorkerSSLCertificateOnMajorExpiry

Table 3-90 SCPWorkerSSLCertificateOnMajorExpiry

Field	Description
Severity	major
Condition	ocscp_metric_ssl_certificate_expire_total == 2
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15011
Description	An alert is raised whenever the SCP SSL certificate is about to expire, based on the configured threshold values.
Recommended Actions	Cause: The SSL certificate expiration is approaching the configured major expiry time. Diagnostic Information: - Whenever this alert is raised, it indicates that the SSL certificates configured for SCP are about to expire. - Verify the certificate expiry date. Recovery: The SCP SSL secret needs to be updated with renewed SSL certificates. For any assistance, contact My Oracle Support.

3.3.87 SCPWorkerSSLCertificateOnMinorExpiry

Table 3-91 SCPWorkerSSLCertificateOnMinorExpiry

Field	Description
Severity	minor
Condition	ocscp_metric_ssl_certificate_expire_total == 3
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15012
Description	An alert is raised whenever the SCP SSL certificate is about to expire, based on the configured threshold values.

Table 3-91 (Cont.) SCPWorkerSSLCertificateOnMinorExpiry

Field	Description
Recommended Actions	Cause: The SSL certificate expiration is approaching the configured minor expiry time. Diagnostic Information: - Whenever this alert is raised, it indicates that the SSL certificates configured for SCP are about to expire. - Verify the certificate expiry date. Recovery: The SCP SSL secret needs to be updated with renewed SSL certificates. For any assistance, contact My Oracle Support.

3.3.88 SCPIngressConnectionEstablishmentFailure

Table 3-92 SCPIngressConnectionEstablishmentFailure

Field	Description
Severity	info
Condition	increase(ocscp_worker_https_ingress_connection_failure_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15007
Description	An alert is raised whenever any ingress HTTPS connection is failed.
Recommended Actions	Cause: Whenever an Ingress HTTPS connection establishment fails. Diagnostic Information: This alert may be raised if any connection establishment fails or if a handshake fails. Recovery: The alert will be cleared if there are no ingress HTTPS connection failures in the next scrape interval (2 minutes). For any assistance, contact My Oracle Support.

3.3.89 SCPEgressConnectionEstablishmentFailure

Table 3-93 SCPEgressConnectionEstablishmentFailure

Field	Description
Severity	info
Condition	increase(ocscp_worker_https_egress_connection_failure_total[2m]) > 0
OID used for SNMP Traps	1.3.6.1.4.1.323.5.3.35.1.2.15008
Description	An alert is raised whenever any egress HTTPS connection is failed.
Recommended Actions	Cause: Whenever an egress HTTPS connection establishment fails to send the request to producer NFs. Diagnostic Information: This alert may be raised if any connection establishment fails or if a handshake fails. Recovery: The alert will be cleared if there are no engress HTTPS connection failures in the next scrape interval (2 minutes). For any assistance, contact My Oracle Support.

3.3.90 SCPNrfProxyOauthQueuesUtilizationAboveMinorThreshold

Table 3-94 SCPNrfProxyOauthQueuesUtilizationAboveMinorThreshold

Field	Details
Description	description: 'instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}: SCP NrfProxyOauth Queues Utilization Above Minor Threshold'
Summary	'SCP NrfProxy Oauth Queues Utilization Above Minor Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when SCP NrfProxy Oauth Queues Utilization is above the minor threshold. ocscp_nrfproxy_oauth_queue_alert{severity="MINOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.14000
Metric Used	ocscp_nrfproxy_oauth_queue_alert
Recommended Action	Cause: When NrfProxy Task queues are filled, the traffic exceeds the limit. Diagnostic Information: - Monitor the traffic toward nrfProxy to pod using the Grafana dashboard. - Refer to the rate of ocscp_nrfproxy_oauth_queue_alert metric on the Grafana dashboard. Recovery: This alert clears automatically when the traffic decreases below the Major threshold. For any assistance, contact My Oracle Support.

3.3.91 SCPNrfProxyOauthQueuesUtilizationAboveMajorThreshold

Table 3-95 SCPNrfProxyOauthQueuesUtilizationAboveMajorThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}: SCP NrfProxyOauth Queues Utilization Above Major Threshold'
Summary	'SCP NrfProxy Oauth Queues Utilization Above Major Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when SCP NrfProxy Oauth Queues Utilization is above the major threshold. ocscp_nrfproxy_oauth_queue_alert{severity="MAJOR"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.14000
Metric Used	ocscp_nrfproxy_oauth_queue_alert

Table 3-95 (Cont.) SCPNrfProxyOauthQueuesUtilizationAboveMajorThreshold

Field	Details
Recommended Action	Cause: When NrfProxy Task queues are filled, the traffic exceeds the limit. Diagnostic Information: - Monitor the traffic toward nrfProxy to pod using the Grafana dashboard. - Refer to the rate of ocscp_nrfproxy_oauth_queue_alert metric on the Grafana dashboard. Recovery: This alert clears automatically when the traffic decreases below the Major threshold. For any assistance, contact My Oracle Support.

3.3.92 SCPNrfProxyOauthQueuesUtilizationAboveCriticalThreshold

Table 3-96 SCPNrfProxyOauthQueuesUtilizationAboveCriticalThreshold

Field	Details
Description	'instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}: SCP NrfProxy Oauth Queues Utilization Above Critical Threshold'
Summary	'SCP NrfProxy Oauth Queues Utilization Above Critical Threshold, instancename: {{\$labels.instance}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when SCP NrfProxy Oauth Queues Utilization is above the critical threshold. ocscp_nrfproxy_oauth_queue_alert{severity="CRITICAL"} == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.14000
Metric Used	ocscp_nrfproxy_oauth_queue_alert
Recommended Action	Cause: When NrfProxy Task queues are filled, the traffic exceeds the limit. Diagnostic Information: - Monitor the traffic toward nrfProxy to pod using the Grafana dashboard. - Refer to the rate of ocscp_nrfproxy_oauth_queue_alert metric on the Grafana dashboard. Recovery: This alert clears automatically when the traffic decreases below the critical threshold. For any assistance, contact My Oracle Support.

3.3.93 ScpSelfOCIThresholdAboveWarn

Table 3-97 ScpSelfOCIThresholdAboveWarn

Field	Details
Description	An alert is raised whenever the configured self-OCI for SCP reaches a warn level.
Summary	SCP load level lead to Warn level. namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}
Severity	warning

Table 3-97 (Cont.) ScpSelfOCIThresholdAboveWarn

Field	Details
Condition	increase(ocscp_worker_self_oci_above_oci_conveyance_threshold_total{ocscp_oci_t hreshold_level="WARN"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15013
Metric Used	ocscp_worker_self_oci_above_oci_conveyance_threshold_total
Recommended Action	Cause: Whenever the configured self-OCI for SCP reaches a warn level. Diagnostic Information: SCP may be experiencing high load. Recovery: The alert will clear automatically when the load on SCP decreases and falls below the configured warn level. For any assistance, contact My Oracle Support .

3.3.94 ScpSelfOCIThresholdAboveMinor

Table 3-98 ScpSelfOCIThresholdAboveMinor

Field	Details
Description	An alert is raised whenever the configured self-OCI for SCP reaches a minor level.
Summary	SCP load level lead to Minor level. namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	minor
Condition	increase(ocscp_worker_self_oci_above_oci_conveyance_threshold_total{ocscp_oci_t hreshold_level="MINOR"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15014
Metric Used	ocscp_worker_self_oci_above_oci_conveyance_threshold_total
Recommended Action	Cause: Whenever the configured self-OCI for SCP reaches a minor level. Diagnostic Information: SCP may be experiencing high load. Recovery: The alert will clear automatically when the load on SCP decreases and falls below the configured minor level. For any assistance, contact My Oracle Support .

3.3.95 ScpSelfOCIThresholdAboveMajor

Table 3-99 ScpSelfOCIThresholdAboveMajor

Field	Details
Description	An alert is raised whenever the configured self-OCI for SCP reaches a major level.
Summary	SCP load level lead to Major level. namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	major
Condition	increase(ocscp_worker_self_oci_above_oci_conveyance_threshold_total{ocscp_oci_t hreshold_level="MAJOR"}[2m]) > 0

Table 3-99 (Cont.) ScpSelfOCIThresholdAboveMajor

Field	Details
OID	1.3.6.1.4.1.323.5.3.35.1.2.15015
Metric Used	ocscp_worker_self_oci_above_oci_conveyance_threshold_total
Recommended Action	Cause: Whenever the configured self-OCI for SCP reaches a major level. Diagnostic Information: SCP may be experiencing high load. Recovery: The alert will clear automatically when the load on SCP decreases and falls below the configured major level. For any assistance, contact My Oracle Support.

3.3.96 ScpSelfOCIThresholdAboveCritical

Table 3-100 ScpSelfOCIThresholdAboveCritical

Field	Details
Description	This alert will be activated whenever the configured self-OCI for SCP reaches a critical level.
Summary	SCP load level lead to Critical level. namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }} {{ . first value humanizeTimestamp }} {{ end }}
Severity	critical
Condition	increase(ocscp_worker_self_oci_above_oci_conveyance_threshold_total{ocscp_oci_threshold_level="CRITICAL"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15016
Metric Used	ocscp_worker_self_oci_above_oci_conveyance_threshold_total
Recommended Action	Cause: Whenever the configured self-OCI for SCP reaches a critical level. Diagnostic Information: SCP may be experiencing high load. Recovery: The alert will clear automatically when the load on SCP decreases and falls below the configured critical level. For any assistance, contact My Oracle Support.

3.3.97 SCPWorkerSSLPartialSecret

Table 3-101 SCPWorkerSSLPartialSecret

Field	Details
Description	SCP Worker SSL Secret is Invalid or Partial'
Summary	'SCP Worker SSL Secret is Invalid or Partial, ocscp_worker_fqdn: {{ \$labels.ocscp_producer_host }}, scpfdn: {{ \$labels.ocscp_fqdn }}, namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }} {{ . first value humanizeTimestamp }} {{ end }}', secretName: {{ \$labels.secretName }} {{ . first value humanizeTimestamp }} {{ end }}
Severity	Major

Table 3-101 (Cont.) SCPWorkerSSLPartialSecret

Field	Details
Condition	This alert is raised when a Kubernetes secret is created or patched with incomplete information. A Kubernetes secret contains private key, certificate, CA bundle, truststore, and keystore password files. If any of these components are missing during creation or patching, this alert is raised. ocscp_worker_ssl_partial_secret == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.15017
Metric Used	ocscp_worker_ssl_partial_secret
Recommended Action	Cause: When a Kubernetes secret is created or patched with incomplete information. Diagnostic Information: Verify the patched secret whether it contains all the required data. Recovery: Alert is cleared automatically when valid secret is patched which contains all the data. For any assistance, contact My Oracle Support .

3.3.98 SCPWorkerAndNFTimeSyncFailure

Table 3-102 SCPWorkerAndNFTimeSyncFailure

Field	Details
Description	'Consumer NF and SCP are not time synchronized'
Summary	Consumer NF and SCP are not time synchronized: consumernfinstanceid = {{ \$labels.ocscp_consumer_nf_type }}, consumernfinstanceid = {{ \$labels.ocscp_consumer_nf_instance_id }}, namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}
Severity	Major
Condition	This alert is raised when there is a time synchronization difference between SCP-Worker and the consumer NF received request with the timestamp_headers_support feature is enabled. increase(ocscp_worker_timestamp_headers_validation_fail_total{ocscp_validation_failure="TIME_SYNC_FAILURE"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15019
Metric Used	ocscp_worker_timestamp_headers_validation_fail_total
Recommended Action	Cause: When there is a time synchronization difference between SCP-Worker and the consumer NF received request with the timestamp_headers_support feature is enabled. Diagnostic Information: SCP may have received a request from the consumer NF which is not in time sync when the timestamp_headers_support is enabled. Recovery: Alert is cleared automatically when no such read failure occurs over the next scrape interval. For any assistance, contact My Oracle Support .

3.3.99 SCPCoherenceCacheHitsRateAboveThreshold

Table 3-103 SCPCoherenceCacheHitsRateAboveThreshold

Field	Details
Description	'SCP Coherence Cache Hits Rate above threshold (i.e. 220000 mps)'
Summary	'namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}: Current Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 220000'
Severity	Critical
Condition	This alert is raised when SCP Coherence cache rate exceeds the threshold (220000 MPS). sum(rate(coherence_cache_hits{app_kubernetes_io_name="scp-cache"}[2m])) by (kubernetes_namespace) > 220000
OID	1.3.6.1.4.1.323.5.3.35.1.2.7034
Metric Used	coherence_cache_hits
Recommended Action	Cause: When SCP Coherence cache rate exceeds the threshold (220000 MPS). Diagnostic Information: SCP Coherence cache rate exceeds the threshold (220000 MPS). Recovery: Alert is cleared automatically when SCP Coherence cache rate goes below the threshold during the next scrape interval. For any assistance, contact My Oracle Support .

3.3.100 SCPNotificationRequestsRateAboveThreshold

Table 3-104 SCPNotificationRequestsRateAboveThreshold

Field	Details
Description	'SCP Notification Requests are above threshold (i.e. 40 mps)'
Summary	'namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}: Current Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 40'
Severity	Critical
Condition	This alert is raised when the number of SCP notification requests exceeds the threshold (40 MPS). sum(rate(ocscp_notification_request_total{app_kubernetes_io_name="scpc-notification"}[2m])) by (kubernetes_namespace) > 40
OID	1.3.6.1.4.1.323.5.3.35.1.2.3012
Metric Used	ocscp_notification_request_total
Recommended Action	Cause: When the number of SCP notification requests exceeds the threshold (40 MPS). Diagnostic Information: SCP notification requests exceed the threshold (40 MPS). Recovery: Alert is cleared automatically when the number of SCP notification requests goes below the threshold during the next scrape interval. For any assistance, contact My Oracle Support .

3.3.101 SCPLciRxRateAboveThreshold

Table 3-105 SCPLciRxRateAboveThreshold

Field	Details
Description	description: 'LCI Rx Rate is above threshold (i.e. 150)'
Summary	'namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}: Current Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 150'
Severity	Critical
Condition	This alert is raised when LCI Rx rate exceeds the threshold, that is, 150. sum(rate(ocscp_metric_nf_lci_rx_total{app_kubernetes_io_name="scp-worker"}[2m])) by (kubernetes_namespace) > 150
OID	1.3.6.1.4.1.323.5.3.35.1.2.7032
Metric Used	ocscp_metric_nf_lci_rx_total
Recommended Action	Cause: When LCI Rx rate exceeds the threshold, that is, 150. Diagnostic Information: LCI Rx rate exceeds the threshold, that is, 150. Recovery: Alert is cleared automatically when LCI Rx rate goes below the threshold during the next scrape interval. For any assistance, contact My Oracle Support .

3.3.102 SCPOciRxRateAboveThreshold

Table 3-106 SCPOciRxRateAboveThreshold

Field	Details
Description	'OCI Rx Rate is above threshold (i.e. 5)'
Summary	'namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}: Current Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 5'
Severity	Critical
Condition	This alert is raised when OCI Rx rate exceeds the threshold, that is, 5. sum(rate(ocscp_metric_nf_oci_rx_total{app_kubernetes_io_name="scp-worker"}[2m])) by (kubernetes_namespace) > 5
OID	1.3.6.1.4.1.323.5.3.35.1.2.7033
Metric Used	ocscp_metric_nf_oci_rx_total
Recommended Action	Cause: When OCI Rx rate exceeds the threshold, that is, 5. Diagnostic Information: OCI Rx rate exceeds the threshold, that is, 5. Recovery: Alert is cleared automatically when OCI Rx rate goes below the threshold during the next scrape interval. For any assistance, contact My Oracle Support .

3.3.103 SCPTrafficRateAboveRatedThreshold

Table 3-107 SCPTrafficRateAboveRatedThreshold

Field	Details
Description	'Traffic Rate at SCP is above threshold (i.e. 730000 mps)'
Summary	'namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}: Current Traffic Rate is {{ \$value printf "%.2f" }} mps which is above 730000'
Severity	Critical
Condition	This alert is raised when the traffic rate at SCP exceeds the threshold, that is, 730000 MPS. (sum(rate(ocscp_metric_http_rx_req_total{app_kubernetes_io_name="scp-worker", ocscp_consumer_host!="OCSCP-INITIATED-MSG"}[2m])) by (kubernetes_namespace) + sum(rate(ocscp_metric_http_tx_req_total{app_kubernetes_io_name="scp-worker"}[2m])) by (kubernetes_namespace)) > 730000
OID	1.3.6.1.4.1.323.5.3.35.1.2.7031
Metric Used	ocscp_metric_http_rx_req_total
Recommended Action	Cause: When the traffic rate at SCP exceeds the threshold, that is, 730000 MPS. Diagnostic Information: The traffic rate at SCP exceeds the threshold, that is, 730000 MPS. Recovery: Alert is cleared automatically when the traffic rate goes below the threshold during the next scrape interval. For any assistance, contact My Oracle Support .

3.3.104 SCPMaxNFProfilesLimitExceeded

Table 3-108 SCPMaxNFProfilesLimitExceeded

Field	Details
Description	'NFs registered are more than the rated capacity (1200 NFs)'
Summary	'NFs registered are more than the rated capacity (1200 NFs) for namespace: {{ \$labels.kubernetes_namespace }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when the number of registered NFs exceeds the rated capacity (1200 NFs). sum(ocscp_notification_nf_profiles_count) by (kubernetes_namespace) > 1200
OID	1.3.6.1.4.1.323.5.3.35.1.2.3013
Metric Used	ocscp_notification_nf_profiles_count
Recommended Action	Cause: When registered NFs are more than the rated capacity (1200 NFs). Diagnostic Information: NFs registered are more than the rated capacity (1200 NFs). Recovery: The alert is cleared automatically when the number of registered NFs falls below the rated capacity (1200 NFs) during the next scrape interval. For any assistance, contact My Oracle Support .

3.3.105 SCPCustomRoamingProxyRoutingNotFound

Table 3-109 SCPCustomRoamingProxyRoutingNotFound

Field	Details
Description	'Custom Roaming Proxy is not configured for PLMN'
Summary	'Custom Roaming Proxy is not configured for PLMN. plmnid: {{\$labels.plmn_id}}, custom_roaming_proxy_nfType: {{\$labels.ocscp_custom_roaming_proxy_nf_type}}, scpfqdn: {{\$labels.scp_fqdn}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod_name}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when Custom Roaming Proxy is not configured for the PLMN ID. ocscp_worker_custom_roaming_proxy_not_found_current == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.15042
Metric Used	ocscp_worker_custom_roaming_proxy_not_found_current
Recommended Action	Cause: Custom Roaming Proxy routing related rules are not configured at SCP for selected PLMN in the inter-PLMN routing. Diagnostic Information: - Custom Roaming Proxy profile is not registered or Custom Roaming Proxy related configuration is not done at SCP. - Verify if the routing rules are created for selected inter-PLMN in the PLMN_ROAMING_PROXY_MAPPING table. Recovery: This alert is cleared automatically when the Custom Roaming Proxy related routing rules are created at SCP for the selected PLMN in the inter-PLMN routing. For any assistance, contact My Oracle Support.

3.3.106 SCPPeerCustomRoamingProxyUnhealthy

Table 3-110 SCPPeerCustomRoamingProxyUnhealthy

Field	Details
Description	Peer Custom Roaming Proxy is marked unhealthy
Summary	'Peer Custom Roaming Proxy is marked unhealthy. custom_roaming_proxy_fqdn: {{\$labels.ocscp_custom_roaming_proxy_fqdn}}, custom_roaming_proxy_nfType: {{\$labels.ocscp_custom_roaming_proxy_nf_type}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when peer Custom Roaming Proxy is marked unhealthy. ocscp_worker_custom_roaming_proxy_unhealthy_current == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.15043
Metric Used	ocscp_worker_custom_roaming_proxy_unhealthy_current

Table 3-110 (Cont.) SCPPeerCustomRoamingProxyUnhealthy

Field	Details
Recommended Action	Cause: The peer Custom Roaming Proxy is marked unhealthy because of consecutive failure responses. Diagnostic Information: • Check transport failures and routing errors on peer SCP. • Use ping command (ping <IPAddress>) from primary or secondary nodes using IP address of the service. • If FQDN of the service is used, run the ping command inside the pod. • If the pod does not support ping command, use the Debug container of SCP pod. • Alternatively, collect tcpdump for ensuring proper connectivity. <pre>tcpdump -w capture.pcap -i <pod interface></pre> Recovery: This alert is automatically after the degradation time is over. Degradation time = number of consecutive degradations * configured base ejection. For any assistance, contact My Oracle Support.

3.3.107

SCPCustomRoamingProxyMaxViaHeaderProcessingCountExceeded

Table 3-111 SCPCustomRoamingProxyMaxViaHeaderProcessingCountExceeded

Field	Details
Description	'Proxy information count in Via header(s) exceeded the maximum via header proxy processing Limit.'
Summary	'Proxy information count in Via header(s) exceeded the maximum via header proxy processing Limit. max_via_header_processing_limit: {{\$labels.max_via_header_processing_limit}}, namespace: {{\$labels.labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when the Via header processing count for Custom Roaming Proxy exceeds the configured limit. <pre>increase(ocscp_worker_custom_roaming_proxy_via_header_limit_exceeded_total[2m]) > 0</pre>
OID	1.3.6.1.4.1.323.5.3.35.1.2.15044
Metric Used	increase(ocscp_worker_custom_roaming_proxy_via_header_limit_exceeded_total)
Recommended Action	Cause: When the Custom Roaming Proxy maximum Via header processing count exceeded the configured limit. Diagnostic Information: Check if the number of Via headers in ingress request is not exceeding the configured limit in configuration for Custom Roaming Proxy. Recovery: This alert is cleared automatically when the request has less Via header than the configured limit per request for Custom Roaming Proxy. For any assistance, contact My Oracle Support.

3.3.108 SCPCustomRoamingProxyRoutingFailed

Table 3-112 SCPCustomRoamingProxyRoutingFailed

Field	Details
Description	'Sending message to Custom Roaming Proxy Failed'
Summary	'Sending message to Custom Roaming Proxy Failed. custom_roaming_proxy_fqdn: {{ \$labels.ocscp_custom_roaming_proxy_fqdn }}, custom_roaming_proxy_nfType: {{ \$labels.ocscp_custom_roaming_proxy_nf_type }}, scpfdn: {{ \$labels.scp_fqdn }}, namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod_name }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when none of the Custom Roaming Proxy NFs in the provided set are reachable. ocscp_worker_custom_roaming_proxy_routing_attempt_fail_current == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.15041
Metric Used	ocscp_worker_custom_roaming_proxy_routing_attempt_fail_current
Recommended Action	Cause: When none of the Custom Roaming Proxy NFs in the provided set are reachable. Diagnostic Information: • Check failure reasons. • Check if Custom Roaming Proxy is configured in the provided set region. • Check if Custom Roaming Proxy is reachable and available. Recovery: This alarm is cleared automatically when the connection is established with the Custom Roaming Proxy NF and the routing is successful. For any assistance, contact My Oracle Support.

3.3.109 SCPWorkerNFSelectionHeaderMatchListDuplicates

Table 3-113 SCPWorkerNFSelectionHeaderMatchListDuplicates

Field	Details
Description	'Duplicate Header Match list in Nf SelectionMode Rule'
Summary	'Duplicate header match list present for nf selection at config: serviceName = {{ \$labels.ocscp_nf_service_type }}, namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when the selected NF SelectionMode rule identifies more than one matching item in the header match list configuration for a request. increase(ocscp_worker_nf_selection_mode_header_match_conflict_total[10m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15029
Metric Used	ocscp_worker_nf_selection_mode_header_match_conflict_total

Table 3-113 (Cont.) SCPWorkerNFSelectionHeaderMatchListDuplicates

Field	Details
Recommended Action	Cause: This alert indicates a misconfiguration in NF selection mode where a single rule contains multiple HeaderMatchList entries that can independently match the same service request. This overlapping header match criteria can result in ambiguity when selecting the appropriate NF selection mode. Example: A rule (rule1) contains the following configurations: Entry A: headerMatchList: [{ "headerName": "header1", "headerStatus": "PRESENT" }], selectionMode: CONSUMER_NF_SELECTION Entry B: headerMatchList: [{ "headerName": "header2", "headerStatus": "PRESENT" }], selectionMode: ENHANCED_CONSUMER_NF_SELECTION If an incoming service request contains both header1 and header2, then: - both Entry A and Entry B become valid candidates. - any one of the entries of NF Selection Mode configuration can be applied. - this results in multiple matching indices, triggering the alert. Diagnostic Information: You must review the NF selection mode configuration for the reported rule where each entry can independently match the same request headers. Recovery: This alert is automatically cleared after 10 minutes. For any assistance, contact My Oracle Support.

3.3.110 ScpWorkerServiceRequestRoutingInfoAbsent

Table 3-114 ScpWorkerServiceRequestRoutingInfoAbsent

Field	Details
Description	'Service request is received without any routing information'
Summary	'Service request is received without any routing information for nfType {{ \$labels.ocscp_nf_type }}, nfServiceType {{ \$labels.ocscp_nf_service_type }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Info
Condition	This alert is raised when an SBI service request is received without routing information. Note: This alert is raised independently of feature parameters and is generated regardless of whether any feature is enabled. increase(ocscp_worker_service_req_routing_information_absent_total[10m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.3014
Metric Used	ocscp_worker_service_req_routing_information_absent_total
Recommended Action	Cause: SCP received SBI service requests without routing information regardless of whether any feature is enabled. Diagnostic Information: NA Recovery: This alert is automatically cleared after 10 minutes. For any assistance, contact My Oracle Support.

3.3.111 ScpWorkerNotificationRequestRoutingInfoAbsent

Table 3-115 ScpWorkerNotificationRequestRoutingInfoAbsent

Field	Details
Description	'Notification request is received without any routing information'
Summary	'Notification request is received without any routing information for nfType {{ \$labels.ocscp_nf_type }}, nfServiceType {{ \$labels.ocscp_nf_service_type }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when a notification request is received without routing information. Note: This alert is raised independently of feature parameters and is generated regardless of whether any feature is enabled. increase(ocscp_worker_notification_req_routing_information_absent_total[10m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.3015
Metric Used	ocscp_worker_notification_req_routing_information_absent_total
Recommended Action	Cause: SCP received Notification request without routing information regardless of whether any feature is enabled. Diagnostic Information: Not applicable. Recovery: This alert is automatically cleared after 10 minutes. For any assistance, contact My Oracle Support .

3.3.112 SCPWorkerNFProfileSyncFailure

Table 3-116 SCPWorkerNFProfileSyncFailure

Field	Details
Description	'NF Profile processinf failed at worker'
Summary	'NF Profile processinf failed at worker: nfType = {{ \$labels.ocscp_nf_type }}, producerNfInstanceId = {{ \$labels.ocscp_producer_nf_instance_id }}, namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpFqdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Info
Condition	This alert is raised when the SCP-Worker microservice is unable to process the NF profile update received from the SCP-Notification microservice. ocscp_worker_nf_profile_sync_failure_status > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15039
Metric Used	ocscp_worker_nf_profile_sync_failure_status
Recommended Action	Cause: SCP-Worker is unable to process the NF profile update received from SCP-Notification. Diagnostic Information: NA Recovery: The alert is automatically cleared when a subsequent NF profile update with the same nfInstanceId is processed by the SCP-Worker microservice. For any assistance, contact My Oracle Support .

3.3.113 SCPNfTopologyForceRefreshAuditFailure

Table 3-117 SCPNfTopologyForceRefreshAuditFailure

Field	Details
Description	'SCP Force Refresh Audit failed for nfType {{\$labels.nfType}}'
Summary	'SCP Force Refresh Audit failed for nfType {{\$labels.nfType}}, nrfRegionOrSetId: {{\$labels.nrfRegionOrSetId}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when a force refresh audit attempt in the SCP-Audit microservice fails for a specific nfType and nrfRegionOrSetId. ocscp_audit_nf_topology_force_refresh_audit_failure > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15045
Metric Used	ocscp_audit_nf_topology_force_refresh_audit_failure
Recommended Action	Cause: When a force refresh audit attempt in the SCP-Audit microservice fails. Diagnostic Information: NA Recovery: Perform the following steps: - The alert is cleared when you re-run the force refresh using the NF Topology On-Demand Audit REST API and it succeeds for the affected nfType and nrfRegionOrSetId. - Changing the Topology Source Info (TSI) to LOCAL for a given nfType clears all alert instances associated with that nfType. - Removing an nfType from SCP using the NFTypes-NFService REST API clears all alert instances associated with that nfType. - Removing an nrfRegionOrSetId clears all the alert instances associated with that nrfRegionOrSetId for the relevant nfType. For any assistance, contact My Oracle Support.

3.3.114 SCPNfGroupIdDiscoveryFailure

Table 3-118 SCPNfGroupIdDiscoveryFailure

Field	Details
Description	'SCP failed to discover SLF (UDR) for GroupId resolution'
Summary	'SCP NF GroupId discovery failed: nfType = {{\$labels.ocscp_nf_type}}, nfservicetype = {{\$labels.ocscp_nf_service_type}}, config = {{\$labels.config_name}}, namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when SCP failed to discover UDR for Group ID resolution. increase(ocscp_worker_nfgroupid_resolution_failure_total{discovery_complete="false"}[2m]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15046
Metric Used	ocscp_worker_nfgroupid_resolution_failure_total

Table 3-118 (Cont.) SCPNfGroupIdDiscoveryFailure

Field	Details
Recommended Action	Cause: When SCP fails to discover UDR for Group ID resolution, no UDR profile with routing functionality is resolved. Diagnostic Information: • Verify that NF Group ID retrieval configuration and rule mapping are configured as required. • Verify that all mandatory discovery parameters defined in the NF Group ID retrieval configuration are present in the incoming requests. • Verify the connection to SCP-NrfProxy microservice and confirm that discovery requests can reach NRF. • Verify that NRF is healthy and able to return discovery responses for the configured UDR discovery parameters. • If the discovery mode is set to SCP for UDR, verify that the correct static rules are configured for the specified UDR discovery parameters. Recovery: • Resolve the configuration or dependency issue identified in diagnostics, and then retest with a matching request. • This alert is automatically cleared when the SCP observes no new failures for a continuous period of 2 minutes. • Previously failed requests are not retried automatically; only the occurrence of new failures ceases. For any assistance, contact My Oracle Support.

3.3.115 SCPAuditDurationExceeded

Table 3-119 SCPAuditDurationExceeded

Field	Details
Description	'Scheduled audit duration exceeded configured interval for nrfRegionOrSetId: {{\$labels.preregistration}}'
Summary	'Audit duration exceeded. nrfRegionOrSetId: {{\$labels.nrfRegionOrSetId}}, namespace: {{\$labels.kubernetes_namespace}}, podname: {{\$labels.kubernetes_pod_name}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when the scheduled audit duration exceeds the configured interval. ocscp_audit_duration_exceeded == 1
OID	1.3.6.1.4.1.323.5.3.35.1.2.4003
Metric Used	ocscp_audit_duration_exceeded

Table 3-119 (Cont.) SCPAuditDurationExceeded

Field	Details
Recommended Action	Cause: Audit duration exceeded the configured audit interval for a specific NRFSetId. Diagnostic Information: - Verify whether audit for the NRFSetId takes longer than the configured audit interval. - Determine whether the increased audit duration is expected. - If the increase is unexpected, verify that NRF responses arrive within the expected time. - If the increase is expected, increase the configured audit interval value. Recovery: - The alert clears in the next audit interval when audit duration takes less time than the configured audit interval. For any assistance, contact My Oracle Support.

3.3.116 SCPSelfProfileUpdateSuccess

Table 3-120 SCPSelfProfileUpdateSuccess

Field	Details
Description	'Self SCP profile runtime update completed successfully'
Summary	'Self SCP profile runtime update completed successfully namespace: {{ \$labels.kubernetes_namespace }}, podname: {{ \$labels.kubernetes_pod_name }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Info
Condition	This alert is raised when the self SCP profile runtime update is completed. increase(ocscp_self_profile_update_success_total[24h]) > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.15047
Metric Used	ocscp_self_profile_update_success_total
Recommended Action	Cause: When the self SCP profile runtime update is completed successfully. This alert is raised when the following parameters in the self SCP profile are updated: priority, capacity, load, plmnList, locality, servingScope, mateScplInfoList, or mateSiteInfo. Diagnostic Information: Verify the namespace, pod name, and SCP FQDN in the summary to confirm the expected instance update. Recovery: No action is required as this is an informational alert. This alert clears automatically if no successful self-profile update is observed in the last 24 hours. For any assistance, contact My Oracle Support.

3.3.117 SCPEgressRateLimitActiveKeyCountAboveCriticalThreshold

Table 3-121 SCPEgressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Description	'Egress Rate Limit active key count excluding Global Egress Rate Limit keys is above 100 Percent of supported limit (3500)'

Table 3-121 (Cont.) SCPEgressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Summary	'Egress Rate Limit active key count excluding Global Egress Rate Limit keys is {{ \$value printf "%.0f" }}, above 100 Percent of supported limit (3500). namespace: {{ \$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }{'
Severity	Critical
Condition	This alert is raised when Egress Rate Limit active key count, excluding Global egress Rate Limit keys, exceeds 100% of the supported limit. (sum(coherence_cache_size{name="egressRateLimitMap"}) by (kubernetes_namespace) - sum(coherence_cache_size{name="localRLConfigDataCache"}) by (kubernetes_namespace)) > 3500
OID	1.3.6.1.4.1.323.5.3.35.1.2.7036
Metric Used	coherence_cache_size
Recommended Action	Cause: Egress rate limit active key count, excluding Global Egress Rate Limit keys, increased above the supported limit, 3500. Diagnostic Information: - Wait 30 minutes for inactive key deletion. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name egressRateLimitMap, and subtract coherence_cache_size for cache name localRLConfigDataCache to identify the number of active egress rate limit bucket keys. - Check ocsdp_metric_message_current_request_rate_total grouped by bucketKey to identify bucket keys receiving ongoing traffic. - If the active bucket key count is high, review the egress rate limiter configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce egress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the critical threshold. For any assistance, contact My Oracle Support.

3.3.118 SCPEgressRateLimitActiveKeyCountAboveMajorThreshold

Table 3-122 SCPEgressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Description	'Egress Rate Limit active key count excluding Global Egress Rate Limit keys is above 90 Percent of supported limit (3500)'
Summary	'Egress Rate Limit active key count excluding Global Egress Rate Limit keys is {{ \$value printf "%.0f" }}, above 90 Percent of supported limit (3500). namespace: {{ \$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }{'
Severity	Major

Table 3-122 (Cont.) SCPEgressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Condition	This alert is raised when the Egress Rate Limit active key count, excluding Global Egress Rate Limit keys, exceeds 90% of the supported limit. <pre>(sum(coherence_cache_size{name="egressRateLimitMap"}) by (kubernetes_namespace) - sum(coherence_cache_size{name="localRLConfigDataCache"}) by (kubernetes_namespace)) > 3150 <= 3500</pre>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7036
Metric Used	coherence_cache_size
Recommended Action	Cause: Egress rate limit active key count, excluding Global Egress Rate Limit keys, increased above 90% of the supported limit, 3500. Diagnostic Information: - Wait 30 minutes for inactive key deletion to complete. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name egressRateLimitMap, and subtract coherence_cache_size for cache name localRLConfigDataCache to identify the number of active egress rate limit bucket keys. - Check ocsdp_metric_message_current_request_rate_total grouped by bucketKey to identify bucket keys receiving ongoing traffic. - If the active bucket key count is high, review the egress rate limiter configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce egress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the major threshold. For any assistance, contact My Oracle Support.

3.3.119

SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold

Table 3-123 SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Description	'Combined Egress and Global Egress Rate Limit active key count is above 100 Percent of supported limit (3500)'
Summary	'Combined Egress and Global Egress Rate Limit active key count is {{ \$value printf "%.0f" }}, above 100 Percent of supported limit (3500). namespace: {{ \$labels.kubernetes_namespace }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical

Table 3-123 (Cont.) SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Condition	This alert is raised when combined Egress and Global Egress Rate Limit active key count exceeds 100% of the supported limit. <pre>sum(coherence_cache_size{name="egressRateLimitMap"}) by (kubernetes_namespace) > 3500</pre>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7038
Metric Used	coherence_cache_size
Recommended Action	Cause: Combined Egress and Global Egress Rate Limit active key count increased above the supported limit, 3500. Diagnostic Information: - Wait 30 minutes for inactive key deletion. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name egressRateLimitMap to identify the combined number of active Egress and Global Egress Rate Limit bucket keys that SCP uses. - Check ocsdp_metric_message_current_request_rate_total grouped by bucketKey to identify bucket keys receiving ongoing traffic. - If the combined active bucket key count is high, review Egress and Global Egress Rate Limit configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce egress and global egress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the critical threshold. For any assistance, contact My Oracle Support.

3.3.120

SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold

Table 3-124 SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Description	'Combined Egress and Global Egress Rate Limit active key count is above 90 Percent of supported limit (3500)'
Summary	'Combined Egress and Global Egress Rate Limit active key count is {{ \$value printf "%.0f" }}, above 90 Percent of supported limit (3500). namespace: {{ \$labels.kubernetes_namespace }}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major

Table 3-124 (Cont.) SCPEgressAndGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Condition	This alert is raised when combined Egress and Global Egress Rate Limit active key count exceeds 90% of the supported limit. <pre>sum(coherence_cache_size{name="egressRateLimitMap"}) by (kubernetes_namespace) > 3150 <= 3500</pre>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7038
Metric Used	coherence_cache_size
Recommended Action	Cause: Combined Egress and Global Egress Rate Limit active key count increased above 90 percent of the supported limit, 3500. Diagnostic Information: - Wait 30 minutes for inactive key deletion. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name egressRateLimitMap to identify the combined number of active Egress and Global Egress Rate Limit bucket keys SCP uses. - Check ocsdp_metric_message_current_request_rate_total grouped by bucketKey to identify bucket keys receiving ongoing traffic. - If the combined active bucket key count is high, review the Egress and Global Egress Rate Limit configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce egress and global egress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the major threshold. For any assistance, contact My Oracle Support.

3.3.121 SCPIngressRateLimitActiveKeyCountAboveMajorThreshold

Table 3-125 SCPIngressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Description	'Ingress Rate Limit active key count for cache {{\$labels.name}} is above 90 Percent of supported limit (900)'
Summary	'Ingress Rate Limit active key count for cache {{\$labels.name}} is {{ \$value printf "%.0f" }}, above 90 Percent of supported limit (900). namespace: {{\$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when the Ingress Rate Limit active key count exceeds 90% of the supported limit. <pre>sum(coherence_cache_size{name="ingressRateLimitMap"}) by (kubernetes_namespace,name) > 810 <= 900</pre>
OID	1.3.6.1.4.1.323.5.3.35.1.2.7035

Table 3-125 (Cont.) SCPIngressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Metric Used	coherence_cache_size
Recommended Action	Cause: Ingress rate limit active key count has increased above 90% of the supported limit, 900. Diagnostic Information: - Wait 30 minutes for inactive key deletion to complete. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name ingressRateLimitMap to identify the number of active ingress rate limit bucket keys SCP uses. - Check ocsdp_metric_message_current_request_rate_total grouped by bucketKey to identify bucket keys receiving ongoing traffic. - If the active bucket key count is high, review the ingress rate limiter configuration and traffic patterns to identify configurations that generate excessive bucket keys. Recovery: - Reduce ingress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count falls below the major threshold. For any assistance, contact My Oracle Support.

3.3.122 SCPIngressRateLimitActiveKeyCountAboveCriticalThreshold

Table 3-126 SCPIngressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Description	'Ingress Rate Limit active key count for cache {{\$labels.name}} is above 100 Percent of supported limit (900)'
Summary	'Ingress Rate Limit active key count for cache {{\$labels.name}} is {{ \$value printf "%.0f" }}, above 100 Percent of supported limit (900). namespace: {{\$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the Ingress Rate Limit active key count exceeds 100% of the supported limit. sum(coherence_cache_size{name="ingressRateLimitMap"}) by (kubernetes_namespace, name) > 900
OID	1.3.6.1.4.1.323.5.3.35.1.2.7035
Metric Used	coherence_cache_size

Table 3-126 (Cont.) SCPIngressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Recommended Action	Cause: Ingress rate limit active key count increased above the supported limit, 900. Diagnostic Information: - Wait 30 minutes for inactive key deletion. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name ingressRateLimitMap to identify the number of active ingress rate limit bucket keys SCP uses. - Check ocscp_metric_message_current_request_rate_total grouped by bucketKey to identify bucket keys receiving ongoing traffic. - If the active bucket key count is high, review the ingress rate limiter configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce ingress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the critical threshold. For any assistance, contact My Oracle Support.

3.3.123 SCPGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold

Table 3-127 SCPGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Description	'Global Egress Rate Limit active key count for cache {{\$labels.name}} is above 100 Percent of supported limit (2400)'
Summary	'Global Egress Rate Limit active key count for cache {{\$labels.name}} is {{ \$value printf "%.0f" }}, above 100 Percent of supported limit (2400). namespace: {{\$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the Global Egress Rate Limit active key count exceeds 100% of the supported limit. sum(coherence_cache_size{name="localRLConfigDataCache"}) by (kubernetes_namespace, name) > 2400
OID	1.3.6.1.4.1.323.5.3.35.1.2.7037
Metric Used	coherence_cache_size

Table 3-127 (Cont.) SCPGlobalEgressRateLimitActiveKeyCountAboveCriticalThreshold

Field	Details
Recommended Action	Cause: Global Egress Rate Limit active key count increased above the supported limit, 2400. Diagnostic Information: - Wait 30 minutes for inactive key deletion. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name localRLConfigDataCache to identify the number of active Global Egress Rate Limit bucket keys SCP uses. - Check ocscp_metric_message_current_request_rate_total grouped by bucketKey to identify Global Egress Rate Limit bucket keys receiving ongoing traffic. - If the active bucket key count is high, review the Global Egress Rate Limit configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce global egress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the critical threshold. For any assistance, contact My Oracle Support.

3.3.124 SCPGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold

Table 3-128 SCPGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Description	'Global Egress Rate Limit active key count for cache {{\$labels.name}} is above 90 Percent of supported limit (2400)'
Summary	'Global Egress Rate Limit active key count for cache {{\$labels.name}} is {{ \$value printf "%.0f" }}, above 90 Percent of supported limit (2400). namespace: {{\$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alerts is raised when the Global Egress Rate Limit active key count exceeds 90% of the supported limit. sum(coherence_cache_size{name="localRLConfigDataCache"}) by (kubernetes_namespace, name) > 2160 <= 2400
OID	1.3.6.1.4.1.323.5.3.35.1.2.7037
Metric Used	coherence_cache_size

Table 3-128 (Cont.) SCPGlobalEgressRateLimitActiveKeyCountAboveMajorThreshold

Field	Details
Recommended Action	Cause: Global Egress Rate Limit active key count increased above 90% of the supported limit, 2400. Diagnostic Information: - Wait 30 minutes for inactive keys to be deleted. If the alert remains active, perform the following steps: - Check the coherence_cache_size metric for cache name localRLConfigDataCache to identify the number of active Global Egress Rate Limit bucket keys used by SCP. - Check ocscp_metric_message_current_request_rate_total grouped by bucketKey to identify Global Egress Rate Limit bucket keys receiving ongoing traffic. - If the active bucket key count is high, review Global Egress Rate Limit configuration and traffic patterns to identify configurations generating excessive bucket keys. Recovery: - Reduce global egress rate limiter key configuration entries, or disable rate limiting for keys that do not require it. - The alert clears automatically when the active key count goes below the major threshold. For any assistance, contact My Oracle Support.

3.3.125 SCPNotificationDBCertificateExpiryAboveMinorThreshold

Table 3-129 SCPNotificationDBCertificateExpiryAboveMinorThreshold

Field	Details
Description	'Notification DB client certificate expiration has reached the minor threshold'
Summary	'Notification DB client certificate expiration is at minor level. namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Minor
Condition	This alert is raised when the client certificate goes beyond minor threshold.ocscp_notification_ssl_certificate_expire{expireLevel="minor"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.3016
Metric Used	ocscp_notification_ssl_certificate_expire

Table 3-129 (Cont.) SCPNotificationDBCertificateExpiryAboveMinorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured minor threshold. Diagnostic Information: - Check the alert labels to identify the affected certificate store. The certStoreType label indicates whether the certificate is from primaryKeystore, primaryTruststore, secondaryKeystore, or secondaryTruststore. - Verify the certificate expiry date in the DB TLS Kubernetes secret configured under global.database.ssl. The global.sslCertExpiryMinorThreshold parameter configures the minor threshold. Recovery: - Plan the certificate renewal before the certificate reaches the major or critical threshold. - Renew or replace the affected DB TLS certificate and update the corresponding Kubernetes secret. - The alert clears automatically when certificate validity is above the configured minor threshold. For any assistance, contact My Oracle Support.

3.3.126 SCPNotificationDBCertificateExpiryAboveCriticalThreshold

Table 3-130 SCPNotificationDBCertificateExpiryAboveCriticalThreshold

Field	Details
Description	'Notification DB client certificate expiration has reached the critical threshold'
Summary	'Notification DB client certificate expiration is at critical level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the client certificate goes beyond critical threshold.ocscp_notification_ssl_certificate_expire{expireLevel="critical"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.3016
Metric Used	ocscp_notification_ssl_certificate_expire

Table 3-130 (Cont.) SCPNotificationDBCertificateExpiryAboveCriticalThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is near expiry or has expired and has reached the configured critical threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under <code>global.database.ssl</code>. - Check the affected microservice logs for certificate switch, reload, or DB TLS connection issues. The <code>global.sslCertExpiryCriticalThreshold</code> parameter configures the critical threshold. Recovery: - Renew or replace the affected DB TLS certificate immediately and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. - If the secondary DB TLS certificates are configured, verify whether traffic has switched to the secondary certificate path. - The alert clears automatically when certificate validity is above the configured critical threshold. For any assistance, contact My Oracle Support.

3.3.127 SCPNotificationKeystoreExpireOutOfTolerance

Table 3-131 SCPNotificationKeystoreExpireOutOfTolerance

Field	Details
Description	'Notification DB keystore expiration is outside the configured tolerance'
Summary	'Notification DB keystore expiration is outside tolerance. namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the keystore is out of tolerance threshold for DB connection. <code>ocscp_notification_keystore_expire_out_of_tolerance > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.3017
Metric Used	ocscp_notification_keystore_expire_out_of_tolerance

Table 3-131 (Cont.) SCPNotificationKeystoreExpireOutOfTolerance

Field	Details
Recommended Action	Cause: The DB TLS keystore or truststore certificate for the SCP microservice is expired or will expire before the configured minimum certificate validity tolerance. Diagnostic Information: - Check the affected microservice pod logs for keystore or truststore validity warnings. - Verify the DB TLS certificate validity for the primary and secondary keystore/truststore configured under <code>global.database.ssl</code>. - Check the ordinality and storeType labels on the alert metric to identify whether the primary or secondary keystore/truststore is outside tolerance. The <code>global.database.ssl.minCertValidity</code> parameter configures the tolerance. Recovery: - Renew or replace the DB TLS certificate and update the corresponding Kubernetes secret for the affected primary or secondary keystore/truststore. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. - The alert clears automatically when the keystore and truststore certificates are valid beyond the configured minimum certificate validity tolerance. For any assistance, contact My Oracle Support.

3.3.128 SCPNotificationDBCertificateExpiryAboveMajorThreshold

Table 3-132 SCPNotificationDBCertificateExpiryAboveMajorThreshold

Field	Details
Description	'Notification DB client certificate expiration has reached the major threshold'
Summary	'Notification DB client certificate expiration is at major level. namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when the client certificate goes beyond major threshold. <code>ocscp_notification_ssl_certificate_expire_{expireLevel="major"} > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.3016
Metric Used	ocscp_notification_ssl_certificate_expire

Table 3-132 (Cont.) SCPNotificationDBCertificateExpiryAboveMajorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured major threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. The global.sslCertExpiryMajorThreshold parameter configures the major threshold. Recovery: - Renew or replace the affected DB TLS certificate as soon as possible and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod, if required, so that it uses the updated certificate. - The alert clears automatically when certificate validity is above the configured major threshold. For any assistance, contact My Oracle Support.

3.3.129 SCPSubscriptionDBCertificateExpiryAboveCriticalThreshold

Table 3-133 SCPSubscriptionDBCertificateExpiryAboveCriticalThreshold

Field	Details
Description	'Subscription DB client certificate expiration has reached the critical threshold'
Summary	'Subscription DB client certificate expiration is at critical level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the client certificate goes beyond critical threshold.ocscp_subscription_ssl_certificate_expire{expireLevel="critical"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.2006
Metric Used	ocscp_subscription_ssl_certificate_expire

Table 3-133 (Cont.) SCPSubscriptionDBCertificateExpiryAboveCriticalThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is near expiry or has expired and has reached the configured critical threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. - Check the affected microservice logs for certificate switch, reload, or DB TLS connection issues. The global.sslCertExpiryCriticalThreshold parameter configures the critical threshold. Recovery: - Renew or replace the affected DB TLS certificate immediately and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. If secondary DB TLS certificates are configured, verify whether traffic has switched to the secondary certificate path. - The alert clears automatically when certificate validity is above the configured critical threshold. For any assistance, contact My Oracle Support.

3.3.130 SCPSubscriptionDBCertificateExpiryAboveMajorThreshold

Table 3-134 SCPSubscriptionDBCertificateExpiryAboveMajorThreshold

Field	Details
Description	'Subscription DB client certificate expiration has reached the major threshold'
Summary	'Subscription DB client certificate expiration is at major level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Major
Condition	This alert is raised when the client certificate goes beyond major threshold.ocscp_subscription_ssl_certificate_expire{expireLevel="major"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.2006
Metric Used	ocscp_subscription_ssl_certificate_expire

Table 3-134 (Cont.) SCPSubscriptionDBCertificateExpiryAboveMajorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured major threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. The global.sslCertExpiryMajorThreshold parameter configures the major threshold. Recovery: - Renew or replace the affected DB TLS certificate as soon as possible and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod, if required, so that it uses the updated certificate. - The alert clears automatically when certificate validity is above the configured major threshold. For any assistance, contact My Oracle Support.

3.3.131 SCPSubscriptionDBCertificateExpiryAboveMinorThreshold

Table 3-135 SCPSubscriptionDBCertificateExpiryAboveMinorThreshold

Field	Details
Description	'Subscription DB client certificate expiration has reached the minor threshold'
Summary	'Subscription DB client certificate expiration is at minor level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Minor
Condition	This alert is raised when the client certificate goes beyond minor threshold. ocscp_subscription_ssl_certificate_expire{expireLevel="minor"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.2006
Metric Used	ocscp_subscription_ssl_certificate_expire

Table 3-135 (Cont.) SCPSubscriptionDBCertificateExpiryAboveMinorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured minor threshold. Diagnostic Information: - Check the alert labels to identify the affected certificate store. The certStoreType label indicates whether the certificate is from primaryKeystore, primaryTruststore, secondaryKeystore, or secondaryTruststore. - Verify the certificate expiry date in the DB TLS Kubernetes secret configured under global.database.ssl. The global.sslCertExpiryMinorThreshold parameter configures the minor threshold. Recovery: - Plan the certificate renewal before the certificate reaches the major or critical threshold. - Renew or replace the affected DB TLS certificate and update the corresponding Kubernetes secret. - The alert clears automatically when certificate validity is above the configured minor threshold. For any assistance, contact My Oracle Support.

3.3.132 SCPSubscriptionKeystoreExpireOutOfTolerance

Table 3-136 SCPSubscriptionKeystoreExpireOutOfTolerance

Field	Details
Description	'Subscription DB keystore expiration is outside the configured tolerance'
Summary	'Subscription DB keystore expiration is outside tolerance. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the keystore is out of tolerance threshold for DB connection.ocscp_subscription_keystore_expire_out_of_tolerance > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.2007
Metric Used	ocscp_subscription_keystore_expire_out_of_tolerance

Table 3-136 (Cont.) SCPSubscriptionKeystoreExpireOutOfTolerance

Field	Details
Recommended Action	Cause: The DB TLS keystore or truststore certificate for the SCP microservice is expired or will expire before the configured minimum certificate validity tolerance. Diagnostic Information: - Check the affected microservice pod logs for keystore or truststore validity warnings. - Verify the DB TLS certificate validity for the primary and secondary keystore/truststore configured under <code>global.database.ssl</code>. - Check the ordinality and storeType labels on the alert metric to identify whether the primary or secondary keystore/truststore is outside tolerance. The <code>global.database.ssl.minCertValidity</code> parameter configures the tolerance. Recovery: - Renew or replace the DB TLS certificate and update the corresponding Kubernetes secret for the affected primary or secondary keystore/truststore. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. - The alert clears automatically when the keystore and truststore certificates are valid beyond the configured minimum certificate validity tolerance. For any assistance, contact My Oracle Support.

3.3.133 SCPAuditDBCertificateExpiryAboveMinorThreshold

Table 3-137 SCPAuditDBCertificateExpiryAboveMinorThreshold

Field	Details
Description	'Audit DB client certificate expiration has reached the minor threshold'
Summary	'Audit DB client certificate expiration is at minor level. namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scpfdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Minor
Condition	This alert is raised when the client certificate goes beyond minor threshold <code>ocscp_audit_ssl_certificate_expire{expireLevel="minor"} > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.4004
Metric Used	ocscp_audit_ssl_certificate_expire

Table 3-137 (Cont.) SCPAuditDBCertificateExpiryAboveMinorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured minor threshold. Diagnostic Information: - Check the alert labels to identify the affected certificate store. The certStoreType label indicates whether the certificate is from primaryKeystore, primaryTruststore, secondaryKeystore, or secondaryTruststore. - Verify the certificate expiry date in the DB TLS Kubernetes secret configured under global.database.ssl. The global.sslCertExpiryMinorThreshold parameter configures the minor threshold. Recovery: - Plan the certificate renewal before the certificate reaches the major or critical threshold. - Renew or replace the affected DB TLS certificate and update the corresponding Kubernetes secret. - The alert clears automatically when certificate validity is above the configured minor threshold. For any assistance, contact My Oracle Support.

3.3.134 SCPAuditDBCertificateExpiryAboveCriticalThreshold

Table 3-138 SCPAuditDBCertificateExpiryAboveCriticalThreshold

Field	Details
Description	'Audit DB client certificate expiration has reached the critical threshold'
Summary	'Audit DB client certificate expiration is at critical level. namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scp fqdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the client certificate goes beyond critical threshold.ocscp_audit_ssl_certificate_expire{expireLevel="critical"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.4004
Metric Used	ocscp_audit_ssl_certificate_expire

Table 3-138 (Cont.) SCPAuditDBCertificateExpiryAboveCriticalThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is near expiry or has expired and has reached the configured critical threshold. Diagnostic Information: • Check certStoreType to identify the affected primary or secondary keystore/truststore. • Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. • Check the affected microservice logs for certificate switch, reload, or DB TLS connection issues. The global.sslCertExpiryCriticalThreshold parameter configures the critical threshold. Recovery: • Renew or replace the affected DB TLS certificate immediately and update the corresponding Kubernetes secret. • Restart or reload the affected SCP microservice pod so that it uses the updated certificate. If secondary DB TLS certificates are configured, verify whether traffic has switched to the secondary certificate path. • The alert clears automatically when certificate validity is above the configured critical threshold. For any assistance, contact My Oracle Support.

3.3.135 SCPAuditKeystoreExpireOutOfTolerance

Table 3-139 SCPAuditKeystoreExpireOutOfTolerance

Field	Details
Description	'Audit DB keystore expiration is outside the configured tolerance'
Summary	'Audit DB keystore expiration is outside tolerance. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is triggered when the keystore is out of tolerance threshold for DB connection.ocscp_audit_keystore_expire_out_of_tolerance > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.4005
Metric Used	ocscp_audit_keystore_expire_out_of_tolerance

Table 3-139 (Cont.) SCPAuditKeystoreExpireOutOfTolerance

Field	Details
Recommended Action	Cause: The DB TLS keystore or truststore certificate for the SCP microservice is expired or will expire before the configured minimum certificate validity tolerance. Diagnostic Information: - Check the affected microservice pod logs for keystore or truststore validity warnings. - Verify the DB TLS certificate validity for the primary and secondary keystore/truststore configured under <code>global.database.ssl</code>. - Check the ordinality and storeType labels on the alert metric to identify whether the primary or secondary keystore/truststore is outside tolerance. The <code>global.database.ssl.minCertValidity</code> parameter configures the tolerance. Recovery: - Renew or replace the DB TLS certificate and update the corresponding Kubernetes secret for the affected primary or secondary keystore/truststore. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. - The alert clears automatically when the keystore and truststore certificates are valid beyond the configured minimum certificate validity tolerance. For any assistance, contact My Oracle Support.

3.3.136 SCPAuditDBCertificateExpiryAboveMajorThreshold

Table 3-140 SCPAuditDBCertificateExpiryAboveMajorThreshold

Field	Details
Description	'Audit DB client certificate expiration has reached the major threshold'
Summary	'Audit DB client certificate expiration is at major level. namespace: {{ \$labels.namespace }}, podname: {{ \$labels.pod }}, scp fqdn: {{ \$labels.ocscp_fqdn }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when the client certificate goes beyond major threshold. <code>threshold.ocscp_audit_ssl_certificate_expire{expireLevel="major"} > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.4004
Metric Used	ocscp_audit_ssl_certificate_expire

Table 3-140 (Cont.) SCPAuditDBCertificateExpiryAboveMajorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured major threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. The global.sslCertExpiryMajorThreshold parameter configures the major threshold. Recovery: - Renew or replace the affected DB TLS certificate as soon as possible and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod, if required, so that it uses the updated certificate. - The alert clears automatically when certificate validity is above the configured major threshold. For any assistance, contact My Oracle Support.

3.3.137 SCPConfigurationDBCertificateExpiryAboveCriticalThreshold

Table 3-141 SCPConfigurationDBCertificateExpiryAboveCriticalThreshold

Field	Details
Description	'Configuration DB client certificate expiration has reached the critical threshold'
Summary	'Configuration DB client certificate expiration is at critical level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Critical
Condition	This alert is raised when the client certificate goes beyond critical threshold. ocscp_configuration_ssl_certificate_expire{expireLevel="critical"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.1000
Metric Used	ocscp_configuration_ssl_certificate_expire

Table 3-141 (Cont.) SCPConfigurationDBCertificateExpiryAboveCriticalThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is near expiry or has expired and has reached the configured critical threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. - Check the affected microservice logs for certificate switch, reload, or DB TLS connection issues. The global.sslCertExpiryCriticalThreshold parameter configures the critical threshold. Recovery: - Renew or replace the affected DB TLS certificate immediately and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. If secondary DB TLS certificates are configured, verify whether traffic has switched to the secondary certificate path. - The alert clears automatically when certificate validity is above the configured critical threshold. For any assistance, contact My Oracle Support.

3.3.138 SCPConfigurationDBCertificateExpiryAboveMajorThreshold

Table 3-142 SCPConfigurationDBCertificateExpiryAboveMajorThreshold

Field	Details
Description	'Configuration DB client certificate expiration has reached the major threshold'
Summary	'Configuration DB client certificate expiration is at major level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	major
Condition	This alert is raised when the client certificate goes beyond major threshold. ocscp_configuration_ssl_certificate_expire{expireLevel="major"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.1000
Metric Used	ocscp_configuration_ssl_certificate_expire

Table 3-142 (Cont.) SCPConfigurationDBCertificateExpiryAboveMajorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured major threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. The global.sslCertExpiryMajorThreshold parameter configures the major threshold. Recovery: - Renew or replace the affected DB TLS certificate as soon as possible and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod, if required, so that it uses the updated certificate. - The alert clears automatically when certificate validity is above the configured major threshold. For any assistance, contact My Oracle Support.

3.3.139 SCPConfigurationDBCertificateExpiryAboveMinorThreshold

Table 3-143 SCPConfigurationDBCertificateExpiryAboveMinorThreshold

Field	Details
Description	'Configuration DB client certificate expiration has reached the minor threshold'
Summary	'Configuration DB client certificate expiration is at minor level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Minor
Condition	This alert is raised when the client certificate goes beyond minor threshold.ocscp_configuration_ssl_certificate_expire{expireLevel="minor"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.1000
Metric Used	ocscp_configuration_ssl_certificate_expire

Table 3-143 (Cont.) SCPConfigurationDBCertificateExpiryAboveMinorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured minor threshold. Diagnostic Information: - Check the alert labels to identify the affected certificate store. The certStoreType label indicates whether the certificate is from primaryKeystore, primaryTruststore, secondaryKeystore, or secondaryTruststore. - Verify the certificate expiry date in the DB TLS Kubernetes secret configured under global.database.ssl. The global.sslCertExpiryMinorThreshold parameter configures the minor threshold. Recovery: - Plan the certificate renewal before the certificate reaches the major or critical threshold. - Renew or replace the affected DB TLS certificate and update the corresponding Kubernetes secret. - The alert clears automatically when certificate validity is above the configured minor threshold. For any assistance, contact My Oracle Support.

3.3.140 SCPConfigurationKeystoreExpireOutOfTolerance

Table 3-144 SCPConfigurationKeystoreExpireOutOfTolerance

Field	Details
Description	'Configuration DB keystore expiration is outside the configured tolerance'
Summary	'Configuration DB keystore expiration is outside tolerance. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the keystore is out of tolerance threshold for DB connection.ocscp_configuration_keystore_expire_out_of_tolerance > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.1001
Metric Used	ocscp_configuration_keystore_expire_out_of_tolerance

Table 3-144 (Cont.) SCPConfigurationKeystoreExpireOutOfTolerance

Field	Details
Recommended Action	Cause: The DB TLS keystore or truststore certificate for the SCP microservice is expired or will expire before the configured minimum certificate validity tolerance. Diagnostic Information: - Check the affected microservice pod logs for keystore or truststore validity warnings. - Verify the DB TLS certificate validity for the primary and secondary keystore/truststore configured under <code>global.database.ssl</code>. - Check the ordinality and storeType labels on the alert metric to identify whether the primary or secondary keystore/truststore is outside tolerance. The <code>global.database.ssl.minCertValidity</code> parameter configures the tolerance. Recovery: - Renew or replace the DB TLS certificate and update the corresponding Kubernetes secret for the affected primary or secondary keystore/truststore. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. - The alert clears automatically when the keystore and truststore certificates are valid beyond the configured minimum certificate validity tolerance. For any assistance, contact My Oracle Support.

3.3.141 SCPAlternateResolutionDBCertificateExpiryAboveMajorThreshold

Table 3-145 SCPAlternateResolutionDBCertificateExpiryAboveMajorThreshold

Field	Details
Description	'Alternate Resolution DB client certificate expiration has reached the major threshold'
Summary	'Alternate Resolution DB client certificate expiration is at major level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Major
Condition	This alert is raised when the client certificate goes beyond major threshold. <code>ocscp_alternate_resolution_ssl_certificate_expire{expireLevel="major"} > 0</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.12000
Metric Used	ocscp_alternate_resolution_ssl_certificate_expire

Table 3-145 (Cont.) SCPAlternateResolutionDBCertificateExpiryAboveMajorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured major threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. The global.sslCertExpiryMajorThreshold parameter configures the major threshold. Recovery: - Renew or replace the affected DB TLS certificate as soon as possible and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod, if required, so that it uses the updated certificate. - The alert clears automatically when certificate validity is above the configured major threshold. For any assistance, contact My Oracle Support.

3.3.142 SCPAlternateResolutionDBCertificateExpiryAboveMinorThreshold

Table 3-146 SCPAlternateResolutionDBCertificateExpiryAboveMinorThreshold

Field	Details
Description	'Alternate Resolution DB client certificate expiration has reached the minor threshold'
Summary	'Alternate Resolution DB client certificate expiration is at minor level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Minor
Condition	This alert is raised when the client certificate goes beyond minor threshold.ocscp_alternate_resolution_ssl_certificate_expire{expireLevel="minor"} > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.12000
Metric Used	ocscp_alternate_resolution_ssl_certificate_expire

Table 3-146 (Cont.) SCPAlternateResolutionDBCertificateExpiryAboveMinorThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is approaching expiry and has reached the configured minor threshold. Diagnostic Information: - Check the alert labels to identify the affected certificate store. The certStoreType label indicates whether the certificate is from primaryKeystore, primaryTruststore, secondaryKeystore, or secondaryTruststore. - Verify the certificate expiry date in the DB TLS Kubernetes secret configured under global.database.ssl. The global.sslCertExpiryMinorThreshold parameter configures the minor threshold. Recovery: - Plan the certificate renewal before the certificate reaches the major or critical threshold. - Renew or replace the affected DB TLS certificate and update the corresponding Kubernetes secret. - The alert clears automatically when certificate validity is above the configured minor threshold. For any assistance, contact My Oracle Support.

3.3.143 SCPAlternateResolutionKeystoreExpireOutOfTolerance

Table 3-147 SCPAlternateResolutionKeystoreExpireOutOfTolerance

Field	Details
Description	'Alternate Resolution DB keystore expiration is outside the configured tolerance'
Summary	'Alternate Resolution DB keystore expiration is outside tolerance. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{{ end }}'
Severity	Critical
Condition	This alert is raised when the keystore is out of tolerance threshold for DB connection. ocscp_alternate_resolution_keystore_expire_out_of_tolerance > 0
OID	1.3.6.1.4.1.323.5.3.35.1.2.12001
Metric Used	ocscp_alternate_resolution_keystore_expire_out_of_tolerance

Table 3-147 (Cont.) SCPAlternateResolutionKeystoreExpireOutOfTolerance

Field	Details
Recommended Action	Cause: The DB TLS keystore or truststore certificate for the SCP microservice is expired or will expire before the configured minimum certificate validity tolerance. Diagnostic Information: - Check the affected microservice pod logs for keystore or truststore validity warnings. - Verify the DB TLS certificate validity for the primary and secondary keystore/truststore configured under <code>global.database.ssl</code>. - Check the ordinality and storeType labels on the alert metric to identify whether the primary or secondary keystore/truststore is outside tolerance. The <code>global.database.ssl.minCertValidity</code> parameter configures the tolerance. Recovery: - Renew or replace the DB TLS certificate and update the corresponding Kubernetes secret for the affected primary or secondary keystore/truststore. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. - The alert clears automatically when the keystore and truststore certificates are valid beyond the configured minimum certificate validity tolerance. For any assistance, contact My Oracle Support.

3.3.144 SCPAlternateResolutionDBCertificateExpiryAboveCriticalThreshold

Table 3-148 SCPAlternateResolutionDBCertificateExpiryAboveCriticalThreshold

Field	Details
Description	'Alternate Resolution DB client certificate expiration has reached the critical threshold'
Summary	'Alternate Resolution DB client certificate expiration is at critical level. namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, scpfqdn: {{\$labels.ocscp_fqdn}}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }'
Severity	Critical
Condition	This alert is raised when the client certificate goes beyond critical threshold. <code>ocscp_alternate_resolution_ssl_certificate_expire{expireLevel="critical"}</code>
OID	1.3.6.1.4.1.323.5.3.35.1.2.12000
Metric Used	ocscp_alternate_resolution_ssl_certificate_expire

Table 3-148 (Cont.) SCPAlternateResolutionDBCertificateExpiryAboveCriticalThreshold

Field	Details
Recommended Action	Cause: The DB client certificate used by the SCP microservice keystore or truststore is near expiry or has expired and has reached the configured critical threshold. Diagnostic Information: - Check certStoreType to identify the affected primary or secondary keystore/truststore. - Verify the certificate expiry in the DB TLS secret configured under global.database.ssl. - Check the affected microservice logs for certificate switch, reload, or DB TLS connection issues. The global.sslCertExpiryCriticalThreshold parameter configures the critical threshold. Recovery: - Renew or replace the affected DB TLS certificate immediately and update the corresponding Kubernetes secret. - Restart or reload the affected SCP microservice pod so that it uses the updated certificate. If secondary DB TLS certificates are configured, verify whether traffic has switched to the secondary certificate path. - The alert clears automatically when certificate validity is above the configured critical threshold. For any assistance, contact My Oracle Support.

3.4 Mediation Alerts

This section provides detailed information on all mediation alerts, including their descriptions, severity, and recommended actions.

3.4.1 NFMediationUserDefinedVariablesMaxSizeLimitExceeded

Table 3-149 NFMediationUserDefinedVariablesMaxSizeLimitExceeded

Field	Details
Description	This alert is raised when the total size of all user-defined variables and their values exceeds the configured size.
Summary	namespace: {{\$labels.namespace}}, podname: {{\$labels.pod}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }}: Currently user defined variables maximum size limit exceeded the configured limit {{ \$value printf "%.2f" }} times
Severity	minor
Condition	increase(ocscp_med_udvs_max_size_limit_exceeded_total[5m]) > 0
OID	1.3.6.1.4.1.323.5.3.47.1.2.1005
Metric Used	ocscp_med_udvs_max_size_limit_exceeded_total
Recommended Action	Cause: This alert is raised when the total size of all user-defined variables and their values exceeds the configured limit. Diagnostic Information: The size of the user-defined variables list (med-user-defined-var-list) during mediation may have exceeded the configured limit. Recovery: Reduce the size or number of user-defined variables to bring the size within the limit. For any assistance, contact My Oracle Support.

