



# OpenShift Container Platform 4.12

## Network Observability

Configuring and using the Network Observability Operator in OpenShift Container Platform



# OpenShift Container Platform 4.12 Network Observability

---

Configuring and using the Network Observability Operator in OpenShift Container Platform

## Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

Use the Network Observability Operator to observe and analyze network traffic flows for OpenShift Container Platform clusters.

## Table of Contents

<b>CHAPTER 1. NETWORK OBSERVABILITY OPERATOR RELEASE NOTES</b> .....	<b>8</b>
1.1. NETWORK OBSERVABILITY OPERATOR 1.6.0	8
1.1.1. New features and enhancements	8
1.1.1.1. Enhanced use of Network Observability Operator without Loki	8
1.1.1.2. Custom metrics API	8
1.1.1.3. eBPF performance enhancements	8
1.1.1.4. eBPF collection rule-based filtering	9
1.1.2. Technology Preview features	9
1.1.2.1. Network Observability CLI	9
1.1.3. Bug fixes	9
1.1.4. Known issues	10
1.2. NETWORK OBSERVABILITY OPERATOR 1.5.0	10
1.2.1. New features and enhancements	10
1.2.1.1. DNS tracking enhancements	10
1.2.1.2. Round-trip time (RTT)	10
1.2.1.3. Metrics, dashboards, and alerts enhancements	10
1.2.1.4. Improvements for Network Observability without Loki	11
1.2.1.5. Availability zones	11
1.2.1.6. Notable enhancements	11
Performance enhancements	11
Web console enhancements:	11
Configuration enhancements:	11
1.2.2. Bug fixes	12
1.2.3. Known issues	12
1.3. NETWORK OBSERVABILITY OPERATOR 1.4.2	13
1.3.1. CVEs	13
1.4. NETWORK OBSERVABILITY OPERATOR 1.4.1	13
1.4.1. CVEs	13
1.4.2. Bug fixes	13
1.5. NETWORK OBSERVABILITY OPERATOR 1.4.0	13
1.5.1. Channel removal	14
1.5.2. New features and enhancements	14
1.5.2.1. Notable enhancements	14
Web console enhancements:	14
Configuration enhancements:	14
1.5.2.2. Network Observability without Loki	14
1.5.2.3. DNS tracking	14
1.5.2.4. SR-IOV support	15
1.5.2.5. IPFIX exporter support	15
1.5.2.6. s390x architecture support	15
1.5.3. Bug fixes	15
1.5.4. Known issues	15
1.6. NETWORK OBSERVABILITY OPERATOR 1.3.0	16
1.6.1. Channel deprecation	16
1.6.2. New features and enhancements	16
1.6.2.1. Multi-tenancy in Network Observability	16
1.6.2.2. Flow-based metrics dashboard	16
1.6.2.3. Troubleshooting with the must-gather tool	16
1.6.2.4. Multiple architectures now supported	17
1.6.3. Deprecated features	17
1.6.3.1. Deprecated configuration parameter setting	17

1.6.4. Bug fixes	17
1.6.5. Known issues	18
1.7. NETWORK OBSERVABILITY OPERATOR 1.2.0	18
1.7.1. Preparing for the next update	18
1.7.2. New features and enhancements	18
1.7.2.1. Histogram in Traffic Flows view	18
1.7.2.2. Conversation tracking	18
1.7.2.3. Network Observability health alerts	18
1.7.3. Bug fixes	19
1.7.4. Known issue	19
1.7.5. Notable technical changes	19
1.8. NETWORK OBSERVABILITY OPERATOR 1.1.0	19
1.8.1. Bug fix	20
<b>CHAPTER 2. ABOUT NETWORK OBSERVABILITY</b>	<b>21</b>
2.1. OPTIONAL DEPENDENCIES OF THE NETWORK OBSERVABILITY OPERATOR	21
2.2. NETWORK OBSERVABILITY OPERATOR	21
2.3. OPENSIFT CONTAINER PLATFORM CONSOLE INTEGRATION	21
2.3.1. Network Observability metrics dashboards	21
2.3.2. Network Observability topology views	21
2.3.3. Traffic flow tables	22
2.4. NETWORK OBSERVABILITY CLI	22
<b>CHAPTER 3. INSTALLING THE NETWORK OBSERVABILITY OPERATOR</b>	<b>23</b>
3.1. NETWORK OBSERVABILITY WITHOUT LOKI	23
3.2. INSTALLING THE LOKI OPERATOR	24
3.2.1. Creating a secret for Loki storage	25
3.2.2. Creating a LokiStack custom resource	25
3.2.3. Creating a new group for the cluster-admin user role	26
3.2.4. Custom admin group access	27
3.2.5. Loki deployment sizing	28
3.2.6. LokiStack ingestion limits and health alerts	28
3.2.7. Enabling multi-tenancy in Network Observability	29
3.3. INSTALLING THE NETWORK OBSERVABILITY OPERATOR	29
3.4. IMPORTANT FLOW COLLECTOR CONFIGURATION CONSIDERATIONS	30
3.5. INSTALLING KAFKA (OPTIONAL)	31
3.6. UNINSTALLING THE NETWORK OBSERVABILITY OPERATOR	31
<b>CHAPTER 4. NETWORK OBSERVABILITY OPERATOR IN OPENSIFT CONTAINER PLATFORM</b>	<b>33</b>
4.1. VIEWING STATUSES	33
4.2. NETWORK OBSERVABILITY OPERATOR ARCHITECTURE	34
4.3. VIEWING NETWORK OBSERVABILITY OPERATOR STATUS AND CONFIGURATION	35
<b>CHAPTER 5. CONFIGURING THE NETWORK OBSERVABILITY OPERATOR</b>	<b>37</b>
5.1. VIEW THE FLOWCOLLECTOR RESOURCE	37
5.2. CONFIGURING THE FLOW COLLECTOR RESOURCE WITH KAFKA	39
5.3. EXPORT ENRICHED NETWORK FLOW DATA	40
5.4. UPDATING THE FLOW COLLECTOR RESOURCE	41
5.5. CONFIGURING QUICK FILTERS	41
5.6. CONFIGURING MONITORING FOR SR-IOV INTERFACE TRAFFIC	42
5.7. RESOURCE MANAGEMENT AND PERFORMANCE CONSIDERATIONS	43
5.7.1. Resource considerations	44
<b>CHAPTER 6. NETWORK POLICY</b>	<b>46</b>

6.1. CREATING A NETWORK POLICY FOR NETWORK OBSERVABILITY	46
6.2. EXAMPLE NETWORK POLICY	46
<b>CHAPTER 7. OBSERVING THE NETWORK TRAFFIC</b>	<b>48</b>
7.1. OBSERVING THE NETWORK TRAFFIC FROM THE OVERVIEW VIEW	48
7.1.1. Working with the Overview view	48
7.1.2. Configuring advanced options for the Overview view	48
7.1.2.1. Managing panels and display	48
7.1.2.2. DNS tracking	49
7.1.3. Round-Trip Time	49
7.1.4. eBPF flow rule filter	50
7.1.4.1. Ingress and egress traffic filtering	50
7.1.4.2. Dashboard and metrics integrations	50
7.1.4.3. Flow filter configuration parameters	50
7.2. OBSERVING THE NETWORK TRAFFIC FROM THE TRAFFIC FLOWS VIEW	52
7.2.1. Working with the Traffic flows view	52
7.2.2. Configuring advanced options for the Traffic flows view	52
7.2.2.1. Managing columns	52
7.2.2.2. Exporting the traffic flow data	52
7.2.3. Working with conversation tracking	53
7.2.4. Working with DNS tracking	54
7.2.5. Working with RTT tracing	55
7.2.5.1. Using the histogram	56
7.2.6. Working with availability zones	56
7.2.7. Filtering eBPF flow data using a global rule	57
7.3. OBSERVING THE NETWORK TRAFFIC FROM THE TOPOLOGY VIEW	58
7.3.1. Working with the Topology view	59
7.3.2. Configuring the advanced options for the Topology view	59
7.3.2.1. Exporting the topology view	59
7.4. FILTERING THE NETWORK TRAFFIC	59
<b>CHAPTER 8. USING METRICS WITH DASHBOARDS AND ALERTS</b>	<b>62</b>
8.1. VIEWING NETWORK OBSERVABILITY METRICS DASHBOARDS	62
8.2. PREDEFINED METRICS	62
8.3. NETWORK OBSERVABILITY METRICS	62
8.4. CUSTOM METRICS	64
8.5. CONFIGURING CUSTOM METRICS BY USING FLOWMETRIC API	64
8.6. CONFIGURING CUSTOM CHARTS USING FLOWMETRIC API	66
<b>CHAPTER 9. MONITORING THE NETWORK OBSERVABILITY OPERATOR</b>	<b>70</b>
9.1. HEALTH DASHBOARDS	70
9.2. HEALTH ALERTS	70
9.3. VIEWING HEALTH INFORMATION	70
9.3.1. Disabling health alerts	71
9.4. USING THE EBPF AGENT ALERT	71
<b>CHAPTER 10. SCHEDULING RESOURCES</b>	<b>73</b>
10.1. NETWORK OBSERVABILITY DEPLOYMENT IN SPECIFIC NODES	73
<b>CHAPTER 11. NETWORK OBSERVABILITY CLI</b>	<b>75</b>
11.1. INSTALLING THE NETWORK OBSERVABILITY CLI	75
11.1.1. About the Network Observability CLI	75
11.1.2. Installing the Network Observability CLI	75
11.2. USING THE NETWORK OBSERVABILITY CLI	76

11.2.1. Capturing flows	76
11.2.2. Capturing packets	78
11.2.3. Cleaning the Network Observability CLI	78
11.3. NETWORK OBSERVABILITY CLI (OC NETOBSERV) REFERENCE	79
11.3.1. oc netobserv CLI reference	79
11.3.1.1. Network Observability enrichment	79
11.3.1.2. Flow capture options	80
11.3.1.3. Packet capture options	81
<b>CHAPTER 12. FLOWCOLLECTOR CONFIGURATION PARAMETERS</b>	<b>82</b>
12.1. FLOWCOLLECTOR API SPECIFICATIONS	82
12.1.1. .metadata	83
12.1.2. .spec	83
12.1.3. .spec.agent	85
12.1.4. .spec.agent.ebpf	85
12.1.5. .spec.agent.ebpf.advanced	88
12.1.6. .spec.agent.ebpf.advanced.scheduling	89
12.1.7. .spec.agent.ebpf.advanced.scheduling.affinity	90
12.1.8. .spec.agent.ebpf.advanced.scheduling.tolerations	90
12.1.9. .spec.agent.ebpf.flowFilter	90
12.1.10. .spec.agent.ebpf.metrics	92
12.1.11. .spec.agent.ebpf.metrics.server	92
12.1.12. .spec.agent.ebpf.metrics.server.tls	93
12.1.13. .spec.agent.ebpf.metrics.server.tls.provided	93
12.1.14. .spec.agent.ebpf.metrics.server.tls.providedCaFile	94
12.1.15. .spec.agent.ebpf.resources	95
12.1.16. .spec.consolePlugin	95
12.1.17. .spec.consolePlugin.advanced	96
12.1.18. .spec.consolePlugin.advanced.scheduling	98
12.1.19. .spec.consolePlugin.advanced.scheduling.affinity	98
12.1.20. .spec.consolePlugin.advanced.scheduling.tolerations	99
12.1.21. .spec.consolePlugin.autoscaler	99
12.1.22. .spec.consolePlugin.portNaming	99
12.1.23. .spec.consolePlugin.quickFilters	99
12.1.24. .spec.consolePlugin.quickFilters[]	100
12.1.25. .spec.consolePlugin.resources	100
12.1.26. .spec.exporters	101
12.1.27. .spec.exporters[]	101
12.1.28. .spec.exporters[].ipfix	101
12.1.29. .spec.exporters[].kafka	102
12.1.30. .spec.exporters[].kafka.sasl	103
12.1.31. .spec.exporters[].kafka.sasl.clientIDReference	103
12.1.32. .spec.exporters[].kafka.sasl.clientSecretReference	104
12.1.33. .spec.exporters[].kafka.tls	104
12.1.34. .spec.exporters[].kafka.tls.caCert	105
12.1.35. .spec.exporters[].kafka.tls.userCert	106
12.1.36. .spec.kafka	107
12.1.37. .spec.kafka.sasl	107
12.1.38. .spec.kafka.sasl.clientIDReference	108
12.1.39. .spec.kafka.sasl.clientSecretReference	108
12.1.40. .spec.kafka.tls	109
12.1.41. .spec.kafka.tls.caCert	109
12.1.42. .spec.kafka.tls.userCert	110



12.1.43. .spec.loki	111
12.1.44. .spec.loki.advanced	113
12.1.45. .spec.loki.lokiStack	114
12.1.46. .spec.loki.manual	114
12.1.47. .spec.loki.manual.statusTls	116
12.1.48. .spec.loki.manual.statusTls.caCert	117
12.1.49. .spec.loki.manual.statusTls.userCert	117
12.1.50. .spec.loki.manual.tls	118
12.1.51. .spec.loki.manual.tls.caCert	119
12.1.52. .spec.loki.manual.tls.userCert	119
12.1.53. .spec.loki.microservices	120
12.1.54. .spec.loki.microservices.tls	121
12.1.55. .spec.loki.microservices.tls.caCert	121
12.1.56. .spec.loki.microservices.tls.userCert	122
12.1.57. .spec.loki.monolithic	123
12.1.58. .spec.loki.monolithic.tls	123
12.1.59. .spec.loki.monolithic.tls.caCert	124
12.1.60. .spec.loki.monolithic.tls.userCert	125
12.1.61. .spec.processor	125
12.1.62. .spec.processor.advanced	128
12.1.63. .spec.processor.advanced.scheduling	129
12.1.64. .spec.processor.advanced.scheduling.affinity	130
12.1.65. .spec.processor.advanced.scheduling.tolerations	130
12.1.66. .spec.processor.kafkaConsumerAutoscaler	131
12.1.67. .spec.processor.metrics	131
12.1.68. .spec.processor.metrics.server	132
12.1.69. .spec.processor.metrics.server.tls	133
12.1.70. .spec.processor.metrics.server.tls.provided	133
12.1.71. .spec.processor.metrics.server.tls.providedCaFile	134
12.1.72. .spec.processor.resources	135
12.1.73. .spec.processor.subnetLabels	135
12.1.74. .spec.processor.subnetLabels.customLabels	136
12.1.75. .spec.processor.subnetLabels.customLabels[]	136
12.1.76. .spec.prometheus	137
12.1.77. .spec.prometheus.querier	137
12.1.78. .spec.prometheus.querier.manual	138
12.1.79. .spec.prometheus.querier.manual.tls	139
12.1.80. .spec.prometheus.querier.manual.tls.caCert	139
12.1.81. .spec.prometheus.querier.manual.tls.userCert	140
<b>CHAPTER 13. FLOWMETRIC CONFIGURATION PARAMETERS</b> .....	<b>142</b>
13.1. FLOWMETRIC [FLOWS.NETOBSERV.IO/VIALPHA1]	142
13.1.1. .metadata	143
13.1.2. .spec	143
13.1.3. .spec.charts	146
13.1.4. .spec.charts[]	146
13.1.5. .spec.charts[].queries	147
13.1.6. .spec.charts[].queries[]	147
13.1.7. .spec.filters	148
13.1.8. .spec.filters[]	148
<b>CHAPTER 14. NETWORK FLOWS FORMAT REFERENCE</b> .....	<b>150</b>
14.1. NETWORK FLOWS FORMAT REFERENCE	150

<b>CHAPTER 15. TROUBLESHOOTING NETWORK OBSERVABILITY .....</b>	<b>155</b>
15.1. USING THE MUST-GATHER TOOL	155
15.2. CONFIGURING NETWORK TRAFFIC MENU ENTRY IN THE OPENSIFT CONTAINER PLATFORM CONSOLE	155
15.3. FLOWLOGS-PIPELINE DOES NOT CONSUME NETWORK FLOWS AFTER INSTALLING KAFKA	157
15.4. FAILING TO SEE NETWORK FLOWS FROM BOTH BR-INT AND BR-EX INTERFACES	157
15.5. NETWORK OBSERVABILITY CONTROLLER MANAGER POD RUNS OUT OF MEMORY	158
15.6. RUNNING CUSTOM QUERIES TO LOKI	158
15.7. TROUBLESHOOTING LOKI RESOURCEEXHAUSTED ERROR	159
15.8. LOKI EMPTY RING ERROR	160
15.9. RESOURCE TROUBLESHOOTING	160
15.10. LOKISTACK RATE LIMIT ERRORS	160



# CHAPTER 1. NETWORK OBSERVABILITY OPERATOR RELEASE NOTES

The Network Observability Operator enables administrators to observe and analyze network traffic flows for OpenShift Container Platform clusters.

These release notes track the development of the Network Observability Operator in the OpenShift Container Platform.

For an overview of the Network Observability Operator, see [About Network Observability Operator](#).

## 1.1. NETWORK OBSERVABILITY OPERATOR 1.6.0

The following advisory is available for the Network Observability Operator 1.6.0:

- [Network Observability Operator 1.6.0](#)

### 1.1.1. New features and enhancements

#### 1.1.1.1. Enhanced use of Network Observability Operator without Loki

You can now use Prometheus metrics and rely less on Loki for storage when using the Network Observability Operator. For more information, see [Network Observability without Loki](#).

#### 1.1.1.2. Custom metrics API

You can create custom metrics out of flowlogs data by using the **FlowMetrics** API. Flowlogs data can be used with Prometheus labels to customize cluster information on your dashboards. You can add custom labels for any subnet that you want to identify in your flows and metrics. This enhancement can also be used to more easily identify external traffic by using the new labels **SrcSubnetLabel** and **DstSubnetLabel**, which exists both in flow logs and in metrics. Those fields are empty when there is external traffic, which gives a way to identify it. For more information, see [Custom metrics](#) and [FlowMetric API reference](#).

#### 1.1.1.3. eBPF performance enhancements

Experience improved performances of the eBPF agent, in terms of CPU and memory, with the following updates:

- The eBPF agent now uses TCX webhooks instead of TC.
- The **NetObserv / Health** dashboard has a new section that shows eBPF metrics.
  - Based on the new eBPF metrics, an alert notifies you when the eBPF agent is dropping flows.
- Loki storage demand decreases significantly now that duplicated flows are removed. Instead of having multiple, individual duplicated flows per network interface, there is one de-duplicated flow with a list of related network interfaces.



## IMPORTANT

With the duplicated flows update, the **Interface** and **Interface Direction** fields in the **Network Traffic** table are renamed to **Interfaces** and **Interface Directions**, so any bookmarked **Quick filter** queries using these fields need to be updated to **interfaces** and **ifdirections**.

For more information, see [Using the eBPF agent alert](#) and [Quick filters](#).

### 1.1.1.4. eBPF collection rule-based filtering

You can use rule-based filtering to reduce the volume of created flows. When this option is enabled, the **Netobserv / Health** dashboard for eBPF agent statistics has the **Filtered flows rate** view. For more information, see [eBPF flow rule filter](#).

## 1.1.2. Technology Preview features

Some features in this release are currently in Technology Preview. These experimental features are not intended for production use. Note the following scope of support on the Red Hat Customer Portal for these features:

[Technology Preview Features Support Scope](#)

### 1.1.2.1. Network Observability CLI

You can debug and troubleshoot network traffic issues without needing to install the Network Observability Operator by using the Network Observability CLI. Capture and visualize flow and packet data in real-time with no persistent storage requirement during the capture. For more information, see [Network Observability CLI](#) and [Network Observability CLI 1.6.0](#)

## 1.1.3. Bug fixes

- Previously, a dead link to the OpenShift container platform documentation was displayed in the Operator Lifecycle Manager (OLM) form for the **FlowMetrics** API creation. Now the link has been updated to point to a valid page. ([NETOBSERV-1607](#))
- Previously, the Network Observability Operator description in the Operator Hub displayed a broken link to the documentation. With this fix, this link is restored. ([NETOBSERV-1544](#))
- Previously, if Loki was disabled and the Loki **Mode** was set to **LokiStack**, or if Loki manual TLS configuration was configured, the Network Observability Operator still tried to read the Loki CA certificates. With this fix, when Loki is disabled, the Loki certificates are not read, even if there are settings in the Loki configuration. ([NETOBSERV-1647](#))
- Previously, the **oc must-gather** plugin for the Network Observability Operator was only working on the **amd64** architecture and failing on all others because the plugin was using **amd64** for the **oc** binary. Now, the Network Observability Operator **oc must-gather** plugin collects logs on any architecture platform.
- Previously, when filtering on IP addresses using **not equal to**, the Network Observability Operator would return a request error. Now, the IP filtering works in both **equal** and **not equal to** cases for IP addresses and ranges. ([NETOBSERV-1630](#))

- Previously, when a user was not an admin, the error messages were not consistent with the selected tab of the **Network Traffic** view in the web console. Now, the **user not admin** error displays on any tab with improved display. ([NETOBSERV-1621](#))

#### 1.1.4. Known issues

- When the eBPF agent **PacketDrop** feature is enabled, and sampling is configured to a value greater than **1**, reported dropped bytes and dropped packets ignore the sampling configuration. While this is done on purpose to not miss any drops, a side effect is that the reported proportion of drops versus non-drops becomes biased. For example, at a very high sampling rate, such as **1:1000**, it is likely that almost all the traffic appears to be dropped when observed from the console plugin. ([NETOBSERV-1676](#))
- In the **Manage panels** pop-up window in the **Overview** tab, filtering on **total**, **bar**, **donut**, or **line** does not show any result. ([NETOBSERV-1540](#))
- The SR-IOV secondary interface is not detected if the interface was created first and then the eBPF agent was deployed. It is only detected if the agent was deployed first and then the SR-IOV interface is created. ([NETOBSERV-1697](#))
- When Loki is disabled, the **Topology** view in the OpenShift web console always shows the **Cluster** and **Zone** aggregation options in the slider beside the network topology diagram, even when the related features are not enabled. There is no specific workaround, besides ignoring these slider options. ([NETOBSERV-1705](#))
- When Loki is disabled, and the OpenShift web console first loads, it might display an error: **Request failed with status code 400 Loki is disabled**. As a workaround, you can continue switching content on the **Network Traffic** page, such as clicking between the **Topology** and the **Overview** tabs. The error should disappear. ([NETOBSERV-1706](#))

## 1.2. NETWORK OBSERVABILITY OPERATOR 1.5.0

The following advisory is available for the Network Observability Operator 1.5.0:

- [Network Observability Operator 1.5.0](#)

### 1.2.1. New features and enhancements

#### 1.2.1.1. DNS tracking enhancements

In 1.5, the TCP protocol is now supported in addition to UDP. New dashboards are also added to the **Overview** view of the Network Traffic page. For more information, see [Configuring DNS tracking](#) and [Working with DNS tracking](#).

#### 1.2.1.2. Round-trip time (RTT)

You can use TCP handshake Round-Trip Time (RTT) captured from the **fentry/tcp\_rcv\_established** Extended Berkeley Packet Filter (eBPF) hookpoint to read smoothed round-trip time (SRTT) and analyze network flows. In the **Overview**, **Network Traffic**, and **Topology** pages in web console, you can monitor network traffic and troubleshoot with RTT metrics, filtering, and edge labeling. For more information, see [RTT Overview](#) and [Working with RTT](#).

#### 1.2.1.3. Metrics, dashboards, and alerts enhancements

The Network Observability metrics dashboards in **Observe** → **Dashboards** → **NetObserv** have new metrics types you can use to create Prometheus alerts. You can now define available metrics in the **includeList** specification. In previous releases, these metrics were defined in the **ignoreTags** specification. For a complete list of these metrics, see [Network Observability Metrics](#).

#### 1.2.1.4. Improvements for Network Observability without Loki

You can create Prometheus alerts for the **Netobserv** dashboard using DNS and RTT metrics, even if you don't use Loki. In the previous version of Network Observability, 1.4, these metrics were only available for querying and analysis in the **Network Traffic**, **Overview**, and **Topology** views, which are not available without Loki. For more information, see [Network Observability Metrics](#).

#### 1.2.1.5. Availability zones

You can configure the **FlowCollector** resource to collect information about the cluster availability zones. This configuration enriches the network flow data with the [topology.kubernetes.io/zone](#) label value applied to the nodes. For more information, see [Working with availability zones](#).

#### 1.2.1.6. Notable enhancements

The 1.5 release of the Network Observability Operator adds improvements and new capabilities to the OpenShift Container Platform web console plugin and the Operator configuration.

##### Performance enhancements

- The **spec.agent.ebpf.kafkaBatchSize** default is changed from **10MB** to **1MB** to enhance eBPF performance when using Kafka.



##### IMPORTANT

When upgrading from an existing installation, this new value is not set automatically in the configuration. If you monitor a performance regression with the eBPF Agent memory consumption after upgrading, you might consider reducing the **kafkaBatchSize** to the new value.

##### Web console enhancements:

- There are new panels added to the **Overview** view for DNS and RTT: Min, Max, P90, P99.
- There are new panel display options added:
  - Focus on one panel while keeping others viewable but with smaller focus.
  - Switch graph type.
  - Show **Top** and **Overall**.
- A collection latency warning is shown in the **Custom time range** pop-up window.
- There is enhanced visibility for the contents of the **Manage panels** and **Manage columns** pop-up windows.
- The Differentiated Services Code Point (DSCP) field for egress QoS is available for filtering QoS DSCP in the web console **Network Traffic** page.

##### Configuration enhancements:

- The **LokiStack** mode in the **spec.loki.mode** specification simplifies installation by automatically setting URLs, TLS, cluster roles and a cluster role binding, as well as the **authToken** value. The **Manual** mode allows more control over configuration of these settings.
- The API version changes from **flows.netobserv.io/v1beta1** to **flows.netobserv.io/v1beta2**.

### 1.2.2. Bug fixes

- Previously, it was not possible to register the console plugin manually in the web console interface if the automatic registration of the console plugin was disabled. If the **spec.console.register** value was set to **false** in the **FlowCollector** resource, the Operator would override and erase the plugin registration. With this fix, setting the **spec.console.register** value to **false** does not impact the console plugin registration or registration removal. As a result, the plugin can be safely registered manually. ([NETOBSERV-1134](#))
- Previously, using the default metrics settings, the **NetObserve/Health** dashboard was showing an empty graph named **Flows Overhead**. This metric was only available by removing "namespaces-flows" and "namespaces" from the **ignoreTags** list. With this fix, this metric is visible when you use the default metrics setting. ([NETOBSERV-1351](#))
- Previously, the node on which the eBPF Agent was running would not resolve with a specific cluster configuration. This resulted in cascading consequences that culminated in a failure to provide some of the traffic metrics. With this fix, the eBPF agent's node IP is safely provided by the Operator, inferred from the pod status. Now, the missing metrics are restored. ([NETOBSERV-1430](#))
- Previously, the Loki error 'Input size too long' error for the Loki Operator did not include additional information to troubleshoot the problem. With this fix, help is directly displayed in the web console next to the error with a direct link for more guidance. ([NETOBSERV-1464](#))
- Previously, the console plugin read timeout was forced to 30s. With the **FlowCollector v1beta2** API update, you can configure the **spec.loki.readTimeout** specification to update this value according to the Loki Operator **queryTimeout** limit. ([NETOBSERV-1443](#))
- Previously, the Operator bundle did not display some of the supported features by CSV annotations as expected, such as **features.operators.openshift.io/...** With this fix, these annotations are set in the CSV as expected. ([NETOBSERV-1305](#))
- Previously, the **FlowCollector** status sometimes oscillated between **DeploymentInProgress** and **Ready** states during reconciliation. With this fix, the status only becomes **Ready** when all of the underlying components are fully ready. ([NETOBSERV-1293](#))

### 1.2.3. Known issues

- When trying to access the web console, cache issues on OCP 4.14.10 prevent access to the **Observe** view. The web console shows the error message: **Failed to get a valid plugin manifest from /api/plugins/monitoring-plugin/**. The recommended workaround is to update the cluster to the latest minor version. If this does not work, you need to apply the workarounds described in this [Red Hat Knowledgebase article](#). ([NETOBSERV-1493](#))
- Since the 1.3.0 release of the Network Observability Operator, installing the Operator causes a warning kernel taint to appear. The reason for this error is that the Network Observability eBPF agent has memory constraints that prevent preallocating the entire hashmap table. The Operator eBPF agent sets the **BPF\_F\_NO\_PREALLOC** flag so that pre-allocation is disabled when the hashmap is too memory expansive.



## 1.3. NETWORK OBSERVABILITY OPERATOR 1.4.2

The following advisory is available for the Network Observability Operator 1.4.2:

- [2023:6787 Network Observability Operator 1.4.2](#)

### 1.3.1. CVEs

- [2023-39325](#)
- [2023-44487](#)

## 1.4. NETWORK OBSERVABILITY OPERATOR 1.4.1

The following advisory is available for the Network Observability Operator 1.4.1:

- [2023:5974 Network Observability Operator 1.4.1](#)

### 1.4.1. CVEs

- [2023-44487](#)
- [2023-39325](#)
- [2023-29406](#)
- [2023-29409](#)
- [2023-39322](#)
- [2023-39318](#)
- [2023-39319](#)
- [2023-39321](#)

### 1.4.2. Bug fixes

- In 1.4, there was a known issue when sending network flow data to Kafka. The Kafka message key was ignored, causing an error with connection tracking. Now the key is used for partitioning, so each flow from the same connection is sent to the same processor. ([NETOBSERV-926](#))
- In 1.4, the **Inner** flow direction was introduced to account for flows between pods running on the same node. Flows with the **Inner** direction were not taken into account in the generated Prometheus metrics derived from flows, resulting in under-evaluated bytes and packets rates. Now, derived metrics are including flows with the **Inner** direction, providing correct bytes and packets rates. ([NETOBSERV-1344](#))

## 1.5. NETWORK OBSERVABILITY OPERATOR 1.4.0

The following advisory is available for the Network Observability Operator 1.4.0:

- [RHSA-2023:5379 Network Observability Operator 1.4.0](#)

## 1.5.1. Channel removal

You must switch your channel from **v1.0.x** to **stable** to receive the latest Operator updates. The **v1.0.x** channel is now removed.

## 1.5.2. New features and enhancements

### 1.5.2.1. Notable enhancements

The 1.4 release of the Network Observability Operator adds improvements and new capabilities to the OpenShift Container Platform web console plugin and the Operator configuration.

#### Web console enhancements:

- In the **Query Options**, the **Duplicate flows** checkbox is added to choose whether or not to show duplicated flows.
- You can now filter source and destination traffic with **↑ One-way**, **↑ ↓ Back-and-forth**, and **Swap** filters.
- The Network Observability metrics dashboards in **Observe** → **Dashboards** → **NetObserv** and **NetObserv / Health** are modified as follows:
  - The **NetObserv** dashboard shows top bytes, packets sent, packets received per nodes, namespaces, and workloads. Flow graphs are removed from this dashboard.
  - The **NetObserv / Health** dashboard shows flows overhead as well as top flow rates per nodes, namespaces, and workloads.
  - Infrastructure and Application metrics are shown in a split-view for namespaces and workloads.

For more information, see [Network Observability metrics](#) and [Quick filters](#).

#### Configuration enhancements:

- You now have the option to specify different namespaces for any configured ConfigMap or Secret reference, such as in certificates configuration.
- The **spec.processor.clusterName** parameter is added so that the name of the cluster appears in the flows data. This is useful in a multi-cluster context. When using OpenShift Container Platform, leave empty to make it automatically determined.

For more information, see [Flow Collector sample resource](#) and [Flow Collector API Reference](#).

### 1.5.2.2. Network Observability without Loki

The Network Observability Operator is now functional and usable without Loki. If Loki is not installed, it can only export flows to KAFKA or IPFIX format and provide metrics in the Network Observability metrics dashboards. For more information, see [Network Observability without Loki](#).

### 1.5.2.3. DNS tracking

In 1.4, the Network Observability Operator makes use of eBPF tracepoint hooks to enable DNS tracking. You can monitor your network, conduct security analysis, and troubleshoot DNS issues in the **Network Traffic** and **Overview** pages in the web console.

For more information, see [Configuring DNS tracking](#) and [Working with DNS tracking](#).

#### 1.5.2.4. SR-IOV support

You can now collect traffic from a cluster with Single Root I/O Virtualization (SR-IOV) device. For more information, see [Configuring the monitoring of SR-IOV interface traffic](#).

#### 1.5.2.5. IPFIX exporter support

You can now export eBPF-enriched network flows to the IPFIX collector. For more information, see [Export enriched network flow data](#).

#### 1.5.2.6. s390x architecture support

Network Observability Operator can now run on **s390x** architecture. Previously it ran on **amd64**, **ppc64le**, or **arm64**.

### 1.5.3. Bug fixes

- Previously, the Prometheus metrics exported by Network Observability were computed out of potentially duplicated network flows. In the related dashboards, from **Observe** → **Dashboards**, this could result in potentially doubled rates. Note that dashboards from the **Network Traffic** view were not affected. Now, network flows are filtered to eliminate duplicates before metrics calculation, which results in correct traffic rates displayed in the dashboards. ([NETOBSERV-1131](#))
- Previously, the Network Observability Operator agents were not able to capture traffic on network interfaces when configured with Multus or SR-IOV, non-default network namespaces. Now, all available network namespaces are recognized and used for capturing flows, allowing capturing traffic for SR-IOV. There are [configurations needed](#) for the **FlowCollector** and **SRIOVnetwork** custom resource to collect traffic. ([NETOBSERV-1283](#))
- Previously, in the Network Observability Operator details from **Operators** → **Installed Operators**, the **FlowCollector Status** field might have reported incorrect information about the state of the deployment. The status field now shows the proper conditions with improved messages. The history of events is kept, ordered by event date. ([NETOBSERV-1224](#))
- Previously, during spikes of network traffic load, certain eBPF pods were OOM-killed and went into a **CrashLoopBackOff** state. Now, the **eBPF** agent memory footprint is improved, so pods are not OOM-killed and entering a **CrashLoopBackOff** state. ([NETOBSERV-975](#))
- Previously when **processor.metrics.tls** was set to **PROVIDED** the **insecureSkipVerify** option value was forced to be **true**. Now you can set **insecureSkipVerify** to **true** or **false**, and provide a CA certificate if needed. ([NETOBSERV-1087](#))

### 1.5.4. Known issues

- Since the 1.2.0 release of the Network Observability Operator, using Loki Operator 5.6, a Loki certificate change periodically affects the **flowlogs-pipeline** pods and results in dropped flows rather than flows written to Loki. The problem self-corrects after some time, but it still causes temporary flow data loss during the Loki certificate change. This issue has only been observed in large-scale environments of 120 nodes or greater. ([NETOBSERV-980](#))
- Currently, when **spec.agent.ebpf.features** includes **DNSTracking**, larger DNS packets require the **eBPF** agent to look for DNS header outside of the 1st socket buffer (SKB) segment. A new

**eBPF** agent helper function needs to be implemented to support it. Currently, there is no workaround for this issue. ([NETOBSERV-1304](#))

- Currently, when **spec.agent.ebpf.features** includes DNSTracking, DNS over TCP packets requires the **eBPF** agent to look for DNS header outside of the 1st SKB segment. A new **eBPF** agent helper function needs to be implemented to support it. Currently, there is no workaround for this issue. ([NETOBSERV-1245](#))
- Currently, when using a **KAFKA** deployment model, if conversation tracking is configured, conversation events might be duplicated across Kafka consumers, resulting in inconsistent tracking of conversations, and incorrect volumetric data. For that reason, it is not recommended to configure conversation tracking when **deploymentModel** is set to **KAFKA**. ([NETOBSERV-926](#))
- Currently, when the **processor.metrics.server.tls.type** is configured to use a **PROVIDED** certificate, the operator enters an unsteady state that might affect its performance and resource consumption. It is recommended to not use a **PROVIDED** certificate until this issue is resolved, and instead using an auto-generated certificate, setting **processor.metrics.server.tls.type** to **AUTO**. ([NETOBSERV-1293](#))
- Since the 1.3.0 release of the Network Observability Operator, installing the Operator causes a warning kernel taint to appear. The reason for this error is that the Network Observability eBPF agent has memory constraints that prevent preallocating the entire hashmap table. The Operator eBPF agent sets the **BPF\_F\_NO\_PREALLOC** flag so that pre-allocation is disabled when the hashmap is too memory expansive.

## 1.6. NETWORK OBSERVABILITY OPERATOR 1.3.0

The following advisory is available for the Network Observability Operator 1.3.0:

- [RHSA-2023:3905 Network Observability Operator 1.3.0](#)

### 1.6.1. Channel deprecation

You must switch your channel from **v1.0.x** to **stable** to receive future Operator updates. The **v1.0.x** channel is deprecated and planned for removal in the next release.

### 1.6.2. New features and enhancements

#### 1.6.2.1. Multi-tenancy in Network Observability

- System administrators can allow and restrict individual user access, or group access, to the flows stored in Loki. For more information, see [Multi-tenancy in Network Observability](#).

#### 1.6.2.2. Flow-based metrics dashboard

- This release adds a new dashboard, which provides an overview of the network flows in your OpenShift Container Platform cluster. For more information, see [Network Observability metrics](#).

#### 1.6.2.3. Troubleshooting with the must-gather tool

- Information about the Network Observability Operator can now be included in the must-gather data for troubleshooting. For more information, see [Network Observability must-gather](#).

### 1.6.2.4. Multiple architectures now supported

- Network Observability Operator can now run on an **amd64**, **ppc64le**, or **arm64** architectures. Previously, it only ran on **amd64**.

### 1.6.3. Deprecated features

#### 1.6.3.1. Deprecated configuration parameter setting

The release of Network Observability Operator 1.3 deprecates the **spec.Loki.authToken HOST** setting. When using the Loki Operator, you must now only use the **FORWARD** setting.

### 1.6.4. Bug fixes

- Previously, when the Operator was installed from the CLI, the **Role** and **RoleBinding** that are necessary for the Cluster Monitoring Operator to read the metrics were not installed as expected. The issue did not occur when the operator was installed from the web console. Now, either way of installing the Operator installs the required **Role** and **RoleBinding**. ([NETOBSERV-1003](#))
- Since version 1.2, the Network Observability Operator can raise alerts when a problem occurs with the flows collection. Previously, due to a bug, the related configuration to disable alerts, **spec.processor.metrics.disableAlerts** was not working as expected and sometimes ineffectual. Now, this configuration is fixed so that it is possible to disable the alerts. ([NETOBSERV-976](#))
- Previously, when Network Observability was configured with **spec.loki.authToken** set to **DISABLED**, only a **kubeadmin** cluster administrator was able to view network flows. Other types of cluster administrators received authorization failure. Now, any cluster administrator is able to view network flows. ([NETOBSERV-972](#))
- Previously, a bug prevented users from setting **spec.consolePlugin.portNaming.enable** to **false**. Now, this setting can be set to **false** to disable port-to-service name translation. ([NETOBSERV-971](#))
- Previously, the metrics exposed by the console plugin were not collected by the Cluster Monitoring Operator (Prometheus), due to an incorrect configuration. Now the configuration has been fixed so that the console plugin metrics are correctly collected and accessible from the OpenShift Container Platform web console. ([NETOBSERV-765](#))
- Previously, when **processor.metrics.tls** was set to **AUTO** in the **FlowCollector**, the **flowlogs-pipeline servicemonitor** did not adapt the appropriate TLS scheme, and metrics were not visible in the web console. Now the issue is fixed for AUTO mode. ([NETOBSERV-1070](#))
- Previously, certificate configuration, such as used for Kafka and Loki, did not allow specifying a namespace field, implying that the certificates had to be in the same namespace where Network Observability is deployed. Moreover, when using Kafka with TLS/mTLS, the user had to manually copy the certificate(s) to the privileged namespace where the **eBPF** agent pods are deployed and manually manage certificate updates, such as in the case of certificate rotation. Now, Network Observability setup is simplified by adding a namespace field for certificates in the **FlowCollector** resource. As a result, users can now install Loki or Kafka in different namespaces without needing to manually copy their certificates in the Network Observability namespace. The original certificates are watched so that the copies are automatically updated when needed. ([NETOBSERV-773](#))

- Previously, the SCTP, ICMPv4 and ICMPv6 protocols were not covered by the Network Observability agents, resulting in a less comprehensive network flows coverage. These protocols are now recognized to improve the flows coverage. ([NETOBSERV-934](#))

### 1.6.5. Known issues

- When `processor.metrics.tls` is set to **PROVIDED** in the **FlowCollector**, the **flowlogs-pipeline servicemonitor** is not adapted to the TLS scheme. ( [NETOBSERV-1087](#) )
- Since the 1.2.0 release of the Network Observability Operator, using Loki Operator 5.6, a Loki certificate change periodically affects the **flowlogs-pipeline** pods and results in dropped flows rather than flows written to Loki. The problem self-corrects after some time, but it still causes temporary flow data loss during the Loki certificate change. This issue has only been observed in large-scale environments of 120 nodes or greater.([NETOBSERV-980](#))
- When you install the Operator, a warning kernel taint can appear. The reason for this error is that the Network Observability eBPF agent has memory constraints that prevent preallocating the entire hashmap table. The Operator eBPF agent sets the **BPF\_F\_NO\_PREALLOC** flag so that pre-allocation is disabled when the hashmap is too memory expansive.

## 1.7. NETWORK OBSERVABILITY OPERATOR 1.2.0

The following advisory is available for the Network Observability Operator 1.2.0:

- [RHSA-2023:1817 Network Observability Operator 1.2.0](#)

### 1.7.1. Preparing for the next update

The subscription of an installed Operator specifies an update channel that tracks and receives updates for the Operator. Until the 1.2 release of the Network Observability Operator, the only channel available was **v1.0.x**. The 1.2 release of the Network Observability Operator introduces the **stable** update channel for tracking and receiving updates. You must switch your channel from **v1.0.x** to **stable** to receive future Operator updates. The **v1.0.x** channel is deprecated and planned for removal in a following release.

### 1.7.2. New features and enhancements

#### 1.7.2.1. Histogram in Traffic Flows view

- You can now choose to show a histogram bar chart of flows over time. The histogram enables you to visualize the history of flows without hitting the Loki query limit. For more information, see [Using the histogram](#).

#### 1.7.2.2. Conversation tracking

- You can now query flows by **Log Type**, which enables grouping network flows that are part of the same conversation. For more information, see [Working with conversations](#).

#### 1.7.2.3. Network Observability health alerts

- The Network Observability Operator now creates automatic alerts if the **flowlogs-pipeline** is dropping flows because of errors at the write stage or if the Loki ingestion rate limit has been reached. For more information, see [Health dashboards](#).

### 1.7.3. Bug fixes

- Previously, after changing the **namespace** value in the FlowCollector spec, **eBPF** agent pods running in the previous namespace were not appropriately deleted. Now, the pods running in the previous namespace are appropriately deleted. ([NETOBSERV-774](#))
- Previously, after changing the **caCert.name** value in the FlowCollector spec (such as in Loki section), FlowLogs-Pipeline pods and Console plug-in pods were not restarted, therefore they were unaware of the configuration change. Now, the pods are restarted, so they get the configuration change. ([NETOBSERV-772](#))
- Previously, network flows between pods running on different nodes were sometimes not correctly identified as being duplicates because they are captured by different network interfaces. This resulted in over-estimated metrics displayed in the console plug-in. Now, flows are correctly identified as duplicates, and the console plug-in displays accurate metrics. ([NETOBSERV-755](#))
- The "reporter" option in the console plug-in is used to filter flows based on the observation point of either source node or destination node. Previously, this option mixed the flows regardless of the node observation point. This was due to network flows being incorrectly reported as Ingress or Egress at the node level. Now, the network flow direction reporting is correct. The "reporter" option filters for source observation point, or destination observation point, as expected. ([NETOBSERV-696](#))
- Previously, for agents configured to send flows directly to the processor as gRPC+protobuf requests, the submitted payload could be too large and is rejected by the processors' GRPC server. This occurred under very-high-load scenarios and with only some configurations of the agent. The agent logged an error message, such as: *grpc: received message larger than max* . As a consequence, there was information loss about those flows. Now, the gRPC payload is split into several messages when the size exceeds a threshold. As a result, the server maintains connectivity. ([NETOBSERV-617](#))

### 1.7.4. Known issue

- In the 1.2.0 release of the Network Observability Operator, using Loki Operator 5.6, a Loki certificate transition periodically affects the **flowlogs-pipeline** pods and results in dropped flows rather than flows written to Loki. The problem self-corrects after some time, but it still causes temporary flow data loss during the Loki certificate transition. ([NETOBSERV-980](#))

### 1.7.5. Notable technical changes

- Previously, you could install the Network Observability Operator using a custom namespace. This release introduces the **conversion webhook** which changes the **ClusterServiceVersion**. Because of this change, all the available namespaces are no longer listed. Additionally, to enable Operator metrics collection, namespaces that are shared with other Operators, like the **openshift-operators** namespace, cannot be used. Now, the Operator must be installed in the **openshift-netobserv-operator** namespace. You cannot automatically upgrade to the new Operator version if you previously installed the Network Observability Operator using a custom namespace. If you previously installed the Operator using a custom namespace, you must delete the instance of the Operator that was installed and re-install your operator in the **openshift-netobserv-operator** namespace. It is important to note that custom namespaces, such as the commonly used **netobserv** namespace, are still possible for the **FlowCollector**, Loki, Kafka, and other plug-ins. ([NETOBSERV-907](#))([NETOBSERV-956](#))

## 1.8. NETWORK OBSERVABILITY OPERATOR 1.1.0

The following advisory is available for the Network Observability Operator 1.1.0:

- [RHSA-2023:0786 Network Observability Operator Security Advisory Update](#)

The Network Observability Operator is now stable and the release channel is upgraded to **v1.1.0**.

### 1.8.1. Bug fix

- Previously, unless the Loki **authToken** configuration was set to **FORWARD** mode, authentication was no longer enforced, allowing any user who could connect to the OpenShift Container Platform console in an OpenShift Container Platform cluster to retrieve flows without authentication. Now, regardless of the Loki **authToken** mode, only cluster administrators can retrieve flows. ([BZ#2169468](#))



## CHAPTER 2. ABOUT NETWORK OBSERVABILITY

Red Hat offers cluster administrators the Network Observability Operator to observe the network traffic for OpenShift Container Platform clusters. The Network Observability Operator uses the eBPF technology to create network flows. The network flows are then enriched with OpenShift Container Platform information. They are available as Prometheus metrics or as logs in Loki. You can view and analyze the stored network flows information in the OpenShift Container Platform console for further insight and troubleshooting.

### 2.1. OPTIONAL DEPENDENCIES OF THE NETWORK OBSERVABILITY OPERATOR

- **Loki Operator:** Loki is the backend that can be used to store all collected flows with a maximal level of details. You can choose to use [Network Observability without Loki](#), but there are some considerations for doing this, as described in the linked section. If you choose to install Loki, it is recommended to use the Loki Operator, which is supported by Red Hat.
- **AMQ Streams Operator:** Kafka provides scalability, resiliency and high availability in the OpenShift Container Platform cluster for large scale deployments. If you choose to use Kafka, it is recommended to use the AMQ Streams Operator, because it is supported by Red Hat.

### 2.2. NETWORK OBSERVABILITY OPERATOR

The Network Observability Operator provides the Flow Collector API custom resource definition. A Flow Collector instance is a cluster-scoped resource that enables configuration of network flow collection. The Flow Collector instance deploys pods and services that form a monitoring pipeline where network flows are then collected and enriched with the Kubernetes metadata before storing in Loki or generating Prometheus metrics. The eBPF agent, which is deployed as a **daemonset** object, creates the network flows.

### 2.3. OPENSIFT CONTAINER PLATFORM CONSOLE INTEGRATION

OpenShift Container Platform console integration offers overview, topology view and traffic flow tables.

#### 2.3.1. Network Observability metrics dashboards

On the **Overview** tab in the OpenShift Container Platform console, you can view the overall aggregated metrics of the network traffic flow on the cluster. You can choose to display the information by zone, node, namespace, owner, pod, and service. Filters and display options can further refine the metrics. For more information, see [Observing the network traffic from the Overview view](#).

In **Observe** → **Dashboards**, the **Netobserv** dashboards provide a quick overview of the network flows in your OpenShift Container Platform cluster. The **Netobserv/Health** dashboard provides metrics about the health of the Operator. For more information, see [Network Observability Metrics](#) and [Viewing health information](#).

#### 2.3.2. Network Observability topology views

The OpenShift Container Platform console offers the **Topology** tab which displays a graphical representation of the network flows and the amount of traffic. The topology view represents traffic between the OpenShift Container Platform components as a network graph. You can refine the graph by using the filters and display options. You can access the information for zone, node, namespace, owner, pod, and service.

### 2.3.3. Traffic flow tables

The traffic flow table view provides a view for raw flows, non aggregated filtering options, and configurable columns. The OpenShift Container Platform console offers the **Traffic flows** tab which displays the data of the network flows and the amount of traffic.

## 2.4. NETWORK OBSERVABILITY CLI

You can quickly debug and troubleshoot networking issues with Network Observability by using the Network Observability CLI (**oc netobserv**). The Network Observability CLI is a flow and packet visualization tool that relies on eBPF agents to stream collected data to an ephemeral collector pod. It requires no persistent storage during the capture. After the run, the output is transferred to your local machine. This enables quick, live insight into packets and flow data without installing the Network Observability Operator.

## CHAPTER 3. INSTALLING THE NETWORK OBSERVABILITY OPERATOR

Installing Loki is a recommended prerequisite for using the Network Observability Operator. You can choose to use [Network Observability without Loki](#), but there are some considerations for doing this, described in the previously linked section.

The Loki Operator integrates a gateway that implements multi-tenancy and authentication with Loki for data flow storage. The **LokiStack** resource manages Loki, which is a scalable, highly-available, multi-tenant log aggregation system, and a web proxy with OpenShift Container Platform authentication. The **LokiStack** proxy uses OpenShift Container Platform authentication to enforce multi-tenancy and facilitate the saving and indexing of data in Loki log stores.



### NOTE

The Loki Operator can also be used for [configuring the LokiStack log store](#). The Network Observability Operator requires a dedicated LokiStack separate from the logging.

### 3.1. NETWORK OBSERVABILITY WITHOUT LOKI

You can use Network Observability without Loki by not performing the Loki installation steps and skipping directly to "Installing the Network Observability Operator". If you only want to export flows to a Kafka consumer or IPFIX collector, or you only need dashboard metrics, then you do not need to install Loki or provide storage for Loki. The following table compares available features with and without Loki.

Table 3.1. Comparison of feature availability with and without Loki

	With Loki	Without Loki
Exporters	✓	✓
Multi-tenancy	✓	✗
Complete filtering and aggregations capabilities <sup>[1]</sup>	✓	✗
Partial filtering and aggregations capabilities <sup>[2]</sup>	✓	✓
Flow-based metrics and dashboards	✓	✓
Traffic flows view overview <sup>[3]</sup>	✓	✓
Traffic flows view table	✓	✗
Topology view	✓	✓

	With Loki	Without Loki
OpenShift Container Platform console Network Traffic tab integration	✓	✓

1. Such as per pod.
2. Such as per workload or namespace.
3. Statistics on packet drops are only available with Loki.

#### Additional resources

- [Export enriched network flow data](#) .

## 3.2. INSTALLING THE LOKI OPERATOR

The [Loki Operator versions 5.7+](#) are the supported Loki Operator versions for Network Observability; these versions provide the ability to create a **LokiStack** instance using the **openshift-network** tenant configuration mode and provide fully-automatic, in-cluster authentication and authorization support for Network Observability. There are several ways you can install Loki. One way is by using the OpenShift Container Platform web console Operator Hub.

#### Prerequisites

- Supported Log Store (AWS S3, Google Cloud Storage, Azure, Swift, Minio, OpenShift Data Foundation)
- OpenShift Container Platform 4.10+
- Linux Kernel 4.18+

#### Procedure

1. In the OpenShift Container Platform web console, click **Operators** → **OperatorHub**.
2. Choose **Loki Operator** from the list of available Operators, and click **Install**.
3. Under **Installation Mode**, select **All namespaces on the cluster**.

#### Verification

1. Verify that you installed the Loki Operator. Visit the **Operators** → **Installed Operators** page and look for **Loki Operator**.
2. Verify that **Loki Operator** is listed with **Status** as **Succeeded** in all the projects.



## IMPORTANT

To uninstall Loki, refer to the uninstallation process that corresponds with the method you used to install Loki. You might have remaining **ClusterRoles** and **ClusterRoleBindings**, data stored in object store, and persistent volume that must be removed.

### 3.2.1. Creating a secret for Loki storage

The Loki Operator supports a few log storage options, such as AWS S3, Google Cloud Storage, Azure, Swift, Minio, OpenShift Data Foundation. The following example shows how to create a secret for AWS S3 storage. The secret created in this example, **loki-s3**, is referenced in "Creating a LokiStack resource". You can create this secret in the web console or CLI.

1. Using the web console, navigate to the **Project** → **All Projects** dropdown and select **Create Project**. Name the project **netobserv** and click **Create**.
2. Navigate to the Import icon, **+**, in the top right corner. Paste your YAML file into the editor. The following shows an example secret YAML file for S3 storage:

```
apiVersion: v1
kind: Secret
metadata:
  name: loki-s3
  namespace: netobserv 1
stringData:
  access_key_id: QUtJQUIPU0ZPRE5ON0VYQU1QTEUK
  access_key_secret:
d0phbHJYVXRuRkVNSS9LN01ERU5HL2JQeFJmaUNZRVhBTVBMRUtFWQo=
  bucketnames: s3-bucket-name
  endpoint: https://s3.eu-central-1.amazonaws.com
  region: eu-central-1
```

- 1** The installation examples in this documentation use the same namespace, **netobserv**, across all components. You can optionally use a different namespace for the different components

#### Verification

- Once you create the secret, you should see it listed under **Workloads** → **Secrets** in the web console.

#### Additional resources

- [Flow Collector API Reference](#)
- [Flow Collector sample resource](#)
- [Loki object storage](#)

### 3.2.2. Creating a LokiStack custom resource

You can deploy a **LokiStack** custom resource (CR) by using the web console or OpenShift CLI ( **oc** ) to create a namespace, or new project.

## Procedure

1. Navigate to **Operators** → **Installed Operators**, viewing **All projects** from the **Project** dropdown.
2. Look for **Loki Operator**. In the details, under **Provided APIs**, select **LokiStack**.
3. Click **Create LokiStack**
4. Ensure the following fields are specified in either **Form View** or **YAML view**:

```

apiVersion: loki.grafana.com/v1
kind: LokiStack
metadata:
  name: loki
  namespace: netobserv 1
spec:
  size: 1x.small 2
  storage:
    schemas:
      - version: v12
        effectiveDate: '2022-06-01'
    secret:
      name: loki-s3
      type: s3
  storageClassName: gp3 3
  tenants:
    mode: openshift-network

```

- 1** The installation examples in this documentation use the same namespace, **netobserv**, across all components. You can optionally use a different namespace.
- 2** Specify the deployment size. In the Loki Operator 5.8 and later versions, the supported size options for production instances of Loki are **1x.extra-small**, **1x.small**, or **1x.medium**.



### IMPORTANT

It is not possible to change the number **1x** for the deployment size.

- 3** Use a storage class name that is available on the cluster for **ReadWriteOnce** access mode. You can use **oc get storageclasses** to see what is available on your cluster.



### IMPORTANT

You must not reuse the same **LokiStack** CR that is used for logging.

5. Click **Create**.

### 3.2.3. Creating a new group for the cluster-admin user role



## IMPORTANT

Querying application logs for multiple namespaces as a **cluster-admin** user, where the sum total of characters of all of the namespaces in the cluster is greater than 5120, results in the error **Parse error: input size too long (XXXX > 5120)**. For better control over access to logs in LokiStack, make the **cluster-admin** user a member of the **cluster-admin** group. If the **cluster-admin** group does not exist, create it and add the desired users to it.

Use the following procedure to create a new group for users with **cluster-admin** permissions.

### Procedure

1. Enter the following command to create a new group:

```
$ oc adm groups new cluster-admin
```

2. Enter the following command to add the desired user to the **cluster-admin** group:

```
$ oc adm groups add-users cluster-admin <username>
```

3. Enter the following command to add **cluster-admin** user role to the group:

```
$ oc adm policy add-cluster-role-to-group cluster-admin cluster-admin
```

### 3.2.4. Custom admin group access

If you have a large deployment with a number of users who require broader permissions, you can create a custom group using the **adminGroup** field. Users who are members of any group specified in the **adminGroups** field of the **LokiStack** CR are considered admins. Admin users have access to all application logs in all namespaces, if they also get assigned the **cluster-logging-application-view** role.


#### Example LokiStack CR

```
apiVersion: loki.grafana.com/v1
kind: LokiStack
metadata:
  name: logging-loki
  namespace: openshift-logging
spec:
  tenants:
    mode: openshift-network 1
  openshift:
    adminGroups: 2
    - cluster-admin
    - custom-admin-group 3
```

- 1 Custom admin groups are only available in this mode.
- 2 Entering an empty list [] value for this field disables admin groups.
- 3 Overrides the default groups (**system:cluster-admins**, **cluster-admin**, **dedicated-admin**)

### 3.2.5. Loki deployment sizing

Sizing for Loki follows the format of **1x.<size>** where the value **1x** is number of instances and **<size>** specifies performance capabilities.



#### IMPORTANT

It is not possible to change the number **1x** for the deployment size.

Table 3.2. Loki sizing

	1x.demo	1x.extra-small	1x.small	1x.medium
Data transfer	Demo use only	100GB/day	500GB/day	2TB/day
Queries per second (QPS)	Demo use only	1-25 QPS at 200ms	25-50 QPS at 200ms	25-75 QPS at 200ms
Replication factor	None	2	2	2
Total CPU requests	None	14 vCPUs	34 vCPUs	54 vCPUs
Total memory requests	None	31Gi	67Gi	139Gi
Total disk requests	40Gi	430Gi	430Gi	590Gi

### 3.2.6. LokiStack ingestion limits and health alerts

The LokiStack instance comes with default settings according to the configured size. It is possible to override some of these settings, such as the ingestion and query limits. You might want to update them if you get Loki errors showing up in the Console plugin, or in **flowlogs-pipeline** logs. An automatic alert in the web console notifies you when these limits are reached.

Here is an example of configured limits:

```
spec:
  limits:
    global:
      ingestion:
        ingestionBurstSize: 40
        ingestionRate: 20
        maxGlobalStreamsPerTenant: 25000
      queries:
        maxChunksPerQuery: 2000000
        maxEntriesLimitPerQuery: 10000
        maxQuerySeries: 3000
```

For more information about these settings, see the [LokiStack API reference](#).



### 3.2.7. Enabling multi-tenancy in Network Observability

Multi-tenancy in the Network Observability Operator allows and restricts individual user access, or group access, to the flows stored in Loki. Access is enabled for project admins. Project admins who have limited access to some namespaces can access flows for only those namespaces.

#### Prerequisite

- You have installed at least [Loki Operator version 5.7](#)
- You must be logged in as a project administrator

#### Procedure

1. Authorize reading permission to **user1** by running the following command:

```
$ oc adm policy add-cluster-role-to-user netobserv-reader user1
```

Now, the data is restricted to only allowed user namespaces. For example, a user that has access to a single namespace can see all the flows internal to this namespace, as well as flows going from and to this namespace. Project admins have access to the Administrator perspective in the OpenShift Container Platform console to access the Network Flows Traffic page.

## 3.3. INSTALLING THE NETWORK OBSERVABILITY OPERATOR

You can install the Network Observability Operator using the OpenShift Container Platform web console Operator Hub. When you install the Operator, it provides the **FlowCollector** custom resource definition (CRD). You can set specifications in the web console when you create the **FlowCollector**.



#### IMPORTANT

The actual memory consumption of the Operator depends on your cluster size and the number of resources deployed. Memory consumption might need to be adjusted accordingly. For more information refer to "Network Observability controller manager pod runs out of memory" in the "Important Flow Collector configuration considerations" section.

#### Prerequisites

- If you choose to use Loki, install the [Loki Operator version 5.7+](#).
- You must have **cluster-admin** privileges.
- One of the following supported architectures is required: **amd64**, **ppc64le**, **arm64**, or **s390x**.
- Any CPU supported by Red Hat Enterprise Linux (RHEL) 9.
- Must be configured with OVN-Kubernetes or OpenShift SDN as the main network plugin, and optionally using secondary interfaces with Multus and SR-IOV.



#### NOTE

Additionally, this installation example uses the **netobserv** namespace, which is used across all components. You can optionally use a different namespace.

## Procedure

1. In the OpenShift Container Platform web console, click **Operators** → **OperatorHub**.
2. Choose **Network Observability Operator** from the list of available Operators in the **OperatorHub**, and click **Install**.
3. Select the checkbox **Enable Operator recommended cluster monitoring on this Namespace**.
4. Navigate to **Operators** → **Installed Operators**. Under Provided APIs for Network Observability, select the **Flow Collector** link.
5. Navigate to the **Flow Collector** tab, and click **Create FlowCollector**. Make the following selections in the form view:
  - a. **spec.agent.ebpf.Sampling**: Specify a sampling size for flows. Lower sampling sizes will have higher impact on resource utilization. For more information, see the "FlowCollector API reference", **spec.agent.ebpf**.
  - b. If you are using Loki, set the following specifications:
    - i. **spec.loki.mode**: Set this to the **LokiStack** mode, which automatically sets URLs, TLS, cluster roles and a cluster role binding, as well as the **authToken** value. Alternatively, the **Manual** mode allows more control over configuration of these settings.
    - ii. **spec.loki.loki.name**: Set this to the name of your **LokiStack** resource. In this documentation, **loki** is used.
  - c. Optional: If you are in a large-scale environment, consider configuring the **FlowCollector** with Kafka for forwarding data in a more resilient, scalable way. See "Configuring the Flow Collector resource with Kafka storage" in the "Important Flow Collector configuration considerations" section.
  - d. Optional: Configure other optional settings before the next step of creating the **FlowCollector**. For example, if you choose not to use Loki, then you can configure exporting flows to Kafka or IPFIX. See "Export enriched network flow data to Kafka and IPFIX" and more in the "Important Flow Collector configuration considerations" section.
6. Click **Create**.

## Verification

To confirm this was successful, when you navigate to **Observe** you should see **Network Traffic** listed in the options.

In the absence of **Application Traffic** within the OpenShift Container Platform cluster, default filters might show that there are "No results", which results in no visual flow. Beside the filter selections, select **Clear all filters** to see the flow.

## 3.4. IMPORTANT FLOW COLLECTOR CONFIGURATION CONSIDERATIONS

Once you create the **FlowCollector** instance, you can reconfigure it, but the pods are terminated and recreated again, which can be disruptive. Therefore, you can consider configuring the following options when creating the **FlowCollector** for the first time:

- [Configuring the Flow Collector resource with Kafka](#)

- [Export enriched network flow data to Kafka or IPFIX](#)
- [Configuring monitoring for SR-IOV interface traffic](#)
- [Working with conversation tracking](#)
- [Working with DNS tracking](#)

### Additional resources

For more general information about Flow Collector specifications and the Network Observability Operator architecture and resource use, see the following resources:

- [Flow Collector API Reference](#)
- [Flow Collector sample resource](#)
- [Resource considerations](#)
- [Troubleshooting Network Observability controller manager pod runs out of memory](#)
- [Network Observability architecture](#)

## 3.5. INSTALLING KAFKA (OPTIONAL)

The Kafka Operator is supported for large scale environments. Kafka provides high-throughput and low-latency data feeds for forwarding network flow data in a more resilient, scalable way. You can install the Kafka Operator as [Red Hat AMQ Streams](#) from the Operator Hub, just as the Loki Operator and Network Observability Operator were installed. Refer to "Configuring the FlowCollector resource with Kafka" to configure Kafka as a storage option.



### NOTE

To uninstall Kafka, refer to the uninstallation process that corresponds with the method you used to install.


### Additional resources



[Configuring the FlowCollector resource with Kafka](#).

## 3.6. UNINSTALLING THE NETWORK OBSERVABILITY OPERATOR

You can uninstall the Network Observability Operator using the OpenShift Container Platform web console Operator Hub, working in the **Operators** → **Installed Operators** area.

### Procedure

1. Remove the **FlowCollector** custom resource.
  - a. Click **Flow Collector**, which is next to the **Network Observability Operator** in the **Provided APIs** column.
  - b. Click the options menu  for the **cluster** and select **Delete FlowCollector**.

2. Uninstall the Network Observability Operator.
  - a. Navigate back to the **Operators → Installed Operators** area.
  - b. Click the options menu  next to the **Network Observability Operator** and select **Uninstall Operator**.
  - c. **Home → Projects** and select **openshift-netobserv-operator**
  - d. Navigate to **Actions** and select **Delete Project**
3. Remove the **FlowCollector** custom resource definition (CRD).
  - a. Navigate to **Administration → CustomResourceDefinitions**.
  - b. Look for **FlowCollector** and click the options menu  .
  - c. Select **Delete CustomResourceDefinition**.



#### IMPORTANT

The Loki Operator and Kafka remain if they were installed and must be removed separately. Additionally, you might have remaining data stored in an object store, and a persistent volume that must be removed.

## CHAPTER 4. NETWORK OBSERVABILITY OPERATOR IN OPENSIFT CONTAINER PLATFORM

Network Observability is an OpenShift operator that deploys a monitoring pipeline to collect and enrich network traffic flows that are produced by the Network Observability eBPF agent.

### 4.1. VIEWING STATUSES

The Network Observability Operator provides the Flow Collector API. When a Flow Collector resource is created, it deploys pods and services to create and store network flows in the Loki log store, as well as to display dashboards, metrics, and flows in the OpenShift Container Platform web console.

#### Procedure

1. Run the following command to view the state of **FlowCollector**:

```
$ oc get flowcollector/cluster
```

#### Example output

NAME	AGENT	SAMPLING (EBPF)	DEPLOYMENT MODEL	STATUS
cluster	EBPF	50	DIRECT	Ready

2. Check the status of pods running in the **netobserv** namespace by entering the following command:

```
$ oc get pods -n netobserv
```

#### Example output

NAME	READY	STATUS	RESTARTS	AGE
flowlogs-pipeline-56hbp	1/1	Running	0	147m
flowlogs-pipeline-9plv	1/1	Running	0	147m
flowlogs-pipeline-h5gkb	1/1	Running	0	147m
flowlogs-pipeline-hh6kf	1/1	Running	0	147m
flowlogs-pipeline-w7vv5	1/1	Running	0	147m
netobserv-plugin-cdd7dc6c-j8ggp	1/1	Running	0	147m

**flowlogs-pipeline** pods collect flows, enriches the collected flows, then send flows to the Loki storage. **netobserv-plugin** pods create a visualization plugin for the OpenShift Container Platform Console.

1. Check the status of pods running in the namespace **netobserv-privileged** by entering the following command:

```
$ oc get pods -n netobserv-privileged
```

#### Example output

NAME	READY	STATUS	RESTARTS	AGE
netobserv-ebpf-agent-4lpp6	1/1	Running	0	151m
netobserv-ebpf-agent-6gbrk	1/1	Running	0	151m

```
netobserv-ebpf-agent-klpl9 1/1 Running 0 151m
netobserv-ebpf-agent-vrcnf 1/1 Running 0 151m
netobserv-ebpf-agent-xf5jh 1/1 Running 0 151m
```

**netobserv-ebpf-agent** pods monitor network interfaces of the nodes to get flows and send them to **flowlogs-pipeline** pods.

1. If you are using the Loki Operator, check the status of pods running in the **openshift-operators-redhat** namespace by entering the following command:

```
$ oc get pods -n openshift-operators-redhat
```

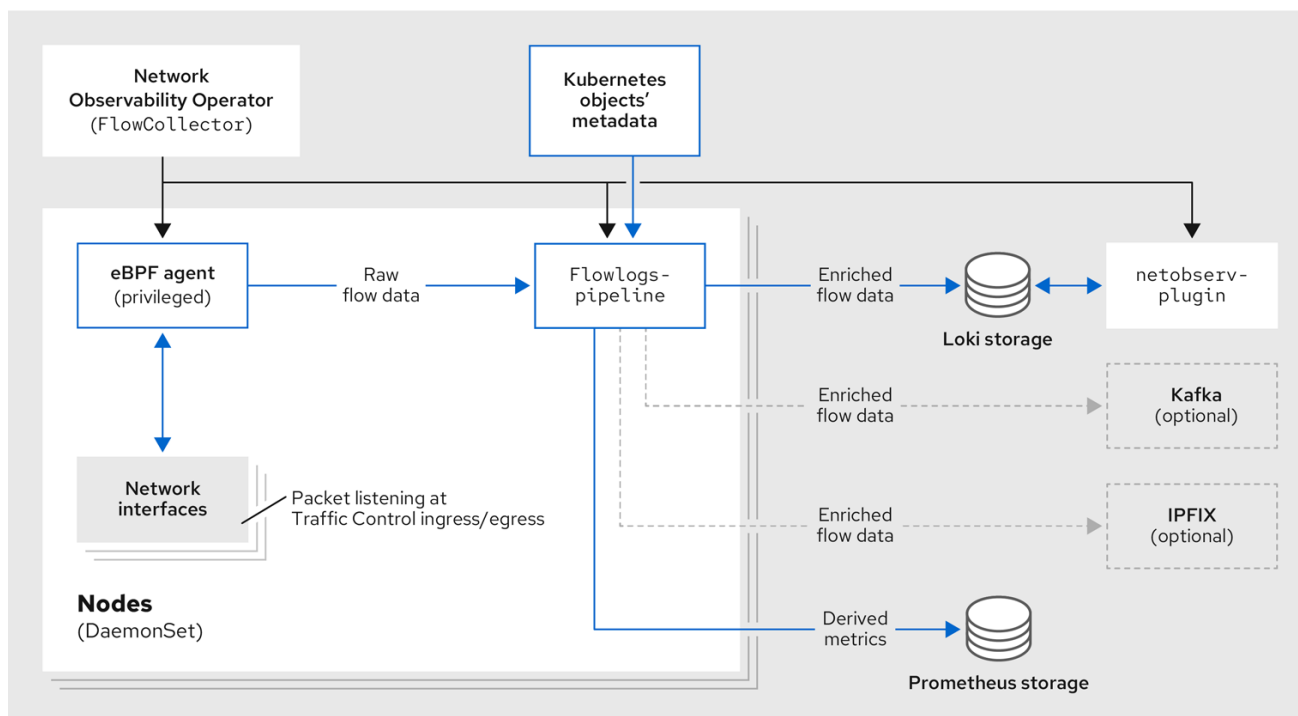
### Example output

```
NAME                                READY STATUS RESTARTS AGE
loki-operator-controller-manager-5f6cff4f9d-jq25h 2/2 Running 0 18h
lokistack-compactor-0                1/1 Running 0 18h
lokistack-distributor-654f87c5bc-qhkhv          1/1 Running 0 18h
lokistack-distributor-654f87c5bc-skxgm          1/1 Running 0 18h
lokistack-gateway-796dc6ff7-c54gz              2/2 Running 0 18h
lokistack-index-gateway-0                 1/1 Running 0 18h
lokistack-index-gateway-1                 1/1 Running 0 18h
lokistack-ingester-0                     1/1 Running 0 18h
lokistack-ingester-1                     1/1 Running 0 18h
lokistack-ingester-2                     1/1 Running 0 18h
lokistack-querier-66747dc666-6vh5x          1/1 Running 0 18h
lokistack-querier-66747dc666-cjr45          1/1 Running 0 18h
lokistack-querier-66747dc666-xh8rq          1/1 Running 0 18h
lokistack-query-frontend-85c6db4fbd-b2xfb     1/1 Running 0 18h
lokistack-query-frontend-85c6db4fbd-jm94f     1/1 Running 0 18h
```

## 4.2. NETWORK OBSERVABILITY OPERATOR ARCHITECTURE

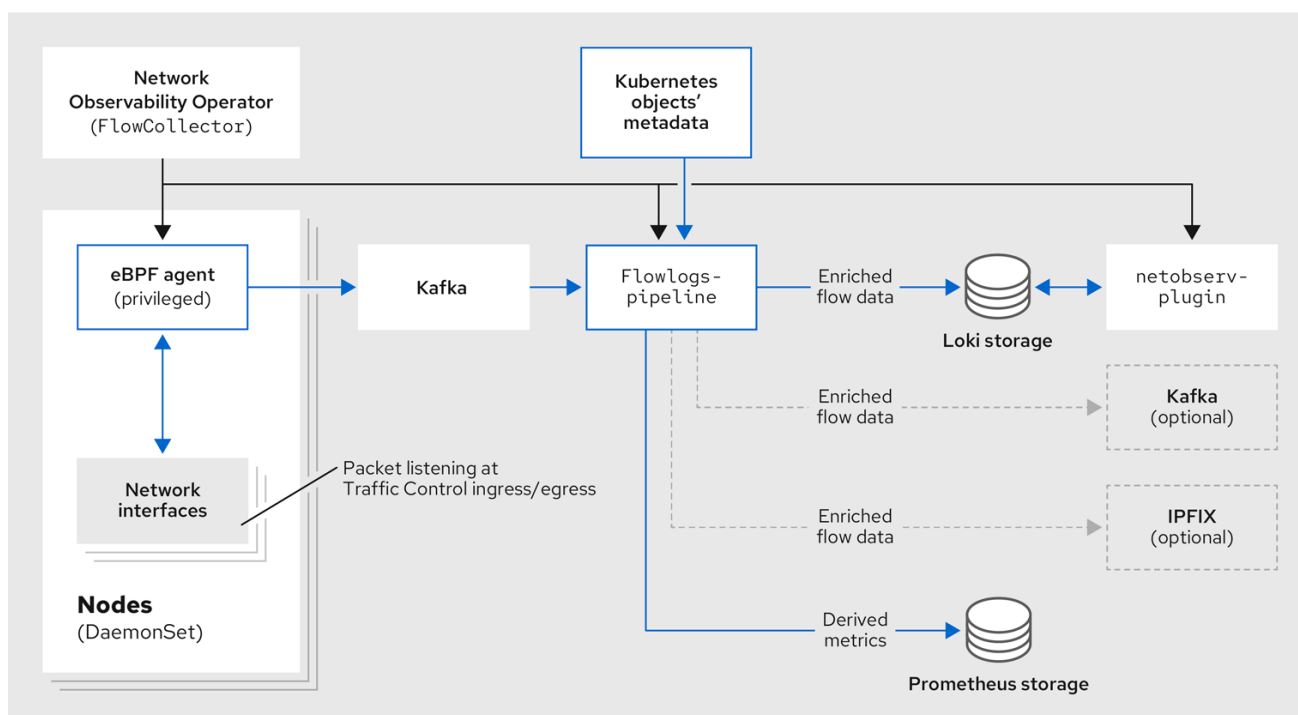
The Network Observability Operator provides the **FlowCollector** API, which is instantiated at installation and configured to reconcile the **eBPF agent**, the **flowlogs-pipeline**, and the **netobserv-plugin** components. Only a single **FlowCollector** per cluster is supported.

The **eBPF agent** runs on each cluster node with some privileges to collect network flows. The **flowlogs-pipeline** receives the network flows data and enriches the data with Kubernetes identifiers. If you are using Loki, the **flowlogs-pipeline** sends flow logs data to Loki for storing and indexing. The **netobserv-plugin**, which is a dynamic OpenShift Container Platform web console plugin, queries Loki to fetch network flows data. Cluster-admins can view the data in the web console.



351\_OpenShift\_0823

If you are using the Kafka option, the eBPF agent sends the network flow data to Kafka, and the **flowlogs-pipeline** reads from the Kafka topic before sending to Loki, as shown in the following diagram.



351\_OpenShift\_0823

### 4.3. VIEWING NETWORK OBSERVABILITY OPERATOR STATUS AND CONFIGURATION

You can inspect the status and view the details of the **FlowCollector** using the **oc describe** command.

## Procedure

1. Run the following command to view the status and configuration of the Network Observability Operator:

```
❯ $ oc describe flowcollector/cluster
```



## CHAPTER 5. CONFIGURING THE NETWORK OBSERVABILITY OPERATOR

You can update the Flow Collector API resource to configure the Network Observability Operator and its managed components. The Flow Collector is explicitly created during installation. Since this resource operates cluster-wide, only a single **FlowCollector** is allowed, and it has to be named **cluster**.

### 5.1. VIEW THE FLOWCOLLECTOR RESOURCE

You can view and edit YAML directly in the OpenShift Container Platform web console.

#### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
3. Select **cluster** then select the **YAML** tab. There, you can modify the **FlowCollector** resource to configure the Network Observability operator.

The following example shows a sample **FlowCollector** resource for OpenShift Container Platform Network Observability operator:

#### Sample FlowCollector resource

```

apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  deploymentModel: Direct
  agent:
    type: eBPF 1
    ebf:
      sampling: 50 2
      logLevel: info
      privileged: false
      resources:
        requests:
          memory: 50Mi
          cpu: 100m
        limits:
          memory: 800Mi
  processor: 3
    logLevel: info
    resources:
      requests:
        memory: 100Mi
        cpu: 100m
      limits:
        memory: 800Mi
  logTypes: Flows
  advanced:

```

```

    conversationEndTimeout: 10s
    conversationHeartbeatInterval: 30s
loki:
    mode: LokiStack
consolePlugin:
    register: true
    logLevel: info
    portNaming:
        enable: true
        portNames:
            "3100": loki
quickFilters:
- name: Applications
  filter:
    src_namespace!: 'openshift-,netobserv'
    dst_namespace!: 'openshift-,netobserv'
    default: true
- name: Infrastructure
  filter:
    src_namespace: 'openshift-,netobserv'
    dst_namespace: 'openshift-,netobserv'
- name: Pods network
  filter:
    src_kind: 'Pod'
    dst_kind: 'Pod'
    default: true
- name: Services network
  filter:
    dst_kind: 'Service'

```

- 1 The Agent specification, **spec.agent.type**, must be **EBPF**. eBPF is the only OpenShift Container Platform supported option.
- 2 You can set the Sampling specification, **spec.agent.ebpf.sampling**, to manage resources. Lower sampling values might consume a large amount of computational, memory and storage resources. You can mitigate this by specifying a sampling ratio value. A value of 100 means 1 flow every 100 is sampled. A value of 0 or 1 means all flows are captured. The lower the value, the increase in returned flows and the accuracy of derived metrics. By default, eBPF sampling is set to a value of 50, so 1 flow every 50 is sampled. Note that more sampled flows also means more storage needed. It is recommend to start with default values and refine empirically, to determine which setting your cluster can manage.
- 3 The Processor specification **spec.processor**. can be set to enable conversation tracking. When enabled, conversation events are queryable in the web console. The **spec.processor.logTypes** value is **Flows**. The **spec.processor.advanced** values are **Conversations**, **EndedConversations**, or **ALL**. Storage requirements are highest for **All** and lowest for **EndedConversations**.
- 4 The Loki specification, **spec.loki**, specifies the Loki client. The default values match the Loki install paths mentioned in the Installing the Loki Operator section. If you used another installation method for Loki, specify the appropriate client information for your install.
- 5 The **LokiStack** mode automatically sets a few configurations: **querierUrl**, **ingesterUrl** and **statusUrl**, **tenantID**, and corresponding TLS configuration. Cluster roles and a cluster role binding are created for reading and writing logs to Loki. And **authToken** is set to **Forward**. You can set these manually using the **Manual** mode.

- 6 The **spec.quickFilters** specification defines filters that show up in the web console. The **Application** filter keys, **src\_namespace** and **dst\_namespace**, are negated (!), so the **Application**

### Additional resources

For more information about conversation tracking, see [Working with conversations](#).

## 5.2. CONFIGURING THE FLOW COLLECTOR RESOURCE WITH KAFKA

You can configure the **FlowCollector** resource to use Kafka for high-throughput and low-latency data feeds. A Kafka instance needs to be running, and a Kafka topic dedicated to OpenShift Container Platform Network Observability must be created in that instance. For more information, see [Kafka documentation with AMQ Streams](#).

### Prerequisites

- Kafka is installed. Red Hat supports Kafka with AMQ Streams Operator.

### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for the Network Observability Operator, select **Flow Collector**.
3. Select the cluster and then click the **YAML** tab.
4. Modify the **FlowCollector** resource for OpenShift Container Platform Network Observability Operator to use Kafka, as shown in the following sample YAML:

### Sample Kafka configuration in **FlowCollector** resource

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  deploymentModel: Kafka
  kafka:
    address: "kafka-cluster-kafka-bootstrap.netobserv"
    topic: network-flows
    tls:
      enable: false
```

- 1 Set **spec.deploymentModel** to **Kafka** instead of **Direct** to enable the Kafka deployment model.
- 2 **spec.kafka.address** refers to the Kafka bootstrap server address. You can specify a port if needed, for instance **kafka-cluster-kafka-bootstrap.netobserv:9093** for using TLS on port 9093.
- 3 **spec.kafka.topic** should match the name of a topic created in Kafka.
- 4

**spec.kafka.tls** can be used to encrypt all communications to and from Kafka with TLS or mTLS. When enabled, the Kafka CA certificate must be available as a ConfigMap or a Secret, both in the

## 5.3. EXPORT ENRICHED NETWORK FLOW DATA

You can send network flows to Kafka, IPFIX, or both at the same time. Any processor or storage that supports Kafka or IPFIX input, such as Splunk, Elasticsearch, or Fluentd, can consume the enriched network flow data.

### Prerequisites

- Your Kafka or IPFIX collector endpoint(s) are available from Network Observability **flowlogs-pipeline** pods.

### Procedure

- In the web console, navigate to **Operators** → **Installed Operators**.
- Under the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
- Select **cluster** and then select the **YAML** tab.
- Edit the **FlowCollector** to configure **spec.exporters** as follows:

```

apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  exporters:
    - type: Kafka 1
      kafka:
        address: "kafka-cluster-kafka-bootstrap.netobserv"
        topic: netobserv-flows-export 2
        tls:
          enable: false 3
    - type: IPFIX 4
      ipfix:
        targetHost: "ipfix-collector.ipfix.svc.cluster.local"
        targetPort: 4739
        transport: tcp or udp 5

```

- The Network Observability Operator exports all flows to the configured Kafka topic.
- You can encrypt all communications to and from Kafka with SSL/TLS or mTLS. When enabled, the Kafka CA certificate must be available as a ConfigMap or a Secret, both in the namespace where the **flowlogs-pipeline** processor component is deployed (default: netobserv). It must be referenced with **spec.exporters.tls.caCert**. When using mTLS, client secrets must be available in these namespaces as well (they can be generated for instance using the AMQ Streams User Operator) and referenced with **spec.exporters.tls.userCert**.
- 1 4 You can export flows to IPFIX instead of or in conjunction with exporting flows to Kafka.

- 5 You have the option to specify transport. The default value is **tcp** but you can also specify **udp**.

5. After configuration, network flows data can be sent to an available output in a JSON format. For more information, see *Network flows format reference*.

### Additional resources

For more information about specifying flow format, see [Network flows format reference](#).

## 5.4. UPDATING THE FLOW COLLECTOR RESOURCE

As an alternative to editing YAML in the OpenShift Container Platform web console, you can configure specifications, such as eBPF sampling, by patching the **flowcollector** custom resource (CR):

### Procedure

1. Run the following command to patch the **flowcollector** CR and update the **spec.agent.ebpf.sampling** value:

```
$ oc patch flowcollector cluster --type=json -p [{"op": "replace", "path":
"/spec/agent/ebpf/sampling", "value": <new value>}] -n netobserv
```

## 5.5. CONFIGURING QUICK FILTERS

You can modify the filters in the **FlowCollector** resource. Exact matches are possible using double-quotes around values. Otherwise, partial matches are used for textual values. The bang (!) character, placed at the end of a key, means negation. See the sample **FlowCollector** resource for more context about modifying the YAML.



### NOTE

The filter matching types "all of" or "any of" is a UI setting that the users can modify from the query options. It is not part of this resource configuration.

Here is a list of all available filter keys:

Table 5.1. Filter keys

Universe*	Source	Destination	Description
namespace	<b>src_namespace</b>	<b>dst_namespace</b>	Filter traffic related to a specific namespace.
name	<b>src_name</b>	<b>dst_name</b>	Filter traffic related to a given leaf resource name, such as a specific pod, service, or node (for host-network traffic).

Universal*	Source	Destination	Description
kind	<b>src_kind</b>	<b>dst_kind</b>	Filter traffic related to a given resource kind. The resource kinds include the leaf resource (Pod, Service or Node), or the owner resource (Deployment and StatefulSet).
owner_name	<b>src_owner_name</b>	<b>dst_owner_name</b>	Filter traffic related to a given resource owner; that is, a workload or a set of pods. For example, it can be a Deployment name, a StatefulSet name, etc.
resource	<b>src_resource</b>	<b>dst_resource</b>	Filter traffic related to a specific resource that is denoted by its canonical name, that identifies it uniquely. The canonical notation is <b>kind.namespace.name</b> for namespaced kinds, or <b>node.name</b> for nodes. For example, <b>Deployment.my-namespace.my-web-server</b> .
address	<b>src_address</b>	<b>dst_address</b>	Filter traffic related to an IP address. IPv4 and IPv6 are supported. CIDR ranges are also supported.
mac	<b>src_mac</b>	<b>dst_mac</b>	Filter traffic related to a MAC address.
port	<b>src_port</b>	<b>dst_port</b>	Filter traffic related to a specific port.
host_addresses	<b>src_host_address</b>	<b>dst_host_address</b>	Filter traffic related to the host IP address where the pods are running.
protocol	N/A	N/A	Filter traffic related to a protocol, such as TCP or UDP.

- Universal keys filter for any of source or destination. For example, filtering **name: 'my-pod'** means all traffic from **my-pod** and all traffic to **my-pod**, regardless of the matching type used, whether **Match all** or **Match any**.

## 5.6. CONFIGURING MONITORING FOR SR-IOV INTERFACE TRAFFIC

In order to collect traffic from a cluster with a Single Root I/O Virtualization (SR-IOV) device, you must set the **FlowCollector spec.agent.ebpf.privileged** field to **true**. Then, the eBPF agent monitors other network namespaces in addition to the host network namespaces, which are monitored by default. When a pod with a virtual functions (VF) interface is created, a new network namespace is created. With **SRIOVNetwork** policy **IPAM** configurations specified, the VF interface is migrated from the host network namespace to the pod network namespace.

### Prerequisites

- Access to an OpenShift Container Platform cluster with a SR-IOV device.

- The **SRIOVNetwork** custom resource (CR) **spec.ipam** configuration must be set with an IP address from the range that the interface lists or from other plugins.

## Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
3. Select **cluster** and then select the **YAML** tab.
4. Configure the **FlowCollector** custom resource. A sample configuration is as follows:

## Configure FlowCollector for SR-IOV monitoring

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  deploymentModel: Direct
  agent:
    type: eBPF
    ebpf:
      privileged: true 1
```

- 1 The **spec.agent.ebpf.privileged** field value must be set to **true** to enable SR-IOV monitoring.

## Additional resources

For more information about creating the **SriovNetwork** custom resource, see [Creating an additional SR-IOV network attachment with the CNI VRF plugin](#).

## 5.7. RESOURCE MANAGEMENT AND PERFORMANCE CONSIDERATIONS

The amount of resources required by Network Observability depends on the size of your cluster and your requirements for the cluster to ingest and store observability data. To manage resources and set performance criteria for your cluster, consider configuring the following settings. Configuring these settings might meet your optimal setup and observability needs.

The following settings can help you manage resources and performance from the outset:

### eBPF Sampling

You can set the Sampling specification, **spec.agent.ebpf.sampling**, to manage resources. Smaller sampling values might consume a large amount of computational, memory and storage resources. You can mitigate this by specifying a sampling ratio value. A value of **100** means 1 flow every 100 is sampled. A value of **0** or **1** means all flows are captured. Smaller values result in an increase in returned flows and the accuracy of derived metrics. By default, eBPF sampling is set to a value of 50, so 1 flow every 50 is sampled. Note that more sampled flows also means more storage needed. Consider starting with the default values and refine empirically, in order to determine which setting your cluster can manage.

## Restricting or excluding interfaces

Reduce the overall observed traffic by setting the values for **spec.agent.ebpf.interfaces** and **spec.agent.ebpf.excludeInterfaces**. By default, the agent fetches all the interfaces in the system, except the ones listed in **excludeInterfaces** and **lo** (local interface). Note that the interface names might vary according to the Container Network Interface (CNI) used.

The following settings can be used to fine-tune performance after the Network Observability has been running for a while:

## Resource requirements and limits

Adapt the resource requirements and limits to the load and memory usage you expect on your cluster by using the **spec.agent.ebpf.resources** and **spec.processor.resources** specifications. The default limits of 800MB might be sufficient for most medium-sized clusters.

## Cache max flows timeout

Control how often flows are reported by the agents by using the eBPF agent's **spec.agent.ebpf.cacheMaxFlows** and **spec.agent.ebpf.cacheActiveTimeout** specifications. A larger value results in less traffic being generated by the agents, which correlates with a lower CPU load. However, a larger value leads to a slightly higher memory consumption, and might generate more latency in the flow collection.

## 5.7.1. Resource considerations

The following table outlines examples of resource considerations for clusters with certain workload sizes.



### IMPORTANT

The examples outlined in the table demonstrate scenarios that are tailored to specific workloads. Consider each example only as a baseline from which adjustments can be made to accommodate your workload needs.

Table 5.2. Resource recommendations

	Extra small (10 nodes)	Small (25 nodes)	Medium (65 nodes) [2]	Large (120 nodes) [2]
Worker Node vCPU and memory	4 vCPUs  16GiB mem [1]	16 vCPUs  64GiB mem [1]	16 vCPUs  64GiB mem [1]	16 vCPUs  64GiB Mem [1]
LokiStack size	<b>1x.extra-small</b>	<b>1x.small</b>	<b>1x.small</b>	<b>1x.medium</b>
Network Observability controller memory limit	400Mi (default)	400Mi (default)	400Mi (default)	400Mi (default)
eBPF sampling rate	50 (default)	50 (default)	50 (default)	50 (default)



	Extra small (10 nodes)	Small (25 nodes)	Medium (65 nodes) [2]	Large (120 nodes) [2]
<b>eBPF memory limit</b>	800Mi (default)	800Mi (default)	800Mi (default)	1600Mi
<b>FLP memory limit</b>	800Mi (default)	800Mi (default)	800Mi (default)	800Mi (default)
<b>FLP Kafka partitions</b>	N/A	48	48	48
<b>Kafka consumer replicas</b>	N/A	24	24	24
<b>Kafka brokers</b>	N/A	3 (default)	3 (default)	3 (default)

1. Tested with AWS M6i instances.
2. In addition to this worker and its controller, 3 infra nodes (size **M6i.12xlarge**) and 1 workload node (size **M6i.8xlarge**) were tested.

## CHAPTER 6. NETWORK POLICY

As a user with the **admin** role, you can create a network policy for the **netobserv** namespace to secure inbound access to the Network Observability Operator.

### 6.1. CREATING A NETWORK POLICY FOR NETWORK OBSERVABILITY

You might need to create a network policy to secure ingress traffic to the **netobserv** namespace. In the web console, you can create a network policy using the form view.

#### Procedure

1. Navigate to **Networking** → **NetworkPolicies**.
2. Select the **netobserv** project from the **Project** dropdown menu.
3. Name the policy. For this example, the policy name is **allow-ingress**.
4. Click **Add ingress rule** three times to create three ingress rules.
5. Specify the following in the form:
  - a. Make the following specifications for the first **Ingress rule**:
    - i. From the **Add allowed source** dropdown menu, select **Allow pods from the same namespace**.
  - b. Make the following specifications for the second **Ingress rule**:
    - i. From the **Add allowed source** dropdown menu, select **Allow pods from inside the cluster**.
    - ii. Click + **Add namespace selector**.
    - iii. Add the label, **kubernetes.io/metadata.name**, and the selector, **openshift-console**.
  - c. Make the following specifications for the third **Ingress rule**:
    - i. From the **Add allowed source** dropdown menu, select **Allow pods from inside the cluster**.
    - ii. Click + **Add namespace selector**.
    - iii. Add the label, **kubernetes.io/metadata.name**, and the selector, **openshift-monitoring**.

#### Verification

1. Navigate to **Observe** → **Network Traffic**.
2. View the **Traffic Flows** tab, or any tab, to verify that the data is displayed.
3. Navigate to **Observe** → **Dashboards**. In the NetObserv/Health selection, verify that the flows are being ingested and sent to Loki, which is represented in the first graph.

### 6.2. EXAMPLE NETWORK POLICY

The following annotates an example **NetworkPolicy** object for the **netobserv** namespace:

### Sample network policy

```
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: allow-ingress
  namespace: netobserv
spec:
  podSelector: {} 1
  ingress:
    - from:
      - podSelector: {} 2
        namespaceSelector: 3
          matchLabels:
            kubernetes.io/metadata.name: openshift-console
      - podSelector: {}
        namespaceSelector:
          matchLabels:
            kubernetes.io/metadata.name: openshift-monitoring
  policyTypes:
    - Ingress
status: {}
```

- 1** A selector that describes the pods to which the policy applies. The policy object can only select pods in the project that defines the **NetworkPolicy** object. In this documentation, it would be the project in which the Network Observability Operator is installed, which is the **netobserv** project.
- 2** A selector that matches the pods from which the policy object allows ingress traffic. The default is that the selector matches pods in the same namespace as the **NetworkPolicy**.
- 3** When the **namespaceSelector** is specified, the selector matches pods in the specified namespace.

### Additional resources

[Creating a network policy using the CLI](#)

## CHAPTER 7. OBSERVING THE NETWORK TRAFFIC

As an administrator, you can observe the network traffic in the OpenShift Container Platform console for detailed troubleshooting and analysis. This feature helps you get insights from different graphical representations of traffic flow. There are several available views to observe the network traffic.

### 7.1. OBSERVING THE NETWORK TRAFFIC FROM THE OVERVIEW VIEW

The **Overview** view displays the overall aggregated metrics of the network traffic flow on the cluster. As an administrator, you can monitor the statistics with the available display options.

#### 7.1.1. Working with the Overview view

As an administrator, you can navigate to the **Overview** view to see the graphical representation of the flow rate statistics.

##### Procedure

1. Navigate to **Observe** → **Network Traffic**.
2. In the **Network Traffic** page, click the **Overview** tab.

You can configure the scope of each flow rate data by clicking the menu icon.

#### 7.1.2. Configuring advanced options for the Overview view

You can customize the graphical view by using advanced options. To access the advanced options, click **Show advanced options**. You can configure the details in the graph by using the **Display options** drop-down menu. The options available are as follows:

- **Scope:** Select to view the components that network traffic flows between. You can set the scope to **Node**, **Namespace**, **Owner**, **Zones**, **Cluster** or **Resource**. **Owner** is an aggregation of resources. **Resource** can be a pod, service, node, in case of host-network traffic, or an unknown IP address. The default value is **Namespace**.
- **Truncate labels:** Select the required width of the label from the drop-down list. The default value is **M**.

##### 7.1.2.1. Managing panels and display

You can select the required panels to be displayed, reorder them, and focus on a specific panel. To add or remove panels, click **Manage panels**.

The following panels are shown by default:

- **Top X average bytes rates**
- **Top X bytes rates stacked with total**

Other panels can be added in **Manage panels**:

- **Top X average packets rates**
- **Top X packets rates stacked with total**

**Query options** allows you to choose whether to show the **Top 5**, **Top 10**, or **Top 15** rates.

### 7.1.2.2. DNS tracking

You can configure graphical representation of Domain Name System (DNS) tracking of network flows in the **Overview** view. Using DNS tracking with extended Berkeley Packet Filter (eBPF) tracepoint hooks can serve various purposes:

- **Network Monitoring:** Gain insights into DNS queries and responses, helping network administrators identify unusual patterns, potential bottlenecks, or performance issues.
- **Security Analysis:** Detect suspicious DNS activities, such as domain name generation algorithms (DGA) used by malware, or identify unauthorized DNS resolutions that might indicate a security breach.
- **Troubleshooting:** Debug DNS-related issues by tracing DNS resolution steps, tracking latency, and identifying misconfigurations.

By default, when DNS tracking is enabled, you can see the following non-empty metrics represented in a donut or line chart in the **Overview**:

- Top X DNS Response Code
- Top X average DNS latencies with overall
- Top X 90th percentile DNS latencies

Other DNS tracking panels can be added in **Manage panels**:

- Bottom X minimum DNS latencies
- Top X maximum DNS latencies
- Top X 99th percentile DNS latencies

This feature is supported for IPv4 and IPv6 UDP and TCP protocols.

See the *Additional Resources* in this section for more information about enabling and working with this view.

#### Additional resources

- [Working with DNS tracking](#)
- [Network Observability metrics](#)

### 7.1.3. Round-Trip Time

You can use TCP smoothed Round-Trip Time (sRTT) to analyze network flow latencies. You can use RTT captured from the **fentry/tcp\_rcv\_established** eBPF hookpoint to read sRTT from the TCP socket to help with the following:

- **Network Monitoring:** Gain insights into TCP latencies, helping network administrators identify unusual patterns, potential bottlenecks, or performance issues.

- Troubleshooting: Debug TCP-related issues by tracking latency and identifying misconfigurations.

By default, when RTT is enabled, you can see the following TCP RTT metrics represented in the **Overview**:

- Top X 90th percentile TCP Round Trip Time with overall
- Top X average TCP Round Trip Time with overall
- Bottom X minimum TCP Round Trip Time with overall

Other RTT panels can be added in **Manage panels**:

- Top X maximum TCP Round Trip Time with overall
- Top X 99th percentile TCP Round Trip Time with overall

See the *Additional Resources* in this section for more information about enabling and working with this view.

#### Additional resources

- [Working with RTT tracing](#)

### 7.1.4. eBPF flow rule filter

You can use rule-based filtering to control the volume of packets cached in the eBPF flow table. For example, a filter can specify that only packets coming from port 100 should be recorded. Then only the packets that match the filter are cached and the rest are not cached.

#### 7.1.4.1. Ingress and egress traffic filtering

CIDR notation efficiently represents IP address ranges by combining the base IP address with a prefix length. For both ingress and egress traffic, the source IP address is first used to match filter rules configured with CIDR notation. If there is a match, then the filtering proceeds. If there is no match, then the destination IP is used to match filter rules configured with CIDR notation.

After matching either the source IP or the destination IP CIDR, you can pinpoint specific endpoints using the **peerIP** to differentiate the destination IP address of the packet. Based on the provisioned action, the flow data is either cached in the eBPF flow table or not cached.

#### 7.1.4.2. Dashboard and metrics integrations

When this option is enabled, the **Netobserv/Health** dashboard for **eBPF agent statistics** now has the **Filtered flows rate** view. Additionally, in **Observe** → **Metrics** you can query **netobserv\_agent\_filtered\_flows\_total** to observe metrics with the reason in **FlowFilterAcceptCounter**, **FlowFilterNoMatchCounter** or **FlowFilterRejectCounter**.

#### 7.1.4.3. Flow filter configuration parameters

The flow filter rules consist of required and optional parameters.

**Table 7.1. Required configuration parameters**

Parameter	Description
<b>enable</b>	Set <b>enable</b> to <b>true</b> to enable the eBPF flow filtering feature.
<b>cidr</b>	Provides the IP address and CIDR mask for the flow filter rule. Supports both IPv4 and IPv6 address format. If you want to match against any IP, you can use <b>0.0.0.0/0</b> for IPv4 or <b>:::0</b> for IPv6.
<b>action</b>	<p>Describes the action that is taken for the flow filter rule. The possible values are <b>Accept</b> or <b>Reject</b>.</p> <ul style="list-style-type: none"> <li>For the <b>Accept</b> action matching rule, the flow data is cached in the eBPF table and updated with the global metric, <b>FlowFilterAcceptCounter</b>.</li> <li>For the <b>Reject</b> action matching rule, the flow data is dropped and not cached in the eBPF table. The flow data is updated with the global metric, <b>FlowFilterRejectCounter</b>.</li> <li>If the rule is not matched, the flow is cached in the eBPF table and updated with the global metric, <b>FlowFilterNoMatchCounter</b>.</li> </ul>

Table 7.2. Optional configuration parameters

Parameter	Description
<b>direction</b>	Defines the direction of the flow filter rule. Possible values are <b>Ingress</b> or <b>Egress</b> .
<b>protocol</b>	Defines the protocol of the flow filter rule. Possible values are <b>TCP</b> , <b>UDP</b> , <b>SCTP</b> , <b>ICMP</b> , and <b>ICMPv6</b> .
<b>ports</b>	Defines the ports to use for filtering flows. It can be used for either source or destination ports. To filter a single port, set a single port as an integer value. For example <b>ports: 80</b> . To filter a range of ports, use a "start-end" range in string format. For example <b>ports: "80-100"</b>
<b>sourcePorts</b>	Defines the source port to use for filtering flows. To filter a single port, set a single port as an integer value, for example <b>sourcePorts: 80</b> . To filter a range of ports, use a "start-end" range, string format, for example <b>sourcePorts: "80-100"</b> .
<b>destPorts</b>	DestPorts defines the destination ports to use for filtering flows. To filter a single port, set a single port as an integer value, for example <b>destPorts: 80</b> . To filter a range of ports, use a "start-end" range in string format, for example <b>destPorts: "80-100"</b> .
<b>icmpType</b>	Defines the ICMP type to use for filtering flows.
<b>icmpCode</b>	Defines the ICMP code to use for filtering flows.

Parameter	Description
<b>peerIP</b>	Defines the IP address to use for filtering flows, for example: <b>10.10.10.10</b> .

### Additional resources

- [Filtering eBPF flow data with rules](#)
- [Network Observability metrics](#)
- [Health dashboards](#)

## 7.2. OBSERVING THE NETWORK TRAFFIC FROM THE TRAFFIC FLOWS VIEW

The **Traffic flows** view displays the data of the network flows and the amount of traffic in a table. As an administrator, you can monitor the amount of traffic across the application by using the traffic flow table.

### 7.2.1. Working with the Traffic flows view

As an administrator, you can navigate to **Traffic flows** table to see network flow information.

#### Procedure

1. Navigate to **Observe** → **Network Traffic**.
2. In the **Network Traffic** page, click the **Traffic flows** tab.

You can click on each row to get the corresponding flow information.

### 7.2.2. Configuring advanced options for the Traffic flows view

You can customize and export the view by using **Show advanced options**. You can set the row size by using the **Display options** drop-down menu. The default value is **Normal**.

#### 7.2.2.1. Managing columns

You can select the required columns to be displayed, and reorder them. To manage columns, click **Manage columns**.

#### 7.2.2.2. Exporting the traffic flow data

You can export data from the **Traffic flows** view.

#### Procedure

1. Click **Export data**.
2. In the pop-up window, you can select the **Export all data** checkbox to export all the data, and clear the checkbox to select the required fields to be exported.



3. Click **Export**.

### 7.2.3. Working with conversation tracking

As an administrator, you can group network flows that are part of the same conversation. A conversation is defined as a grouping of peers that are identified by their IP addresses, ports, and protocols, resulting in a unique **Conversation Id**. You can query conversation events in the web console. These events are represented in the web console as follows:

- **Conversation start**: This event happens when a connection is starting or TCP flag intercepted
- **Conversation tick**: This event happens at each specified interval defined in the **FlowCollector spec.processor.conversationHeartbeatInterval** parameter while the connection is active.
- **Conversation end**: This event happens when the **FlowCollector spec.processor.conversationEndTimeout** parameter is reached or the TCP flag is intercepted.
- **Flow**: This is the network traffic flow that occurs within the specified interval.

#### Procedure


1. In the web console, navigate to **Operators → Installed Operators**.
2. Under the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
3. Select **cluster** then select the **YAML** tab.
4. Configure the **FlowCollector** custom resource so that **spec.processor.logTypes**, **conversationEndTimeout**, and **conversationHeartbeatInterval** parameters are set according to your observation needs. A sample configuration is as follows:

#### Configure FlowCollector for conversation tracking

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  processor:
    logTypes: Flows
  advanced:
    conversationEndTimeout: 10s
    conversationHeartbeatInterval: 30s
```

- 1 When **logTypes** is set to **Flows**, only the **Flow** event is exported. If you set the value to **All**, both conversation and flow events are exported and visible in the **Network Traffic** page. To focus only on conversation events, you can specify **Conversations** which exports the **Conversation start**, **Conversation tick** and **Conversation end** events; or **EndedConversations** exports only the **Conversation end** events. Storage requirements are highest for **All** and lowest for **EndedConversations**.
- 2 The **Conversation end** event represents the point when the **conversationEndTimeout** is reached or the TCP flag is intercepted.

- 3** The **Conversation tick** event represents each specified interval defined in the **FlowCollector conversationHeartbeatInterval** parameter while the network connection is



#### NOTE

If you update the **logType** option, the flows from the previous selection do not clear from the console plugin. For example, if you initially set **logType** to **Conversations** for a span of time until 10 AM and then move to **EndedConversations**, the console plugin shows all conversation events before 10 AM and only ended conversations after 10 AM.

- Refresh the **Network Traffic** page on the **Traffic flows** tab. Notice there are two new columns, **Event/Type** and **Conversation Id**. All the **Event/Type** fields are **Flow** when **Flow** is the selected query option.
- Select **Query Options** and choose the **Log Type, Conversation**. Now the **Event/Type** shows all of the desired conversation events.
- Next you can filter on a specific conversation ID or switch between the **Conversation** and **Flow** log type options from the side panel.

### 7.2.4. Working with DNS tracking

Using DNS tracking, you can monitor your network, conduct security analysis, and troubleshoot DNS issues. You can track DNS by editing the **FlowCollector** to the specifications in the following YAML example.



#### IMPORTANT

CPU and memory usage increases are observed in the eBPF agent when this feature is enabled.

#### Procedure

- In the web console, navigate to **Operators → Installed Operators**.
- Under the **Provided APIs** heading for **Network Observability**, select **Flow Collector**.
- Select **cluster** then select the **YAML** tab.
- Configure the **FlowCollector** custom resource. A sample configuration is as follows:

#### Configure FlowCollector for DNS tracking

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  agent:
    type: eBPF
    ebpf:
```

features:

- DNSTracking 1

sampling: 1 2

- 1 You can set the **spec.agent.ebpf.features** parameter list to enable DNS tracking of each network flow in the web console.
  - 2 You can set **sampling** to a value of **1** for more accurate metrics and to capture **DNS latency**. For a **sampling** value greater than 1, you can observe flows with **DNS Response Code** and **DNS Id**, and it is unlikely that **DNS Latency** can be observed.
5. When you refresh the **Network Traffic** page, there are new DNS representations you can choose to view in the **Overview** and **Traffic Flow** views and new filters you can apply.
    - a. Select new DNS choices in **Manage panels** to display graphical visualizations and DNS metrics in the **Overview**.
    - b. Select new choices in **Manage columns** to add DNS columns to the **Traffic Flows** view.
    - c. Filter on specific DNS metrics, such as **DNS Id**, **DNS Error**, **DNS Latency** and **DNS Response Code**, and see more information from the side panel. The **DNS Latency** and **DNS Response Code** columns are shown by default.



#### NOTE

TCP handshake packets do not have DNS headers. TCP protocol flows without DNS headers are shown in the traffic flow data with **DNS Latency**, **ID**, and **Response code** values of "n/a". You can filter out flow data to view only flows that have DNS headers using the **Common** filter "DNSError" equal to "0".

### 7.2.5. Working with RTT tracing

You can track RTT by editing the **FlowCollector** to the specifications in the following YAML example.

#### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. In the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
3. Select **cluster**, and then select the **YAML** tab.
4. Configure the **FlowCollector** custom resource for RTT tracing, for example:

#### Example FlowCollector configuration

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  agent:
    type: eBPF
```

```
ebpf:
  features:
    - FlowRTT 1
```

- 1** You can start tracing RTT network flows by listing the **FlowRTT** parameter in the **spec.agent.ebpf.features** specification list.

## Verification

When you refresh the **Network Traffic** page, the **Overview**, **Traffic Flow**, and **Topology** views display new information about RTT:

- a. In the **Overview**, select new choices in **Manage panels** to choose which graphical visualizations of RTT to display.
- b. In the **Traffic flows** table, the **Flow RTT** column can be seen, and you can manage display in **Manage columns**.
- c. In the **Traffic Flows** view, you can also expand the side panel to view more information about RTT.

## Example filtering

- i. Click the **Common filters** → **Protocol**.
- ii. Filter the network flow data based on **TCP**, **Ingress** direction, and look for **FlowRTT** values greater than 10,000,000 nanoseconds (10ms).
- iii. Remove the **Protocol** filter.
- iv. Filter for **Flow RTT** values greater than 0 in the **Common** filters.
- d. In the **Topology** view, click the Display option dropdown. Then click **RTT** in the **edge labels** drop-down list.

### 7.2.5.1. Using the histogram

You can click **Show histogram** to display a toolbar view for visualizing the history of flows as a bar chart. The histogram shows the number of logs over time. You can select a part of the histogram to filter the network flow data in the table that follows the toolbar.

### 7.2.6. Working with availability zones

You can configure the **FlowCollector** to collect information about the cluster availability zones. This allows you to enrich network flow data with the [topology.kubernetes.io/zone](#) label value applied to the nodes.

## Procedure

1. In the web console, go to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
3. Select **cluster** then select the **YAML** tab.

4. Configure the **FlowCollector** custom resource so that the **spec.processor.addZone** parameter is set to **true**. A sample configuration is as follows:

### Configure FlowCollector for availability zones collection

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  # ...
  processor:
    addZone: true
  # ...
```

### Verification

When you refresh the **Network Traffic** page, the **Overview**, **Traffic Flow**, and **Topology** views display new information about availability zones:

1. In the **Overview** tab, you can see **Zones** as an available **Scope**.
2. In **Network Traffic** → **Traffic flows**, **Zones** are viewable under the **SrcK8S\_Zone** and **DstK8S\_Zone** fields.
3. In the **Topology** view, you can set **Zones** as **Scope** or **Group**.

### 7.2.7. Filtering eBPF flow data using a global rule

You can configure the **FlowCollector** to filter eBPF flows using a global rule to control the flow of packets cached in the eBPF flow table.

#### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for **Network Observability**, select **Flow Collector**.
3. Select **cluster**, then select the **YAML** tab.
4. Configure the **FlowCollector** custom resource, similar to the following sample configurations:

#### Example 7.1. Filter Kubernetes service traffic to a specific Pod IP endpoint

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  deploymentModel: Direct
  agent:
    type: eBPF
  ebpf:
    flowFilter:
      action: Accept 1
```

```

cidr: 172.210.150.1/24 2
protocol: SCTP
direction: Ingress
destPortRange: 80-100
peerIP: 10.10.10.10
enable: true 3

```

- 1 The required **action** parameter describes the action that is taken for the flow filter rule. Possible values are **Accept** or **Reject**.
- 2 The required **cidr** parameter provides the IP address and CIDR mask for the flow filter rule and supports IPv4 and IPv6 address formats. If you want to match against any IP address, you can use **0.0.0.0/0** for IPv4 or **::/0** for IPv6.
- 3 You must set **spec.agent.ebpf.flowFilter.enable** to **true** to enable this feature.

### Example 7.2. See flows to any addresses outside the cluster

```

apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  deploymentModel: Direct
agent:
  type: eBPF
  ebpf:
    flowFilter:
      action: Accept 1
      cidr: 0.0.0.0/0 2
      protocol: TCP
      direction: Egress
      sourcePort: 100
      peerIP: 192.168.127.12 3
      enable: true 4

```

- 1 You can **Accept** flows based on the criteria in the **flowFilter** specification.
- 2 The **cidr** value of **0.0.0.0/0** matches against any IP address.
- 3 See flows after **peerIP** is configured with **192.168.127.12**.
- 4 You must set **spec.agent.ebpf.flowFilter.enable** to **true** to enable the feature.

## 7.3. OBSERVING THE NETWORK TRAFFIC FROM THE TOPOLOGY VIEW

The **Topology** view provides a graphical representation of the network flows and the amount of traffic. As an administrator, you can monitor the traffic data across the application by using the **Topology** view.

### 7.3.1. Working with the Topology view

As an administrator, you can navigate to the **Topology** view to see the details and metrics of the component.

#### Procedure

1. Navigate to **Observe** → **Network Traffic**.
2. In the **Network Traffic** page, click the **Topology** tab.

You can click each component in the **Topology** to view the details and metrics of the component.

### 7.3.2. Configuring the advanced options for the Topology view

You can customize and export the view by using **Show advanced options**. The advanced options view has the following features:

- **Find in view:** To search the required components in the view.
- **Display options:** To configure the following options:
  - **Edge labels:** To show the specified measurements as edge labels. The default is to show the **Average rate** in **Bytes**.
  - **Scope:** To select the scope of components between which the network traffic flows. The default value is **Namespace**.
  - **Groups:** To enhance the understanding of ownership by grouping the components. The default value is **None**.
  - **Layout:** To select the layout of the graphical representation. The default value is **ColaNoForce**.
  - **Show:** To select the details that need to be displayed. All the options are checked by default. The options available are: **Edges**, **Edges label**, and **Badges**.
  - **Truncate labels:** To select the required width of the label from the drop-down list. The default value is **M**.
  - **Collapse groups:** To expand or collapse the groups. The groups are expanded by default. This option is disabled if **Groups** has the value of **None**.

#### 7.3.2.1. Exporting the topology view

To export the view, click **Export topology view**. The view is downloaded in PNG format.

## 7.4. FILTERING THE NETWORK TRAFFIC

By default, the Network Traffic page displays the traffic flow data in the cluster based on the default filters configured in the **FlowCollector** instance. You can use the filter options to observe the required data by changing the preset filter.

#### Query Options

You can use **Query Options** to optimize the search results, as listed below:

- **Log Type:** The available options **Conversation** and **Flows** provide the ability to query flows by log type, such as flow log, new conversation, completed conversation, and a heartbeat, which is a periodic record with updates for long conversations. A conversation is an aggregation of flows between the same peers.
- **Match filters:** You can determine the relation between different filter parameters selected in the advanced filter. The available options are **Match all** and **Match any**. **Match all** provides results that match all the values, and **Match any** provides results that match any of the values entered. The default value is **Match all**.
- **Datasource:** You can choose the datasource to use for queries: **Loki**, **Prometheus**, or **Auto**. Notable performance improvements can be realized when using Prometheus as a datasource rather than Loki, but Prometheus supports a limited set of filters and aggregations. The default datasource is **Auto**, which uses Prometheus on supported queries or uses Loki if the query does not support Prometheus.
- **Drops filter:** You can view different levels of dropped packets with the following query options:
  - **Fully dropped** shows flow records with fully dropped packets.
  - **Containing drops** shows flow records that contain drops but can be sent.
  - **Without drops** shows records that contain sent packets.
  - **All** shows all the aforementioned records.
- **Limit:** The data limit for internal backend queries. Depending upon the matching and the filter settings, the number of traffic flow data is displayed within the specified limit.

### Quick filters

The default values in **Quick filters** drop-down menu are defined in the **FlowCollector** configuration. You can modify the options from console.

### Advanced filters

You can set the advanced filters, **Common**, **Source**, or **Destination**, by selecting the parameter to be filtered from the dropdown list. The flow data is filtered based on the selection. To enable or disable the applied filter, you can click on the applied filter listed below the filter options.

You can toggle between **↑ One way** and **↑ ↓ Back and forth** filtering. The **↑ One way** filter shows only **Source** and **Destination** traffic according to your filter selections. You can use **Swap** to change the directional view of the **Source** and **Destination** traffic. The **↑ ↓ Back and forth** filter includes return traffic with the **Source** and **Destination** filters. The directional flow of network traffic is shown in the **Direction** column in the Traffic flows table as **Ingress** or **Egress** for inter-node traffic and **Inner** for traffic inside a single node.

You can click **Reset defaults** to remove the existing filters, and apply the filter defined in **FlowCollector** configuration.



#### NOTE

To understand the rules of specifying the text value, click **Learn More**.



Alternatively, you can access the traffic flow data in the **Network Traffic** tab of the **Namespaces**, **Services**, **Routes**, **Nodes**, and **Workloads** pages which provide the filtered data of the corresponding aggregations.

### **Additional resources**

For more information about configuring quick filters in the **FlowCollector**, see [Configuring Quick Filters](#) and the [Flow Collector sample resource](#).

## CHAPTER 8. USING METRICS WITH DASHBOARDS AND ALERTS

The Network Observability Operator uses the **flowlogs-pipeline** to generate metrics from flow logs. You can utilize these metrics by setting custom alerts and viewing dashboards.

### 8.1. VIEWING NETWORK OBSERVABILITY METRICS DASHBOARDS

On the **Overview** tab in the OpenShift Container Platform console, you can view the overall aggregated metrics of the network traffic flow on the cluster. You can choose to display the information by node, namespace, owner, pod, and service. You can also use filters and display options to further refine the metrics.

#### Procedure

1. In the web console **Observe** → **Dashboards**, select the **Netobserv** dashboard.
2. View network traffic metrics in the following categories, with each having the subset per node, namespace, source, and destination:
  - **Byte rates**
  - **Packet drops**
  - **DNS**
  - **RTT**
3. Select the **Netobserv/Health** dashboard.
4. View metrics about the health of the Operator in the following categories, with each having the subset per node, namespace, source, and destination.
  - **Flows**
  - **Flows Overhead**
  - **Flow rates**
  - **Agents**
  - **Processor**
  - **Operator**

**Infrastructure** and **Application** metrics are shown in a split-view for namespace and workloads.

### 8.2. PREDEFINED METRICS

Metrics generated by the **flowlogs-pipeline** are configurable in the **spec.processor.metrics.includeList** of the **FlowCollector** custom resource to add or remove metrics.

### 8.3. NETWORK OBSERVABILITY METRICS

You can also create alerts by using the **includeList** metrics in Prometheus rules, as shown in the example "Creating alerts".

When looking for these metrics in Prometheus, such as in the Console through **Observe** → **Metrics**, or when defining alerts, all the metrics names are prefixed with **netobserv\_**. For example, **netobserv\_namespace\_flows\_total**. Available metrics names are as follows:

#### includeList metrics names

Names followed by an asterisk \* are enabled by default.

- **namespace\_egress\_bytes\_total**
- **namespace\_egress\_packets\_total**
- **namespace\_ingress\_bytes\_total**
- **namespace\_ingress\_packets\_total**
- **namespace\_flows\_total** \*
- **node\_egress\_bytes\_total**
- **node\_egress\_packets\_total**
- **node\_ingress\_bytes\_total** \*
- **node\_ingress\_packets\_total**
- **node\_flows\_total**
- **workload\_egress\_bytes\_total**
- **workload\_egress\_packets\_total**
- **workload\_ingress\_bytes\_total** \*
- **workload\_ingress\_packets\_total**
- **workload\_flows\_total**

#### PacketDrop metrics names

When the **PacketDrop** feature is enabled in **spec.agent.ebpf.features** (with **privileged** mode), the following additional metrics are available:

- **namespace\_drop\_bytes\_total**
- **namespace\_drop\_packets\_total** \*
- **node\_drop\_bytes\_total**
- **node\_drop\_packets\_total**
- **workload\_drop\_bytes\_total**
- **workload\_drop\_packets\_total**

#### DNS metrics names

When the **DNSTracking** feature is enabled in **spec.agent.ebpf.features**, the following additional metrics are available:

- **namespace\_dns\_latency\_seconds** \*
- **node\_dns\_latency\_seconds**
- **workload\_dns\_latency\_seconds**

#### FlowRTT metrics names

When the **FlowRTT** feature is enabled in **spec.agent.ebpf.features**, the following additional metrics are available:

- **namespace\_rtt\_seconds** \*
- **node\_rtt\_seconds**
- **workload\_rtt\_seconds**

## 8.4. CUSTOM METRICS

You can create custom metrics out of the flowlogs data using the **FlowMetric** API. In every flowlogs data that is collected, there are a number of fields labeled per log, such as source name and destination name. These fields can be leveraged as Prometheus labels to enable the customization of cluster information on your dashboard.

## 8.5. CONFIGURING CUSTOM METRICS BY USING FLOWMETRIC API

You can configure the **FlowMetric** API to create custom metrics by using flowlogs data fields as Prometheus labels. You can add multiple **FlowMetric** resources to a project to see multiple dashboard views.

#### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. In the **Provided APIs** heading for the **NetObserv Operator**, select **FlowMetric**.
3. In the **Project:** dropdown list, select the project of the Network Observability Operator instance.
4. Click **Create FlowMetric**.
5. Configure the **FlowMetric** resource, similar to the following sample configurations:

**Example 8.1. Generate a metric that tracks ingress bytes received from cluster external sources**

```
apiVersion: flows.netobserv.io/v1alpha1
kind: FlowMetric
metadata:
  name: flowmetric-cluster-external-ingress-traffic
  namespace: netobserv 1
spec:
  metricName: cluster_external_ingress_bytes_total 2
```

```

type: Counter 3
valueField: Bytes
direction: Ingress 4
labels:
[DstK8S_HostName,DstK8S_Namespace,DstK8S_OwnerName,DstK8S_OwnerType] 5
filters: 6
- field: SrcSubnetLabel
  matchType: Absence

```

- 1 The **FlowMetric** resources need to be created in the namespace defined in the **FlowCollector spec.namespace**, which is **netobserv** by default.
- 2 The name of the Prometheus metric, which in the web console appears with the prefix **netobserv-<metricName>**.
- 3 The **type** specifies the type of metric. The **Counter type** is useful for counting bytes or packets.
- 4 The direction of traffic to capture. If not specified, both ingress and egress are captured, which can lead to duplicated counts.
- 5 Labels define what the metrics look like and the relationship between the different entities and also define the metrics cardinality. For example, **SrcK8S\_Name** is a high cardinality metric.
- 6 Refines results based on the listed criteria. In this example, selecting only the cluster external traffic is done by matching only flows where **SrcSubnetLabel** is absent. This assumes the subnet labels feature is enabled (via **spec.processor.subnetLabels**), which is done by default.

### Verification

1. Once the pods refresh, navigate to **Observe → Metrics**.
2. In the **Expression** field, type the metric name to view the corresponding result. You can also enter an expression, such as **topk(5, sum(rate(netobserv\_cluster\_external\_ingress\_bytes\_total{DstK8S\_Namespace="my-namespace"}[2m])) by (DstK8S\_HostName, DstK8S\_OwnerName, DstK8S\_OwnerType))**

### Example 8.2. Show RTT latency for cluster external ingress traffic

```

apiVersion: flows.netobserv.io/v1 alpha1
kind: FlowMetric
metadata:
  name: flowmetric-cluster-external-ingress-rtt
  namespace: netobserv 1
spec:
  metricName: cluster_external_ingress_rtt_seconds
  type: Histogram 2
  valueField: TimeFlowRttNs
  direction: Ingress
  labels:


```

```
[DstK8S_HostName,DstK8S_Namespace,DstK8S_OwnerName,DstK8S_OwnerType]
filters:
- field: SrcSubnetLabel
  matchType: Absence
- field: TimeFlowRttNs
  matchType: Presence
divider: "1000000000" 3
buckets: [".001", ".005", ".01", ".02", ".03", ".04", ".05", ".075", ".1", ".25", "1"] 4
```

- 1** The **FlowMetric** resources need to be created in the namespace defined in the **FlowCollector spec.namespace**, which is **netobserv** by default.
- 2** The **type** specifies the type of metric. The **Histogram type** is useful for a latency value (**TimeFlowRttNs**).
- 3** Since the Round-trip time (RTT) is provided as nanos in flows, use a divider of 1 billion to convert into seconds, which is standard in Prometheus guidelines.
- 4** The custom buckets specify precision on RTT, with optimal precision ranging between 5ms and 250ms.

### Verification

1. Once the pods refresh, navigate to **Observe → Metrics**.
2. In the **Expression** field, you can type the metric name to view the corresponding result.



### IMPORTANT

High cardinality can affect the memory usage of Prometheus. You can check whether specific labels have high cardinality in the [Network Flows format reference](#).

## 8.6. CONFIGURING CUSTOM CHARTS USING FLOWMETRIC API

You can generate charts for dashboards in the OpenShift Container Platform web console, which you can view as an administrator in the **Dashboard** menu by defining the **charts** section of the **FlowMetric** resource.

### Procedure

1. In the web console, navigate to **Operators → Installed Operators**.
2. In the **Provided APIs** heading for the **NetObserv Operator**, select **FlowMetric**.
3. In the **Project:** dropdown list, select the project of the Network Observability Operator instance.
4. Click **Create FlowMetric**.
5. Configure the **FlowMetric** resource, similar to the following sample configurations:

#### Example 8.3. Chart for tracking ingress bytes received from cluster external sources

```
apiVersion: flows.netobserv.io/v1alpha1
```

```

kind: FlowMetric
metadata:
  name: flowmetric-cluster-external-ingress-traffic
  namespace: netobserv 1
# ...
charts:
- dashboardName: Main 2
  title: External ingress traffic
  unit: Bps
  type: SingleStat
  queries:
  - promQL: "sum(rate($METRIC[2m]))"
    legend: ""
- dashboardName: Main 3
  sectionName: External
  title: Top external ingress traffic per workload
  unit: Bps
  type: StackArea
  queries:
  - promQL: "sum(rate($METRIC{DstK8S_Namespace!=\"\"}[2m])) by (DstK8S_Namespace,
    DstK8S_OwnerName)"
    legend: "{{DstK8S_Namespace}} / {{DstK8S_OwnerName}}"
# ...

```

- 1** The **FlowMetric** resources need to be created in the namespace defined in the **FlowCollector spec.namespace**, which is **netobserv** by default.

## Verification

1. Once the pods refresh, navigate to **Observe → Dashboards**.
2. Search for the **NetObserv / Main** dashboard. View two panels under the **NetObserv / Main** dashboard, or optionally a dashboard name that you create:
  - A textual single statistic showing the global external ingress rate summed across all dimensions
  - A timeseries graph showing the same metric per destination workload

For more information about the query language, refer to the [Prometheus documentation](#).

## Example 8.4. Chart for RTT latency for cluster external ingress traffic

```

apiVersion: flows.netobserv.io/v1alpha1
kind: FlowMetric
metadata:
  name: flowmetric-cluster-external-ingress-traffic
  namespace: netobserv 1
# ...
charts:
- dashboardName: Main 2
  title: External ingress TCP latency
  unit: seconds
  type: SingleStat

```

```

queries:
- promQL: "histogram_quantile(0.99, sum(rate($METRIC_bucket[2m])) by (le)) > 0"
  legend: "p99"
- dashboardName: Main 3
  sectionName: External
  title: "Top external ingress sRTT per workload, p50 (ms)"
  unit: seconds
  type: Line
  queries:
- promQL: "histogram_quantile(0.5, sum(rate($METRIC_bucket{DstK8S_Namespace!=\"\"}
[2m])) by (le,DstK8S_Namespace,DstK8S_OwnerName))*1000 > 0"
  legend: "{{DstK8S_Namespace}} / {{DstK8S_OwnerName}}"
- dashboardName: Main 4
  sectionName: External
  title: "Top external ingress sRTT per workload, p99 (ms)"
  unit: seconds
  type: Line
  queries:
- promQL: "histogram_quantile(0.99, sum(rate($METRIC_bucket{DstK8S_Namespace!=\"\"}
[2m])) by (le,DstK8S_Namespace,DstK8S_OwnerName))*1000 > 0"
  legend: "{{DstK8S_Namespace}} / {{DstK8S_OwnerName}}"
# ...

```

**1** The **FlowMetric** resources need to be created in the namespace defined in the **FlowCollector spec.namespace**, which is **netobserv** by default.

**2 3 4** Using a different **dashboardName** creates a new dashboard that is prefixed with **Netobserv**. For example, **Netobserv / <dashboard\_name>**.

This example uses the **histogram\_quantile** function to show **p50** and **p99**.

You can show averages of histograms by dividing the metric, **\$METRIC\_sum**, by the metric **\$METRIC\_count**, which are automatically generated when you create a histogram. With the preceding example, the Prometheus query to do this is as follows:

```

promQL: "(sum(rate($METRIC_sum{DstK8S_Namespace!=\"\"}[2m])) by
(DstK8S_Namespace,DstK8S_OwnerName) /
sum(rate($METRIC_count{DstK8S_Namespace!=\"\"}[2m])) by
(DstK8S_Namespace,DstK8S_OwnerName))*1000"

```

## Verification

1. Once the pods refresh, navigate to **Observe → Dashboards**.
2. Search for the **NetObserv / Main** dashboard. View the new panel under the **NetObserv / Main** dashboard, or optionally a dashboard name that you create.

For more information about the query language, refer to the [Prometheus documentation](#).

## Additional resources

- [Creating alerting rules for user-defined projects](#).



- Troubleshooting high cardinality metrics- Determining why Prometheus is consuming a lot of disk space

## CHAPTER 9. MONITORING THE NETWORK OBSERVABILITY OPERATOR

You can use the web console to monitor alerts related to the health of the Network Observability Operator.

### 9.1. HEALTH DASHBOARDS

Metrics about health and resource usage of the Network Observability Operator are located in the **Observe** → **Dashboards** page in the web console. You can view metrics about the health of the Operator in the following categories:

- **Flows per second**
- **Sampling**
- **Errors last minute**
- **Dropped flows per second**
- **Flowlogs-pipeline statistics**
- **Flowlogs-pipeline statistics views**
- **eBPF agent statistics views**
- **Operator statistics**
- **Resource usage**

### 9.2. HEALTH ALERTS

A health alert banner that directs you to the dashboard can appear on the **Network Traffic** and **Home** pages if an alert is triggered. Alerts are generated in the following cases:

- The **NetObservLokiError** alert occurs if the **flowlogs-pipeline** workload is dropping flows because of Loki errors, such as if the Loki ingestion rate limit has been reached.
- The **NetObservNoFlows** alert occurs if no flows are ingested for a certain amount of time.
- The **NetObservFlowsDropped** alert occurs if the Network Observability eBPF agent hashmap table is full, and the eBPF agent processes flows with degraded performance, or when the capacity limiter is triggered.

### 9.3. VIEWING HEALTH INFORMATION

You can access metrics about health and resource usage of the Network Observability Operator from the **Dashboards** page in the web console.

#### Prerequisites

- You have the Network Observability Operator installed.

- You have access to the cluster as a user with the **cluster-admin** role or with view permissions for all projects.

### Procedure

1. From the **Administrator** perspective in the web console, navigate to **Observe** → **Dashboards**.
2. From the **Dashboards** dropdown, select **Netobserv/Health**.
3. View the metrics about the health of the Operator that are displayed on the page.

### 9.3.1. Disabling health alerts

You can opt out of health alerting by editing the **FlowCollector** resource:


1. In the web console, navigate to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for the **NetObserv Operator**, select **Flow Collector**.
3. Select **cluster** then select the **YAML** tab.
4. Add **spec.processor.metrics.disableAlerts** to disable health alerts, as in the following YAML sample:

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  processor:
    metrics:
      disableAlerts: [NetObservLokiError, NetObservNoFlows] 1
```

- 1 You can specify one or a list with both types of alerts to disable.

## 9.4. USING THE EBPF AGENT ALERT

An alert, **NetObservAgentFlowsDropped**, is triggered when the Network Observability eBPF agent hashmap table is full or when the capacity limiter is triggered. If you see this alert, consider increasing the **cacheMaxFlows** in the **FlowCollector**, as shown in the following example.



### NOTE

Increasing the **cacheMaxFlows** might increase the memory usage of the eBPF agent.

### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**.
2. Under the **Provided APIs** heading for the **Network Observability Operator**, select **Flow Collector**.
3. Select **cluster**, and then select the **YAML** tab.

4. Increase the **spec.agent.ebpf.cacheMaxFlows** value, as shown in the following YAML sample:

```
apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
  namespace: netobserv
  deploymentModel: Direct
  agent:
    type: eBPF
    ebpf:
      cacheMaxFlows: 200000 1
```

- 1 Increase the **cacheMaxFlows** value from its value at the time of the **NetObservAgentFlowsDropped** alert.

## CHAPTER 10. SCHEDULING RESOURCES

Taints and tolerations allow the node to control which pods should (or should not) be scheduled on them.

A node selector specifies a map of key/value pairs that are defined using custom labels on nodes and selectors specified in pods.

For the pod to be eligible to run on a node, the pod must have the same key/value node selector as the label on the node.

### 10.1. NETWORK OBSERVABILITY DEPLOYMENT IN SPECIFIC NODES

You can configure the **FlowCollector** to control the deployment of Network Observability components in specific nodes. The **spec.agent.ebpf.advanced.scheduling**, **spec.processor.advanced.scheduling**, and **spec.consolePlugin.advanced.scheduling** specifications have the following configurable settings:

- **NodeSelector**
- **Tolerations**
- **Affinity**
- **PriorityClassName**

Sample **FlowCollector** resource for **spec.<component>.advanced.scheduling**

```

apiVersion: flows.netobserv.io/v1beta2
kind: FlowCollector
metadata:
  name: cluster
spec:
# ...
advanced:
  scheduling:
    tolerations:
      - key: "<taint key>"
        operator: "Equal"
        value: "<taint value>"
        effect: "<taint effect>"
      nodeSelector:
        <key>: <value>
    affinity:
      nodeAffinity:
        requiredDuringSchedulingIgnoredDuringExecution:
          nodeSelectorTerms:
            - matchExpressions:
                - key: name
                  operator: In
                  values:
                    - app-worker-node
          priorityClassName: ""
# ...

```

### Additional resources

- [Understanding taints and tolerations](#)
- [Assign Pods to Nodes](#) (Kubernetes documentation)
- [Pod Priority and Preemption](#) (Kubernetes documentation)

## CHAPTER 11. NETWORK OBSERVABILITY CLI

### 11.1. INSTALLING THE NETWORK OBSERVABILITY CLI

The Network Observability CLI (**oc netobserv**) is deployed separately from the Network Observability Operator. The CLI is available as an OpenShift CLI (**oc**) plugin. It provides a lightweight way to quickly debug and troubleshoot with network observability.



#### IMPORTANT

Network Observability CLI (**oc netobserv**) is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

#### 11.1.1. About the Network Observability CLI

You can quickly debug and troubleshoot networking issues by using the Network Observability CLI (**oc netobserv**). The Network Observability CLI is a flow and packet visualization tool that relies on eBPF agents to stream collected data to an ephemeral collector pod. It requires no persistent storage during the capture. After the run, the output is transferred to your local machine. This enables quick, live insight into packets and flow data without installing the Network Observability Operator.



#### IMPORTANT

CLI capture is meant to run only for short durations, such as 8-10 minutes. If it runs for too long, it can be difficult to delete the running process.

#### 11.1.2. Installing the Network Observability CLI

Installing the Network Observability CLI (**oc netobserv**) is a separate procedure from the Network Observability Operator installation. This means that, even if you have the Operator installed from OperatorHub, you need to install the CLI separately.



#### NOTE

You can optionally use Krew to install the **netobserv** CLI plugin. For more information, see "Installing a CLI plugin with Krew".

#### Prerequisites

- You must install the OpenShift CLI (**oc**).
- You must have a macOS or Linux operating system.

#### Procedure

1. Download the [oc netobserv CLI tar file](#).

2. Unpack the archive:

```
$ tar xvf netobserv-cli.tar.gz
```

3. Make the file executable:

```
$ chmod +x ./build/oc-netobserv
```

4. Move the extracted **netobserv-cli** binary to a directory that is on your **PATH**, such as **/usr/local/bin/**:

```
$ sudo mv ./build/oc-netobserv /usr/local/bin/
```

### Verification

- Verify that **oc netobserv** is available:

```
$ oc netobserv version
```

### Example output

```
Netobserv CLI version <version>
```

### Additional resources

- [Installing and using CLI plugins](#)
- [Installing a CLI plugin with Krew](#)

## 11.2. USING THE NETWORK OBSERVABILITY CLI

You can visualize and filter the flows and packets data directly in the terminal to see specific usage, such as identifying who is using a specific port. The Network Observability CLI collects flows as JSON and database files or packets as a PCAP file, which you can use with third-party tools.

### 11.2.1. Capturing flows

You can capture flows and filter on any resource or zone in the data to solve use cases, such as displaying Round-Trip Time (RTT) between two zones. Table visualization in the CLI provides viewing and flow search capabilities.

#### Prerequisites

- Install the OpenShift CLI (**oc**).
- Install the Network Observability CLI (**oc netobserv**) plugin.

#### Procedure

1. Capture flows with filters enabled by running the following command:



```
$ oc netobserv flows --enable_filter=true --action=Accept --cidr=0.0.0.0/0 --protocol=TCP --port=49051
```

2. Add filters to the **live table filter** prompt in the terminal to further refine the incoming flows. For example:

```
live table filter: [SrcK8S_Zone:us-west-1b] press enter to match multiple regular expressions at once
```

3. To stop capturing, press **Ctrl+C**. The data that was captured is written to two separate files in an **.output** directory located in the same path used to install the CLI.
4. View the captured data in the **./output/flow/<capture\_date\_time>.json** JSON file, which contains JSON arrays of the captured data.

### Example JSON file

```
{
  "AgentIP": "10.0.1.76",
  "Bytes": 561,
  "DnsErrno": 0,
  "Dscp": 20,
  "DstAddr": "f904:ece9:ba63:6ac7:8018:1e5:7130:0",
  "DstMac": "0A:58:0A:80:00:37",
  "DstPort": 9999,
  "Duplicate": false,
  "Etype": 2048,
  "Flags": 16,
  "FlowDirection": 0,
  "IfDirection": 0,
  "Interface": "ens5",
  "K8S_FlowLayer": "infra",
  "Packets": 1,
  "Proto": 6,
  "SrcAddr": "3e06:6c10:6440:2:a80:37:b756:270f",
  "SrcMac": "0A:58:0A:80:00:01",
  "SrcPort": 46934,
  "TimeFlowEndMs": 1709741962111,
  "TimeFlowRttNs": 121000,
  "TimeFlowStartMs": 1709741962111,
  "TimeReceived": 1709741964
}
```

5. You can use SQLite to inspect the **./output/flow/<capture\_date\_time>.db** database file. For example:
  - a. Open the file by running the following command:

```
$ sqlite3 ./output/flow/<capture_date_time>.db
```

- b. Query the data by running a SQLite **SELECT** statement, for example:

```
sqlite> SELECT DnsLatencyMs, DnsFlagsResponseCode, DnsId, DstAddr, DstPort,
Interface, Proto, SrcAddr, SrcPort, Bytes, Packets FROM flow WHERE DnsLatencyMs
>10 LIMIT 10;
```

### Example output

```
12|NoError|58747|10.128.0.63|57856||17|172.30.0.10|53|284|1
11|NoError|20486|10.128.0.52|56575||17|169.254.169.254|53|225|1
11|NoError|59544|10.128.0.103|51089||17|172.30.0.10|53|307|1
13|NoError|32519|10.128.0.52|55241||17|169.254.169.254|53|254|1
12|NoError|32519|10.0.0.3|55241||17|169.254.169.254|53|254|1
15|NoError|57673|10.128.0.19|59051||17|172.30.0.10|53|313|1
13|NoError|35652|10.0.0.3|46532||17|169.254.169.254|53|183|1
32|NoError|37326|10.0.0.3|52718||17|169.254.169.254|53|169|1
14|NoError|14530|10.0.0.3|58203||17|169.254.169.254|53|246|1
15|NoError|40548|10.0.0.3|45933||17|169.254.169.254|53|174|1
```

## 11.2.2. Capturing packets

You can capture packets using the Network Observability CLI.

### Prerequisites

- Install the OpenShift CLI (**oc**).
- Install the Network Observability CLI (**oc netobserv**) plugin.

### Procedure

1. Run the packet capture with filters enabled:

```
$ oc netobserv packets tcp,80
```

2. Add filters to the **live table filter** prompt in the terminal to refine the incoming packets. An example filter is as follows:

```
live table filter: [SrcK8S_Zone:us-west-1b] press enter to match multiple regular expressions
at once
```

3. To stop capturing, press **Ctrl+C**.
4. View the captured data, which is written to a single file in an **./output/pcap** directory located in the same path that was used to install the CLI:
  - a. The **./output/pcap/<capture\_date\_time>.pcap** file can be opened with Wireshark.

## 11.2.3. Cleaning the Network Observability CLI

You can manually clean the CLI workload by running **oc netobserv cleanup**. This command removes all the CLI components from your cluster.

When you end a capture, this command is run automatically by the client. You might be required to manually run it if you experience connectivity issues.

## Procedure

- Run the following command:

```
$ oc netobserv cleanup
```

## Additional resources

- [Network Observability CLI reference](#)

## 11.3. NETWORK OBSERVABILITY CLI (OC NETOBSERV) REFERENCE

The Network Observability CLI (**oc netobserv**) has most features and filtering options that are available for the Network Observability Operator. You can pass command line arguments to enable features or filtering options.

### 11.3.1. oc netobserv CLI reference

The Network Observability CLI (**oc netobserv**) is a CLI tool for capturing flow data and packet data for further analysis.

#### oc netobserv syntax

```
$ oc netobserv [<command>] [<feature_option>] [<command_options>] 1
```

- 1** Feature options can only be used with the **oc netobserv flows** command. They cannot be used with the **oc netobserv packets** command.

Table 11.1. Basic commands

Command	Description
<b>flows</b>	Capture flows information. For subcommands, see the "Flow capture subcommands" table.
<b>packets</b>	Capture packets from a specific protocol or port pair, such as <b>netobserv packets --filter=tcp,80</b> . For more information about packet capture, see the "Packet capture subcommand" table.
<b>cleanup</b>	Remove the Network Observability CLI components.
<b>version</b>	Print the software version.
<b>help</b>	Show help.

#### 11.3.1.1. Network Observability enrichment

The Network Observability enrichment to display zone, node, owner and resource names including optional features about DNS latencies and Round-trip time can only be enabled when capturing flows. These do not appear in packet capture pcap output file.

## Network Observability enrichment syntax

```
$ oc netobserv flows [<enrichment_options>] [<subcommands>]
```

Table 11.2. Network Observability enrichment options

Option	Description	Possible values	Default
<b>--enable_rtt</b>	Enable round trip time.	<b>true, false</b>	<b>false</b>
<b>--enable_dns</b>	Enable DNS tracking.	<b>true, false</b>	<b>false</b>
<b>--help</b>	Show help.	-	-
<b>--interfaces</b>	Interfaces to match on the flow. For example, <b>"eth0,eth1"</b> .	<b>"&lt;interface&gt;"</b>	-

### 11.3.1.2. Flow capture options

Flow capture has mandatory commands as well as additional options, such as enabling extra features about DNS latencies, Round-trip time, and filtering.

#### oc netobserv flows syntax

```
$ oc netobserv flows [<feature_option>] [<command_options>]
```

Table 11.3. Flow capture filter options

Option	Description	Possible values	Mandatory	Default
<b>--enable_filter</b>	Enable flow filter.	<b>true, false</b>	Yes	<b>false</b>
<b>--action</b>	Action to apply on the flow.	<b>Accept, Reject</b>	Yes	<b>Accept</b>
<b>--cidr</b>	CIDR to match on the flow.	<b>1.1.1.0/24, 1::100/64, or 0.0.0.0/0</b>	Yes	<b>0.0.0.0/0</b>
<b>--protocol</b>	Protocol to match on the flow	<b>TCP, UDP, SCTP, ICMP, or ICMPv6</b>	No	-
<b>--direction</b>	Direction to match on the flow	<b>Ingress, Egress</b>	No	-
<b>--dport</b>	Destination port to match on the flow.	<b>80, 443, or 49051</b>	no	-

Option	Description	Possible values	Mandatory	Default
<b>--sport</b>	Source port to match on the flow.	<b>80, 443, or 49051</b>	No	-
<b>--port</b>	Port to match on the flow.	<b>80, 443, or 49051</b>	No	-
<b>--sport_range</b>	Source port range to match on the flow.	<b>80-100 or 443-445</b>	No	-
<b>--dport_range</b>	Destination port range to match on the flow.	<b>80-100</b>	No	-
<b>--port_range</b>	Port range to match on the flow.	<b>80-100 or 443-445</b>	No	-
<b>--icmp_type</b>	ICMP type to match on the flow.	<b>8 or 13</b>	No	-
<b>--icmp_code</b>	ICMP code to match on the flow.	<b>0 or 1</b>	No	-
<b>--peer_ip</b>	Peer IP to match on the flow.	<b>1.1.1.1 or 1::1</b>	No	-

### 11.3.1.3. Packet capture options

You can filter on port and protocol for packet capture data.

#### oc netobserv packets syntax

```
$ oc netobserv packets [<option>]
```

Table 11.4. Packet capture filter option

Option	Description	Mandatory	Default
<b>&lt;protocol&gt;,&lt;port&gt;</b>	Capture packets from a specific protocol and port pair. Use a comma as delimiter. For example, <b>tcp,80</b> specifies the <b>tcp</b> protocol and port <b>80</b> .	Yes	-

## CHAPTER 12. FLOWCOLLECTOR CONFIGURATION PARAMETERS

FlowCollector is the Schema for the network flows collection API, which pilots and configures the underlying deployments.

### 12.1. FLOWCOLLECTOR API SPECIFICATIONS

#### Description

**FlowCollector** is the schema for the network flows collection API, which pilots and configures the underlying deployments.

#### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and might reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers might infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>object</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>spec</b>	<b>object</b>	<p>Defines the desired state of the FlowCollector resource.</p> <p>*: the mention of "unsupported" or "deprecated" for a feature throughout this document means that this feature is not officially supported by Red Hat. It might have been, for example, contributed by the community and accepted without a formal agreement for maintenance. The product maintainers might provide some support for these features as a best effort only.</p>

### 12.1.1. .metadata

#### Description

Standard object's metadata. More info: <https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata>

#### Type

**object**

### 12.1.2. .spec

#### Description

Defines the desired state of the FlowCollector resource.

\*: the mention of "unsupported" or "deprecated" for a feature throughout this document means that this feature is not officially supported by Red Hat. It might have been, for example, contributed by the community and accepted without a formal agreement for maintenance. The product maintainers might provide some support for these features as a best effort only.

#### Type

**object**

Property	Type	Description
<b>agent</b>	<b>object</b>	Agent configuration for flows extraction.

Property	Type	Description
<b>consolePlugin</b>	<b>object</b>	<b>consolePlugin</b> defines the settings related to the OpenShift Container Platform Console plugin, when available.
<b>deploymentModel</b>	<b>string</b>	<p><b>deploymentModel</b> defines the desired type of deployment for flow processing. Possible values are:</p> <ul style="list-style-type: none"> <li>- <b>Direct</b> (default) to make the flow processor listen directly from the agents.</li> <li>- <b>Kafka</b> to make flows sent to a Kafka pipeline before consumption by the processor.</li> </ul> <p>Kafka can provide better scalability, resiliency, and high availability (for more details, see <a href="https://www.redhat.com/en/topic/s/integration/what-is-apache-kafka">https://www.redhat.com/en/topic/s/integration/what-is-apache-kafka</a>).</p>
<b>exporters</b>	<b>array</b>	<b>exporters</b> define additional optional exporters for custom consumption or storage.
<b>kafka</b>	<b>object</b>	Kafka configuration, allowing to use Kafka as a broker as part of the flow collection pipeline. Available when the <b>spec.deploymentModel</b> is <b>Kafka</b> .
<b>loki</b>	<b>object</b>	<b>loki</b> , the flow store, client settings.
<b>namespace</b>	<b>string</b>	Namespace where Network Observability pods are deployed.
<b>processor</b>	<b>object</b>	<b>processor</b> defines the settings of the component that receives the flows from the agent, enriches them, generates metrics, and forwards them to the Loki persistence layer and/or any available exporter.



Property	Type	Description
<b>prometheus</b>	<b>object</b>	<b>prometheus</b> defines Prometheus settings, such as querier configuration used to fetch metrics from the Console plugin.

### 12.1.3. .spec.agent

#### Description

Agent configuration for flows extraction.

#### Type

**object**

Property	Type	Description
<b>ebpf</b>	<b>object</b>	<b>ebpf</b> describes the settings related to the eBPF-based flow reporter when <b>spec.agent.type</b> is set to <b>eBPF</b> .
<b>type</b>	<b>string</b>	<b>type</b> [deprecated (*)] selects the flows tracing agent. Previously, this field allowed to select between <b>eBPF</b> or <b>IPFIX</b> . Only <b>eBPF</b> is allowed now, so this field is deprecated and is planned for removal in a future version of the API.

### 12.1.4. .spec.agent.ebpf

#### Description

**ebpf** describes the settings related to the eBPF-based flow reporter when **spec.agent.type** is set to **eBPF**.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>advanced</b>	<b>object</b>	<b>advanced</b> allows setting some aspects of the internal configuration of the eBPF agent. This section is aimed mostly for debugging and fine-grained performance optimizations, such as <b>GOGC</b> and <b>GOMAXPROCS</b> env vars. Set these values at your own risk.
<b>cacheActiveTimeout</b>	<b>string</b>	<b>cacheActiveTimeout</b> is the max period during which the reporter aggregates flows before sending. Increasing <b>cacheMaxFlows</b> and <b>cacheActiveTimeout</b> can decrease the network traffic overhead and the CPU load, however you can expect higher memory consumption and an increased latency in the flow collection.
<b>cacheMaxFlows</b>	<b>integer</b>	<b>cacheMaxFlows</b> is the max number of flows in an aggregate; when reached, the reporter sends the flows. Increasing <b>cacheMaxFlows</b> and <b>cacheActiveTimeout</b> can decrease the network traffic overhead and the CPU load, however you can expect higher memory consumption and an increased latency in the flow collection.
<b>excludeInterfaces</b>	<b>array (string)</b>	<b>excludeInterfaces</b> contains the interface names that are excluded from flow tracing. An entry enclosed by slashes, such as <b>/br-</b> , is matched as a regular expression. Otherwise it is matched as a case-sensitive string.

Property	Type	Description
<b>features</b>	<b>array (string)</b>	<p>List of additional features to enable. They are all disabled by default. Enabling additional features might have performance impacts. Possible values are:</p> <ul style="list-style-type: none"> <li>- <b>PacketDrop</b>: enable the packets drop flows logging feature. This feature requires mounting the kernel debug filesystem, so the eBPF pod has to run as privileged. If the <b>spec.agent.ebpf.privileged</b> parameter is not set, an error is reported.</li> <li>- <b>DNSTracking</b>: enable the DNS tracking feature.</li> <li>- <b>FlowRTT</b>: enable flow latency (sRTT) extraction in the eBPF agent from TCP traffic.</li> </ul>
<b>flowFilter</b>	<b>object</b>	<b>flowFilter</b> defines the eBPF agent configuration regarding flow filtering.
<b>imagePullPolicy</b>	<b>string</b>	<b>imagePullPolicy</b> is the Kubernetes pull policy for the image defined above
<b>interfaces</b>	<b>array (string)</b>	<b>interfaces</b> contains the interface names from where flows are collected. If empty, the agent fetches all the interfaces in the system, excepting the ones listed in <b>excludeInterfaces</b> . An entry enclosed by slashes, such as <b>/br-/,</b> is matched as a regular expression. Otherwise it is matched as a case-sensitive string.
<b>kafkaBatchSize</b>	<b>integer</b>	<b>kafkaBatchSize</b> limits the maximum size of a request in bytes before being sent to a partition. Ignored when not using Kafka. Default: 1MB.

Property	Type	Description
<b>logLevel</b>	<b>string</b>	<b>logLevel</b> defines the log level for the Network Observability eBPF Agent
<b>metrics</b>	<b>object</b>	<b>metrics</b> defines the eBPF agent configuration regarding metrics.
<b>privileged</b>	<b>boolean</b>	Privileged mode for the eBPF Agent container. When ignored or set to <b>false</b> , the operator sets granular capabilities (BPF, PERFMON, NET_ADMIN, SYS_RESOURCE) to the container. If for some reason these capabilities cannot be set, such as if an old kernel version not knowing CAP_BPF is in use, then you can turn on this mode for more global privileges. Some agent features require the privileged mode, such as packet drops tracking (see <b>features</b> ) and SR-IOV support.
<b>resources</b>	<b>object</b>	<b>resources</b> are the compute resources required by this container. For more information, see <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>sampling</b>	<b>integer</b>	Sampling rate of the flow reporter. 100 means one flow on 100 is sent. 0 or 1 means all flows are sampled.

### 12.1.5. .spec.agent.ebpf.advanced

#### Description

**advanced** allows setting some aspects of the internal configuration of the eBPF agent. This section is aimed mostly for debugging and fine-grained performance optimizations, such as **GOGC** and **GOMAXPROCS** env vars. Set these values at your own risk.

#### Type

**object**

Property	Type	Description
<b>env</b>	<b>object (string)</b>	<b>env</b> allows passing custom environment variables to underlying components. Useful for passing some very concrete performance-tuning options, such as <b>GOGC</b> and <b>GOMAXPROCS</b> , that should not be publicly exposed as part of the FlowCollector descriptor, as they are only useful in edge debug or support scenarios.
<b>scheduling</b>	<b>object</b>	scheduling controls how the pods are scheduled on nodes.

### 12.1.6. .spec.agent.ebpf.advanced.scheduling

#### Description

scheduling controls how the pods are scheduled on nodes.

#### Type

**object**

Property	Type	Description
<b>affinity</b>	<b>object</b>	If specified, the pod's scheduling constraints. For documentation, refer to <a href="https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling">https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling</a> .
<b>nodeSelector</b>	<b>object (string)</b>	<b>nodeSelector</b> allows scheduling of pods only onto nodes that have each of the specified labels. For documentation, refer to <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a> .

Property	Type	Description
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. For documentation, refer to <a href="https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/#how-to-use-priority-and-preemption">https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/#how-to-use-priority-and-preemption</a> . If not specified, default priority is used, or zero if there is no default.
<b>tolerations</b>	<b>array</b>	<b>tolerations</b> is a list of tolerations that allow the pod to schedule onto nodes with matching taints. For documentation, refer to <a href="https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling">https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling</a> .

### 12.1.7. .spec.agent.ebpf.advanced.scheduling.affinity

#### Description

If specified, the pod's scheduling constraints. For documentation, refer to <https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling>.

#### Type

**object**

### 12.1.8. .spec.agent.ebpf.advanced.scheduling.tolerations

#### Description

**tolerations** is a list of tolerations that allow the pod to schedule onto nodes with matching taints. For documentation, refer to <https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling>.

#### Type

**array**

### 12.1.9. .spec.agent.ebpf.flowFilter

#### Description

**flowFilter** defines the eBPF agent configuration regarding flow filtering.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>action</b>	<b>string</b>	<b>action</b> defines the action to perform on the flows that match the filter.
<b>cidr</b>	<b>string</b>	<b>cidr</b> defines the IP CIDR to filter flows by. Examples: <b>10.10.10.0/24</b> or <b>100:100:100:100::/64</b>
<b>destPorts</b>	<b>integer-or-string</b>	<b>destPorts</b> defines the destination ports to filter flows by. To filter a single port, set a single port as an integer value. For example: <b>destPorts: 80</b> . To filter a range of ports, use a "start-end" range in string format. For example: <b>destPorts: "80-100"</b> .
<b>direction</b>	<b>string</b>	<b>direction</b> defines the direction to filter flows by.
<b>enable</b>	<b>boolean</b>	Set <b>enable</b> to <b>true</b> to enable the eBPF flow filtering feature.
<b>icmpCode</b>	<b>integer</b>	<b>icmpCode</b> , for Internet Control Message Protocol (ICMP) traffic, defines the ICMP code to filter flows by.
<b>icmpType</b>	<b>integer</b>	<b>icmpType</b> , for ICMP traffic, defines the ICMP type to filter flows by.
<b>peerIP</b>	<b>string</b>	<b>peerIP</b> defines the IP address to filter flows by. Example: <b>10.10.10.10</b> .
<b>ports</b>	<b>integer-or-string</b>	<b>ports</b> defines the ports to filter flows by. It is used both for source and destination ports. To filter a single port, set a single port as an integer value. For example: <b>ports: 80</b> . To filter a range of ports, use a "start-end" range in string format. For example: <b>ports: "80-100"</b> .
<b>protocol</b>	<b>string</b>	<b>protocol</b> defines the protocol to filter flows by.

Property	Type	Description
<b>sourcePorts</b>	<b>integer-or-string</b>	<b>sourcePorts</b> defines the source ports to filter flows by. To filter a single port, set a single port as an integer value. For example: <b>sourcePorts: 80</b> . To filter a range of ports, use a "start-end" range in string format. For example: <b>sourcePorts: "80-100"</b> .

### 12.1.10. .spec.agent.ebpf.metrics

#### Description

**metrics** defines the eBPF agent configuration regarding metrics.

#### Type

**object**

Property	Type	Description
<b>disableAlerts</b>	<b>array (string)</b>	<b>disableAlerts</b> is a list of alerts that should be disabled. Possible values are:  <b>NetObservDroppedFlows</b> , which is triggered when the eBPF agent is dropping flows, such as when the BPF hashmap is full or the capacity limiter is being triggered.
<b>enable</b>	<b>boolean</b>	Set <b>enable</b> to <b>false</b> to disable eBPF agent metrics collection. It is enabled by default.
<b>server</b>	<b>object</b>	Metrics server endpoint configuration for the Prometheus scraper.

### 12.1.11. .spec.agent.ebpf.metrics.server

#### Description

Metrics server endpoint configuration for the Prometheus scraper.

#### Type

**object**



Property	Type	Description
<b>port</b>	<b>integer</b>	The metrics server HTTP port.
<b>tls</b>	<b>object</b>	TLS configuration.

### 12.1.12. .spec.agent.ebpf.metrics.server.tls

#### Description

TLS configuration.

#### Type

**object**

Property	Type	Description
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the provided certificate. If set to <b>true</b> , the <b>providedCaFile</b> field is ignored.
<b>provided</b>	<b>object</b>	TLS configuration when <b>type</b> is set to <b>Provided</b> .
<b>providedCaFile</b>	<b>object</b>	Reference to the CA file when <b>type</b> is set to <b>Provided</b> .
<b>type</b>	<b>string</b>	Select the type of TLS configuration: <ul style="list-style-type: none"> <li>- <b>Disabled</b> (default) to not configure TLS for the endpoint.</li> <li>- <b>Provided</b> to manually provide cert file and a key file.</li> <li>[Unsupported (*)].</li> <li>- <b>Auto</b> to use OpenShift Container Platform auto generated certificate using annotations.</li> </ul>

### 12.1.13. .spec.agent.ebpf.metrics.server.tls.provided

#### Description

TLS configuration when **type** is set to **Provided**.

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

#### 12.1.14. .spec.agent.ebpf.metrics.server.tls.providedCaFile

##### Description

Reference to the CA file when **type** is set to **Provided**.

##### Type

**object**

Property	Type	Description
<b>file</b>	<b>string</b>	File name within the config map or secret.
<b>name</b>	<b>string</b>	Name of the config map or secret containing the file.

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing the file. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the file reference: "configmap" or "secret".

### 12.1.15. .spec.agent.ebpf.resources

#### Description

**resources** are the compute resources required by this container. For more information, see <https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/>

#### Type

**object**

Property	Type	Description
<b>limits</b>	<b>integer-or-string</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>integer-or-string</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

### 12.1.16. .spec.consolePlugin

#### Description

**consolePlugin** defines the settings related to the OpenShift Container Platform Console plugin, when available.

Type  
object

Property	Type	Description
<b>advanced</b>	<b>object</b>	<b>advanced</b> allows setting some aspects of the internal configuration of the console plugin. This section is aimed mostly for debugging and fine-grained performance optimizations, such as <b>GOGC</b> and <b>GOMAXPROCS</b> env vars. Set these values at your own risk.
<b>autoscaler</b>	<b>object</b>	<b>autoscaler</b> spec of a horizontal pod autoscaler to set up for the plugin Deployment. Refer to HorizontalPodAutoscaler documentation (autoscaling/v2).
<b>enable</b>	<b>boolean</b>	Enables the console plugin deployment.
<b>imagePullPolicy</b>	<b>string</b>	<b>imagePullPolicy</b> is the Kubernetes pull policy for the image defined above
<b>logLevel</b>	<b>string</b>	<b>logLevel</b> for the console plugin backend
<b>portNaming</b>	<b>object</b>	<b>portNaming</b> defines the configuration of the port-to-service name translation
<b>quickFilters</b>	<b>array</b>	<b>quickFilters</b> configures quick filter presets for the Console plugin
<b>replicas</b>	<b>integer</b>	<b>replicas</b> defines the number of replicas (pods) to start.
<b>resources</b>	<b>object</b>	<b>resources</b> , in terms of compute resources, required by this container. For more information, see <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

## 12.1.17. .spec.consolePlugin.advanced

## Description

**advanced** allows setting some aspects of the internal configuration of the console plugin. This section is aimed mostly for debugging and fine-grained performance optimizations, such as **GOGC** and **GOMAXPROCS** env vars. Set these values at your own risk.

## Type

**object**

Property	Type	Description
<b>args</b>	<b>array (string)</b>	<b>args</b> allows passing custom arguments to underlying components. Useful for overriding some parameters, such as a URL or a configuration path, that should not be publicly exposed as part of the FlowCollector descriptor, as they are only useful in edge debug or support scenarios.
<b>env</b>	<b>object (string)</b>	<b>env</b> allows passing custom environment variables to underlying components. Useful for passing some very concrete performance-tuning options, such as <b>GOGC</b> and <b>GOMAXPROCS</b> , that should not be publicly exposed as part of the FlowCollector descriptor, as they are only useful in edge debug or support scenarios.
<b>port</b>	<b>integer</b>	<b>port</b> is the plugin service port. Do not use 9002, which is reserved for metrics.
<b>register</b>	<b>boolean</b>	<b>register</b> allows, when set to <b>true</b> , to automatically register the provided console plugin with the OpenShift Container Platform Console operator. When set to <b>false</b> , you can still register it manually by editing <code>console.operator.openshift.io/cluster</code> with the following command: <b>oc patch console.operator.openshift.io cluster --type='json' -p '[{"op": "add", "path": "/spec/plugins/-", "value": "netobserv-plugin"}]'</b>

Property	Type	Description
<b>scheduling</b>	<b>object</b>	<b>scheduling</b> controls how the pods are scheduled on nodes.

### 12.1.18. .spec.consolePlugin.advanced.scheduling

#### Description

**scheduling** controls how the pods are scheduled on nodes.

#### Type

**object**

Property	Type	Description
<b>affinity</b>	<b>object</b>	If specified, the pod's scheduling constraints. For documentation, refer to <a href="https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling">https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling</a> .
<b>nodeSelector</b>	<b>object (string)</b>	<b>nodeSelector</b> allows scheduling of pods only onto nodes that have each of the specified labels. For documentation, refer to <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a> .
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. For documentation, refer to <a href="https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/#how-to-use-priority-and-preemption">https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/#how-to-use-priority-and-preemption</a> . If not specified, default priority is used, or zero if there is no default.
<b>tolerations</b>	<b>array</b>	<b>tolerations</b> is a list of tolerations that allow the pod to schedule onto nodes with matching taints. For documentation, refer to <a href="https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling">https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling</a> .

### 12.1.19. .spec.consolePlugin.advanced.scheduling.affinity

**Description**

If specified, the pod's scheduling constraints. For documentation, refer to <https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling>.

**Type**

**object**

**12.1.20. .spec.consolePlugin.advanced.scheduling.tolerations****Description**

**tolerations** is a list of tolerations that allow the pod to schedule onto nodes with matching taints. For documentation, refer to <https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling>.

**Type**

**array**

**12.1.21. .spec.consolePlugin.autoscaler****Description**

**autoscaler** spec of a horizontal pod autoscaler to set up for the plugin Deployment. Refer to HorizontalPodAutoscaler documentation (autoscaling/v2).

**Type**

**object**

**12.1.22. .spec.consolePlugin.portNaming****Description**

**portNaming** defines the configuration of the port-to-service name translation

**Type**

**object**

Property	Type	Description
<b>enable</b>	<b>boolean</b>	Enable the console plugin port-to-service name translation
<b>portNames</b>	<b>object (string)</b>	<b>portNames</b> defines additional port names to use in the console, for example, <b>portNames: {"3100": "loki"}</b> .

**12.1.23. .spec.consolePlugin.quickFilters****Description**

**quickFilters** configures quick filter presets for the Console plugin

**Type**

**array**

### 12.1.24. .spec.consolePlugin.quickFilters[]

#### Description

**QuickFilter** defines preset configuration for Console's quick filters

#### Type

**object**

#### Required

- **filter**
- **name**

Property	Type	Description
<b>default</b>	<b>boolean</b>	<b>default</b> defines whether this filter should be active by default or not
<b>filter</b>	<b>object (string)</b>	<b>filter</b> is a set of keys and values to be set when this filter is selected. Each key can relate to a list of values using a coma-separated string, for example, <b>filter:</b> <b>{"src_namespace": "namespace1,namespace2"}</b> .
<b>name</b>	<b>string</b>	Name of the filter, that is displayed in the Console

### 12.1.25. .spec.consolePlugin.resources

#### Description

**resources**, in terms of compute resources, required by this container. For more information, see <https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/>

#### Type

**object**

Property	Type	Description
<b>limits</b>	<b>integer-or-string</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>



Property	Type	Description
<b>requests</b>	<b>integer-or-string</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manager-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manager-resources-containers/</a>

### 12.1.26. .spec.exporters

#### Description

**exporters** define additional optional exporters for custom consumption or storage.

#### Type

**array**

### 12.1.27. .spec.exporters[]

#### Description

**FlowCollectorExporter** defines an additional exporter to send enriched flows to.

#### Type

**object**

#### Required

- **type**

Property	Type	Description
<b>ipfix</b>	<b>object</b>	IPFIX configuration, such as the IP address and port to send enriched IPFIX flows to.
<b>kafka</b>	<b>object</b>	Kafka configuration, such as the address and topic, to send enriched flows to.
<b>type</b>	<b>string</b>	<b>type</b> selects the type of exporters. The available options are <b>Kafka</b> and <b>IPFIX</b> .

### 12.1.28. .spec.exporters[].ipfix

**Description**

IPFIX configuration, such as the IP address and port to send enriched IPFIX flows to.

**Type**

**object**

**Required**

- **targetHost**
- **targetPort**

Property	Type	Description
<b>targetHost</b>	<b>string</b>	Address of the IPFIX external receiver
<b>targetPort</b>	<b>integer</b>	Port for the IPFIX external receiver
<b>transport</b>	<b>string</b>	Transport protocol ( <b>TCP</b> or <b>UDP</b> ) to be used for the IPFIX connection, defaults to <b>TCP</b> .

**12.1.29. .spec.exporters[].kafka****Description**

Kafka configuration, such as the address and topic, to send enriched flows to.

**Type**

**object**

**Required**

- **address**
- **topic**

Property	Type	Description
<b>address</b>	<b>string</b>	Address of the Kafka server
<b>sasl</b>	<b>object</b>	SASL authentication configuration. [Unsupported (*)].
<b>tls</b>	<b>object</b>	TLS client configuration. When using TLS, verify that the address matches the Kafka port used for TLS, generally 9093.

Property	Type	Description
<b>topic</b>	<b>string</b>	Kafka topic to use. It must exist. Network Observability does not create it.

### 12.1.30. .spec.exporters[].kafka.sasl

#### Description

SASL authentication configuration. [Unsupported (\*)].

#### Type

**object**

Property	Type	Description
<b>clientIDReference</b>	<b>object</b>	Reference to the secret or config map containing the client ID
<b>clientSecretReference</b>	<b>object</b>	Reference to the secret or config map containing the client secret
<b>type</b>	<b>string</b>	Type of SASL authentication to use, or <b>Disabled</b> if SASL is not used

### 12.1.31. .spec.exporters[].kafka.sasl.clientIDReference

#### Description

Reference to the secret or config map containing the client ID

#### Type

**object**

Property	Type	Description
<b>file</b>	<b>string</b>	File name within the config map or secret.
<b>name</b>	<b>string</b>	Name of the config map or secret containing the file.

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing the file. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the file reference: "configmap" or "secret".

### 12.1.32. `.spec.exporters[].kafka.sasl.clientSecretReference`

#### Description

Reference to the secret or config map containing the client secret

#### Type

**object**

Property	Type	Description
<b>file</b>	<b>string</b>	File name within the config map or secret.
<b>name</b>	<b>string</b>	Name of the config map or secret containing the file.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing the file. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the file reference: "configmap" or "secret".

### 12.1.33. `.spec.exporters[].kafka.tls`

#### Description

TLS client configuration. When using TLS, verify that the address matches the Kafka port used for TLS, generally 9093.

Type

object

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### 12.1.34. .spec.exporters[.kafka.tls.caCert

Description

**caCert** defines the reference of the certificate for the Certificate Authority

Type

object

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.35. .spec.exporters[].kafka.tls.userCert

#### Description

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.

Property	Type	Description
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.36. .spec.kafka

#### Description

Kafka configuration, allowing to use Kafka as a broker as part of the flow collection pipeline. Available when the **spec.deploymentModel** is **Kafka**.

#### Type

**object**

#### Required

- **address**
- **topic**

Property	Type	Description
<b>address</b>	<b>string</b>	Address of the Kafka server
<b>sasl</b>	<b>object</b>	SASL authentication configuration. [Unsupported (*)].
<b>tls</b>	<b>object</b>	TLS client configuration. When using TLS, verify that the address matches the Kafka port used for TLS, generally 9093.
<b>topic</b>	<b>string</b>	Kafka topic to use. It must exist. Network Observability does not create it.

### 12.1.37. .spec.kafka.sasl

#### Description

SASL authentication configuration. [Unsupported (\*)].

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>clientIDReference</b>	<b>object</b>	Reference to the secret or config map containing the client ID
<b>clientSecretReference</b>	<b>object</b>	Reference to the secret or config map containing the client secret
<b>type</b>	<b>string</b>	Type of SASL authentication to use, or <b>Disabled</b> if SASL is not used

### 12.1.38. .spec.kafka.sasl.clientIDReference

#### Description

Reference to the secret or config map containing the client ID

#### Type

**object**

Property	Type	Description
<b>file</b>	<b>string</b>	File name within the config map or secret.
<b>name</b>	<b>string</b>	Name of the config map or secret containing the file.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing the file. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the file reference: "configmap" or "secret".

### 12.1.39. .spec.kafka.sasl.clientSecretReference

#### Description

Reference to the secret or config map containing the client secret

#### Type

**object**



Property	Type	Description
<b>file</b>	<b>string</b>	File name within the config map or secret.
<b>name</b>	<b>string</b>	Name of the config map or secret containing the file.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing the file. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the file reference: "configmap" or "secret".

#### 12.1.40. .spec.kafka.tls

##### Description

TLS client configuration. When using TLS, verify that the address matches the Kafka port used for TLS, generally 9093.

##### Type

**object**

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### 12.1.41. .spec.kafka.tls.caCert

**Description**

**caCert** defines the reference of the certificate for the Certificate Authority

**Type**

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

**12.1.42. .spec.kafka.tls.userCert****Description**

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

**Type**

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.

Property	Type	Description
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.43. .spec.loki

#### Description

**loki**, the flow store, client settings.

#### Type

**object**

Property	Type	Description
<b>advanced</b>	<b>object</b>	<b>advanced</b> allows setting some aspects of the internal configuration of the Loki clients. This section is aimed mostly for debugging and fine-grained performance optimizations.

Property	Type	Description
<b>enable</b>	<b>boolean</b>	Set <b>enable</b> to <b>true</b> to store flows in Loki. The Console plugin can use either Loki or Prometheus as a data source for metrics (see also <b>spec.prometheus.querier</b> ), or both. Not all queries are transposable from Loki to Prometheus. Hence, if Loki is disabled, some features of the plugin are disabled as well, such as getting per-pod information or viewing raw flows. If both Prometheus and Loki are enabled, Prometheus takes precedence and Loki is used as a fallback for queries that Prometheus cannot handle. If they are both disabled, the Console plugin is not deployed.
<b>lokiStack</b>	<b>object</b>	Loki configuration for <b>LokiStack</b> mode. This is useful for an easy Loki Operator configuration. It is ignored for other modes.
<b>manual</b>	<b>object</b>	Loki configuration for <b>Manual</b> mode. This is the most flexible configuration. It is ignored for other modes.
<b>microservices</b>	<b>object</b>	Loki configuration for <b>Microservices</b> mode. Use this option when Loki is installed using the microservices deployment mode ( <a href="https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#microservices-mode">https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#microservices-mode</a> ). It is ignored for other modes.

Property	Type	Description
<b>mode</b>	<b>string</b>	<p><b>mode</b> must be set according to the installation mode of Loki:</p> <ul style="list-style-type: none"> <li>- Use <b>LokiStack</b> when Loki is managed using the Loki Operator</li> <li>- Use <b>Monolithic</b> when Loki is installed as a monolithic workload</li> <li>- Use <b>Microservices</b> when Loki is installed as microservices, but without Loki Operator</li> <li>- Use <b>Manual</b> if none of the options above match your setup</li> </ul>
<b>monolithic</b>	<b>object</b>	<p>Loki configuration for <b>Monolithic</b> mode. Use this option when Loki is installed using the monolithic deployment mode (<a href="https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#monolithic-mode">https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#monolithic-mode</a>). It is ignored for other modes.</p>
<b>readTimeout</b>	<b>string</b>	<p><b>readTimeout</b> is the maximum console plugin loki query total time limit. A timeout of zero means no timeout.</p>
<b>writeBatchSize</b>	<b>integer</b>	<p><b>writeBatchSize</b> is the maximum batch size (in bytes) of Loki logs to accumulate before sending.</p>
<b>writeBatchWait</b>	<b>string</b>	<p><b>writeBatchWait</b> is the maximum time to wait before sending a Loki batch.</p>
<b>writeTimeout</b>	<b>string</b>	<p><b>writeTimeout</b> is the maximum Loki time connection / request limit. A timeout of zero means no timeout.</p>

#### 12.1.44. .spec.loki.advanced

##### Description

**advanced** allows setting some aspects of the internal configuration of the Loki clients. This section is aimed mostly for debugging and fine-grained performance optimizations.

## Type

object

Property	Type	Description
<b>staticLabels</b>	<b>object (string)</b>	<b>staticLabels</b> is a map of common labels to set on each flow in Loki storage.
<b>writeMaxBackoff</b>	<b>string</b>	<b>writeMaxBackoff</b> is the maximum backoff time for Loki client connection between retries.
<b>writeMaxRetries</b>	<b>integer</b>	<b>writeMaxRetries</b> is the maximum number of retries for Loki client connections.
<b>writeMinBackoff</b>	<b>string</b>	<b>writeMinBackoff</b> is the initial backoff time for Loki client connection between retries.

## 12.1.45. .spec.loki.lokiStack

## Description

Loki configuration for **LokiStack** mode. This is useful for an easy Loki Operator configuration. It is ignored for other modes.

## Type

object

Property	Type	Description
<b>name</b>	<b>string</b>	Name of an existing LokiStack resource to use.
<b>namespace</b>	<b>string</b>	Namespace where this <b>LokiStack</b> resource is located. If omitted, it is assumed to be the same as <b>spec.namespace</b> .

## 12.1.46. .spec.loki.manual

## Description

Loki configuration for **Manual** mode. This is the most flexible configuration. It is ignored for other modes.

## Type

object

Property	Type	Description
<b>authToken</b>	<b>string</b>	<p><b>authToken</b> describes the way to get a token to authenticate to Loki.</p> <ul style="list-style-type: none"> <li>- <b>Disabled</b> does not send any token with the request.</li> <li>- <b>Forward</b> forwards the user token for authorization.</li> <li>- <b>Host</b> [deprecated (*)] - uses the local pod service account to authenticate to Loki.</li> </ul> <p>When using the Loki Operator, this must be set to <b>Forward</b>.</p>
<b>ingesterUrl</b>	<b>string</b>	<p><b>ingesterUrl</b> is the address of an existing Loki ingester service to push the flows to. When using the Loki Operator, set it to the Loki gateway service with the <b>network</b> tenant set in path, for example <a href="https://loki-gateway-http.netobserv.svc:8080/api/logs/v1/network">https://loki-gateway-http.netobserv.svc:8080/api/logs/v1/network</a>.</p>
<b>querierUrl</b>	<b>string</b>	<p><b>querierUrl</b> specifies the address of the Loki querier service. When using the Loki Operator, set it to the Loki gateway service with the <b>network</b> tenant set in path, for example <a href="https://loki-gateway-http.netobserv.svc:8080/api/logs/v1/network">https://loki-gateway-http.netobserv.svc:8080/api/logs/v1/network</a>.</p>
<b>statusTls</b>	<b>object</b>	<p>TLS client configuration for Loki status URL.</p>

Property	Type	Description
<b>statusUrl</b>	<b>string</b>	<b>statusUrl</b> specifies the address of the Loki <b>/ready</b> , <b>/metrics</b> and <b>/config</b> endpoints, in case it is different from the Loki querier URL. If empty, the <b>querierUrl</b> value is used. This is useful to show error messages and some context in the frontend. When using the Loki Operator, set it to the Loki HTTP query frontend service, for example <a href="https://loki-query-frontend-http.netobserv.svc:3100/">https://loki-query-frontend-http.netobserv.svc:3100/</a> . <b>statusTLS</b> configuration is used when <b>statusUrl</b> is set.
<b>tenantID</b>	<b>string</b>	<b>tenantID</b> is the Loki <b>X-Scope-OrgID</b> that identifies the tenant for each request. When using the Loki Operator, set it to <b>network</b> , which corresponds to a special tenant mode.
<b>tls</b>	<b>object</b>	TLS client configuration for Loki URL.

### 12.1.47. .spec.loki.manual.statusTls

#### Description

TLS client configuration for Loki status URL.

#### Type

**object**

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.



Property	Type	Description
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

### 12.1.48. .spec.loki.manual.statusTls.caCert

#### Description

**caCert** defines the reference of the certificate for the Certificate Authority

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.49. .spec.loki.manual.statusTls.userCert

#### Description

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.50. .spec.loki.manual.tls

Description

TLS client configuration for Loki URL.

Type

**object**

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS

Property	Type	Description
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

### 12.1.51. .spec.loki.manual.tls.caCert

#### Description

**caCert** defines the reference of the certificate for the Certificate Authority

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.52. .spec.loki.manual.tls.userCert

**Description**

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

**Type**

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

**12.1.53. .spec.loki.microservices****Description**

Loki configuration for **Microservices** mode. Use this option when Loki is installed using the microservices deployment mode

(<https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#microservices-mode>). It is ignored for other modes.

**Type**

**object**

Property	Type	Description
<b>ingesterUrl</b>	<b>string</b>	<b>ingesterUrl</b> is the address of an existing Loki ingester service to push the flows to.

Property	Type	Description
<b>querierUrl</b>	<b>string</b>	<b>querierURL</b> specifies the address of the Loki querier service.
<b>tenantID</b>	<b>string</b>	<b>tenantID</b> is the Loki <b>X-Scope-OrgID</b> header that identifies the tenant for each request.
<b>tls</b>	<b>object</b>	TLS client configuration for Loki URL.

#### 12.1.54. .spec.loki.microservices.tls

##### Description

TLS client configuration for Loki URL.

##### Type

**object**

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### 12.1.55. .spec.loki.microservices.tls.caCert

##### Description

**caCert** defines the reference of the certificate for the Certificate Authority

##### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.56. .spec.loki.microservices.tls.userCert

#### Description

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.

Property	Type	Description
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.57. .spec.loki.monolithic

#### Description

Loki configuration for **Monolithic** mode. Use this option when Loki is installed using the monolithic deployment mode (<https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#monolithic-mode>). It is ignored for other modes.

#### Type

**object**

Property	Type	Description
<b>tenantID</b>	<b>string</b>	<b>tenantID</b> is the Loki <b>X-Scope-OrgID</b> header that identifies the tenant for each request.
<b>tls</b>	<b>object</b>	TLS client configuration for Loki URL.
<b>url</b>	<b>string</b>	<b>url</b> is the unique address of an existing Loki service that points to both the ingester and the querier.

### 12.1.58. .spec.loki.monolithic.tls

#### Description

TLS client configuration for Loki URL.

#### Type

**object**

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

### 12.1.59. .spec.loki.monolithic.tls.caCert

#### Description

**caCert** defines the reference of the certificate for the Certificate Authority

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.



Property	Type	Description
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.60. .spec.loki.monolithic.tls.userCert

#### Description

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.61. .spec.processor

#### Description

**processor** defines the settings of the component that receives the flows from the agent, enriches them, generates metrics, and forwards them to the Loki persistence layer and/or any available exporter.

#### Type

**object**

Property	Type	Description
<b>addZone</b>	<b>boolean</b>	<b>addZone</b> allows availability zone awareness by labelling flows with their source and destination zones. This feature requires the "topology.kubernetes.io/zone" label to be set on nodes.
<b>advanced</b>	<b>object</b>	<b>advanced</b> allows setting some aspects of the internal configuration of the flow processor. This section is aimed mostly for debugging and fine-grained performance optimizations, such as <b>GOGC</b> and <b>GOMAXPROCS</b> env vars. Set these values at your own risk.
<b>clusterName</b>	<b>string</b>	<b>clusterName</b> is the name of the cluster to appear in the flows data. This is useful in a multi-cluster context. When using OpenShift Container Platform, leave empty to make it automatically determined.
<b>imagePullPolicy</b>	<b>string</b>	<b>imagePullPolicy</b> is the Kubernetes pull policy for the image defined above
<b>kafkaConsumerAutoscaler</b>	<b>object</b>	<b>kafkaConsumerAutoscaler</b> is the spec of a horizontal pod autoscaler to set up for <b>flowlogs-pipeline-transformer</b> , which consumes Kafka messages. This setting is ignored when Kafka is disabled. Refer to HorizontalPodAutoscaler documentation (autoscaling/v2).
<b>kafkaConsumerBatchSize</b>	<b>integer</b>	<b>kafkaConsumerBatchSize</b> indicates to the broker the maximum batch size, in bytes, that the consumer accepts. Ignored when not using Kafka. Default: 10MB.

Property	Type	Description
<b>kafkaConsumerQueueCapacity</b>	<b>integer</b>	<b>kafkaConsumerQueueCapacity</b> defines the capacity of the internal message queue used in the Kafka consumer client. Ignored when not using Kafka.
<b>kafkaConsumerReplicas</b>	<b>integer</b>	<b>kafkaConsumerReplicas</b> defines the number of replicas (pods) to start for <b>flowlogs-pipeline-transformer</b> , which consumes Kafka messages. This setting is ignored when Kafka is disabled.
<b>logLevel</b>	<b>string</b>	<b>logLevel</b> of the processor runtime
<b>logTypes</b>	<b>string</b>	<p><b>logTypes</b> defines the desired record types to generate. Possible values are:</p> <ul style="list-style-type: none"> <li>- <b>Flows</b> (default) to export regular network flows</li> <li>- <b>Conversations</b> to generate events for started conversations, ended conversations as well as periodic "tick" updates</li> <li>- <b>EndedConversations</b> to generate only ended conversations events</li> <li>- <b>All</b> to generate both network flows and all conversations events</li> </ul>
<b>metrics</b>	<b>object</b>	<b>Metrics</b> define the processor configuration regarding metrics
<b>multiClusterDeployment</b>	<b>boolean</b>	Set <b>multiClusterDeployment</b> to <b>true</b> to enable multi clusters feature. This adds <b>clusterName</b> label to flows data

Property	Type	Description
<b>resources</b>	<b>object</b>	<b>resources</b> are the compute resources required by this container. For more information, see <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>subnetLabels</b>	<b>object</b>	<b>subnetLabels</b> allows to define custom labels on subnets and IPs or to enable automatic labelling of recognized subnets in OpenShift Container Platform, which is used to identify cluster external traffic. When a subnet matches the source or destination IP of a flow, a corresponding field is added: <b>SrcSubnetLabel</b> or <b>DstSubnetLabel</b> .

## 12.1.62. .spec.processor.advanced

### Description

**advanced** allows setting some aspects of the internal configuration of the flow processor. This section is aimed mostly for debugging and fine-grained performance optimizations, such as **GOGC** and **GOMAXPROCS** env vars. Set these values at your own risk.

### Type

**object**

Property	Type	Description
<b>conversationEndTimeout</b>	<b>string</b>	<b>conversationEndTimeout</b> is the time to wait after a network flow is received, to consider the conversation ended. This delay is ignored when a FIN packet is collected for TCP flows (see <b>conversationTerminatingTimeout</b> instead).
<b>conversationHeartbeatInterval</b>	<b>string</b>	<b>conversationHeartbeatInterval</b> is the time to wait between "tick" events of a conversation

Property	Type	Description
<b>conversationTerminatingTimeout</b>	<b>string</b>	<b>conversationTerminatingTimeout</b> is the time to wait from detected FIN flag to end a conversation. Only relevant for TCP flows.
<b>dropUnusedFields</b>	<b>boolean</b>	<b>dropUnusedFields</b> [deprecated (*)] this setting is not used anymore.
<b>enableKubeProbes</b>	<b>boolean</b>	<b>enableKubeProbes</b> is a flag to enable or disable Kubernetes liveness and readiness probes
<b>env</b>	<b>object (string)</b>	<b>env</b> allows passing custom environment variables to underlying components. Useful for passing some very concrete performance-tuning options, such as <b>GOGC</b> and <b>GOMAXPROCS</b> , that should not be publicly exposed as part of the FlowCollector descriptor, as they are only useful in edge debug or support scenarios.
<b>healthPort</b>	<b>integer</b>	<b>healthPort</b> is a collector HTTP port in the Pod that exposes the health check API
<b>port</b>	<b>integer</b>	Port of the flow collector (host port). By convention, some values are forbidden. It must be greater than 1024 and different from 4500, 4789 and 6081.
<b>profilePort</b>	<b>integer</b>	<b>profilePort</b> allows setting up a Go pprof profiler listening to this port
<b>scheduling</b>	<b>object</b>	scheduling controls how the pods are scheduled on nodes.

### 12.1.63. .spec.processor.advanced.scheduling

#### Description

scheduling controls how the pods are scheduled on nodes.

#### Type

**object**

Property	Type	Description
<b>affinity</b>	<b>object</b>	If specified, the pod's scheduling constraints. For documentation, refer to <a href="https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling">https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling</a> .
<b>nodeSelector</b>	<b>object (string)</b>	<b>nodeSelector</b> allows scheduling of pods only onto nodes that have each of the specified labels. For documentation, refer to <a href="https://kubernetes.io/docs/concepts/configuration/assign-pod-node/">https://kubernetes.io/docs/concepts/configuration/assign-pod-node/</a> .
<b>priorityClassName</b>	<b>string</b>	If specified, indicates the pod's priority. For documentation, refer to <a href="https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/#how-to-use-priority-and-preemption">https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/#how-to-use-priority-and-preemption</a> . If not specified, default priority is used, or zero if there is no default.
<b>tolerations</b>	<b>array</b>	<b>tolerations</b> is a list of tolerations that allow the pod to schedule onto nodes with matching taints. For documentation, refer to <a href="https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling">https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling</a> .

**12.1.64. .spec.processor.advanced.scheduling.affinity****Description**

If specified, the pod's scheduling constraints. For documentation, refer to <https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling>.

**Type**

**object**

**12.1.65. .spec.processor.advanced.scheduling.tolerations****Description**

**tolerations** is a list of tolerations that allow the pod to schedule onto nodes with matching taints. For documentation, refer to <https://kubernetes.io/docs/reference/kubernetes-api/workload-resources/pod-v1/#scheduling>.

Type

array

### 12.1.66. .spec.processor.kafkaConsumerAutoscaler

Description

**kafkaConsumerAutoscaler** is the spec of a horizontal pod autoscaler to set up for **flowlogs-pipeline-transformer**, which consumes Kafka messages. This setting is ignored when Kafka is disabled. Refer to HorizontalPodAutoscaler documentation (autoscaling/v2).

Type

object

### 12.1.67. .spec.processor.metrics

Description

**Metrics** define the processor configuration regarding metrics

Type

object

Property	Type	Description
<b>disableAlerts</b>	<b>array (string)</b>	<p><b>disableAlerts</b> is a list of alerts that should be disabled. Possible values are:</p> <p><b>NetObservNoFlows</b>, which is triggered when no flows are being observed for a certain period.</p> <p><b>NetObservLokiError</b>, which is triggered when flows are being dropped due to Loki errors.</p>

Property	Type	Description
<b>includeList</b>	<b>array (string)</b>	<p><b>includeList</b> is a list of metric names to specify which ones to generate. The names correspond to the names in Prometheus without the prefix. For example, <b>namespace_egress_packets_total</b> shows up as <b>netobserv_namespace_egress_packets_total</b> in Prometheus. Note that the more metrics you add, the bigger is the impact on Prometheus workload resources. Metrics enabled by default are:</p> <p><b>namespace_flows_total</b>,  <b>node_ingress_bytes_total</b>,  <b>workload_ingress_bytes_total</b>,  <b>namespace_drop_packets_total</b> (when <b>PacketDrop</b> feature is enabled),  <b>namespace_rtt_seconds</b> (when <b>FlowRTT</b> feature is enabled),  <b>namespace_dns_latency_seconds</b> (when <b>DNSTracking</b> feature is enabled). More information, with full list of available metrics: <a href="https://github.com/netobserv/network-observability-operator/blob/main/docs/Metrics.md">https://github.com/netobserv/network-observability-operator/blob/main/docs/Metrics.md</a></p>
<b>server</b>	<b>object</b>	Metrics server endpoint configuration for Prometheus scraper

### 12.1.68. .spec.processor.metrics.server

#### Description

Metrics server endpoint configuration for Prometheus scraper

#### Type

**object**

Property	Type	Description
<b>port</b>	<b>integer</b>	The metrics server HTTP port.



Property	Type	Description
<b>tls</b>	<b>object</b>	TLS configuration.

### 12.1.69. .spec.processor.metrics.server.tls

#### Description

TLS configuration.

#### Type

**object**

Property	Type	Description
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the provided certificate. If set to <b>true</b> , the <b>providedCaFile</b> field is ignored.
<b>provided</b>	<b>object</b>	TLS configuration when <b>type</b> is set to <b>Provided</b> .
<b>providedCaFile</b>	<b>object</b>	Reference to the CA file when <b>type</b> is set to <b>Provided</b> .
<b>type</b>	<b>string</b>	Select the type of TLS configuration: <ul style="list-style-type: none"> <li>- <b>Disabled</b> (default) to not configure TLS for the endpoint.</li> <li>- <b>Provided</b> to manually provide cert file and a key file.</li> <li>[Unsupported (*)].</li> <li>- <b>Auto</b> to use OpenShift Container Platform auto generated certificate using annotations.</li> </ul>

### 12.1.70. .spec.processor.metrics.server.tls.provided

#### Description

TLS configuration when **type** is set to **Provided**.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.71. .spec.processor.metrics.server.tls.providedCaFile

#### Description

Reference to the CA file when **type** is set to **Provided**.

#### Type

**object**

Property	Type	Description
<b>file</b>	<b>string</b>	File name within the config map or secret.
<b>name</b>	<b>string</b>	Name of the config map or secret containing the file.

Property	Type	Description
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing the file. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the file reference: "configmap" or "secret".

### 12.1.72. .spec.processor.resources

#### Description

**resources** are the compute resources required by this container. For more information, see <https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/>

#### Type

**object**

Property	Type	Description
<b>limits</b>	<b>integer-or-string</b>	Limits describes the maximum amount of compute resources allowed. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>
<b>requests</b>	<b>integer-or-string</b>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. Requests cannot exceed Limits. More info: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>

### 12.1.73. .spec.processor.subnetLabels

#### Description

**subnetLabels** allows to define custom labels on subnets and IPs or to enable automatic labelling of recognized subnets in OpenShift Container Platform, which is used to identify cluster external traffic. When a subnet matches the source or destination IP of a flow, a corresponding field is added:

**SrcSubnetLabel** or **DstSubnetLabel**.

Type

**object**

Property	Type	Description
<b>customLabels</b>	<b>array</b>	<b>customLabels</b> allows to customize subnets and IPs labelling, such as to identify cluster-external workloads or web services. If you enable <b>openShiftAutoDetect</b> , <b>customLabels</b> can override the detected subnets in case they overlap.
<b>openShiftAutoDetect</b>	<b>boolean</b>	<b>openShiftAutoDetect</b> allows, when set to <b>true</b> , to detect automatically the machines, pods and services subnets based on the OpenShift Container Platform install configuration and the Cluster Network Operator configuration. Indirectly, this is a way to accurately detect external traffic: flows that are not labeled for those subnets are external to the cluster. Enabled by default on OpenShift Container Platform.

#### 12.1.74. .spec.processor.subnetLabels.customLabels

Description

**customLabels** allows to customize subnets and IPs labelling, such as to identify cluster-external workloads or web services. If you enable **openShiftAutoDetect**, **customLabels** can override the detected subnets in case they overlap.

Type

**array**

#### 12.1.75. .spec.processor.subnetLabels.customLabels[]

Description

SubnetLabel allows to label subnets and IPs, such as to identify cluster-external workloads or web services.

Type

**object**

Property	Type	Description
<b>cidrs</b>	<b>array (string)</b>	List of CIDRs, such as <b>["1.2.3.4/32"]</b> .
<b>name</b>	<b>string</b>	Label name, used to flag matching flows.

### 12.1.76. .spec.prometheus

#### Description

**prometheus** defines Prometheus settings, such as querier configuration used to fetch metrics from the Console plugin.

#### Type

**object**

Property	Type	Description
<b>querier</b>	<b>object</b>	Prometheus querying configuration, such as client settings, used in the Console plugin.

### 12.1.77. .spec.prometheus.querier

#### Description

Prometheus querying configuration, such as client settings, used in the Console plugin.

#### Type

**object**

Property	Type	Description
----------	------	-------------

Property	Type	Description
<b>enable</b>	<b>boolean</b>	When <b>enable</b> is <b>true</b> , the Console plugin queries flow metrics from Prometheus instead of Loki whenever possible. It is enabled by default: set it to <b>false</b> to disable this feature. The Console plugin can use either Loki or Prometheus as a data source for metrics (see also <b>spec.loki</b> ), or both. Not all queries are transposable from Loki to Prometheus. Hence, if Loki is disabled, some features of the plugin are disabled as well, such as getting per-pod information or viewing raw flows. If both Prometheus and Loki are enabled, Prometheus takes precedence and Loki is used as a fallback for queries that Prometheus cannot handle. If they are both disabled, the Console plugin is not deployed.
<b>manual</b>	<b>object</b>	Prometheus configuration for <b>Manual</b> mode.
<b>mode</b>	<b>string</b>	<b>mode</b> must be set according to the type of Prometheus installation that stores Network Observability metrics: <ul style="list-style-type: none"> <li>- Use <b>Auto</b> to try configuring automatically. In OpenShift Container Platform, it uses the Thanos querier from OpenShift Container Platform Cluster Monitoring</li> <li>- Use <b>Manual</b> for a manual setup</li> </ul>
<b>timeout</b>	<b>string</b>	<b>timeout</b> is the read timeout for console plugin queries to Prometheus. A timeout of zero means no timeout.

### 12.1.78. .spec.prometheus.querier.manual

#### Description

Prometheus configuration for **Manual** mode.

Type  
object

Property	Type	Description
<b>forwardUserToken</b>	<b>boolean</b>	Set <b>true</b> to forward logged in user token in queries to Prometheus
<b>tls</b>	<b>object</b>	TLS client configuration for Prometheus URL.
<b>url</b>	<b>string</b>	<b>url</b> is the address of an existing Prometheus service to use for querying metrics.

### 12.1.79. .spec.prometheus.querier.manual.tls

Description

TLS client configuration for Prometheus URL.

Type  
object

Property	Type	Description
<b>caCert</b>	<b>object</b>	<b>caCert</b> defines the reference of the certificate for the Certificate Authority
<b>enable</b>	<b>boolean</b>	Enable TLS
<b>insecureSkipVerify</b>	<b>boolean</b>	<b>insecureSkipVerify</b> allows skipping client-side verification of the server certificate. If set to <b>true</b> , the <b>caCert</b> field is ignored.
<b>userCert</b>	<b>object</b>	<b>userCert</b> defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

### 12.1.80. .spec.prometheus.querier.manual.tls.caCert

Description

**caCert** defines the reference of the certificate for the Certificate Authority

Type  
object

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

### 12.1.81. .spec.prometheus.querier.manual.tls.userCert

#### Description

**userCert** defines the user certificate reference and is used for mTLS (you can ignore it when using one-way TLS)

#### Type

**object**

Property	Type	Description
<b>certFile</b>	<b>string</b>	<b>certFile</b> defines the path to the certificate file name within the config map or secret.
<b>certKey</b>	<b>string</b>	<b>certKey</b> defines the path to the certificate private key file name within the config map or secret. Omit when the key is not necessary.



Property	Type	Description
<b>name</b>	<b>string</b>	Name of the config map or secret containing certificates.
<b>namespace</b>	<b>string</b>	Namespace of the config map or secret containing certificates. If omitted, the default is to use the same namespace as where Network Observability is deployed. If the namespace is different, the config map or the secret is copied so that it can be mounted as required.
<b>type</b>	<b>string</b>	Type for the certificate reference: <b>configmap</b> or <b>secret</b> .

## CHAPTER 13. FLOWMETRIC CONFIGURATION PARAMETERS

**FlowMetric** is the API allowing to create custom metrics from the collected flow logs.

### 13.1. FLOWMETRIC [FLOWS.NETOBSERV.IO/V1ALPHA1]

#### Description

FlowMetric is the API allowing to create custom metrics from the collected flow logs.

#### Type

**object**

Property	Type	Description
<b>apiVersion</b>	<b>string</b>	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and might reject unrecognized values. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources</a>
<b>kind</b>	<b>string</b>	Kind is a string value representing the REST resource this object represents. Servers might infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds</a>
<b>metadata</b>	<b>object</b>	Standard object's metadata. More info: <a href="https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata">https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</a>

Property	Type	Description
<b>spec</b>	<b>object</b>	<p>FlowMetricSpec defines the desired state of FlowMetric The provided API allows you to customize these metrics according to your needs.</p> <p>When adding new metrics or modifying existing labels, you must carefully monitor the memory usage of Prometheus workloads as this could potentially have a high impact. Cf <a href="https://rhobs-handbook.netlify.app/products/openshiftmonitoring/telemetry.md/#what-is-the-cardinality-of-a-metric">https://rhobs-handbook.netlify.app/products/openshiftmonitoring/telemetry.md/#what-is-the-cardinality-of-a-metric</a></p> <p>To check the cardinality of all Network Observability metrics, run as <b>promql</b>:  <b>count({name=~"netobserv.*"}) by (name).</b></p>

### 13.1.1. .metadata

#### Description

Standard object's metadata. More info: <https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata>

#### Type

**object**

### 13.1.2. .spec

#### Description

FlowMetricSpec defines the desired state of FlowMetric The provided API allows you to customize these metrics according to your needs.

When adding new metrics or modifying existing labels, you must carefully monitor the memory usage of Prometheus workloads as this could potentially have a high impact. Cf <https://rhobs-handbook.netlify.app/products/openshiftmonitoring/telemetry.md/#what-is-the-cardinality-of-a-metric>

To check the cardinality of all Network Observability metrics, run as **promql**:  
**count({name=~"netobserv.\*"}) by (name).**

#### Type

**object**

#### Required

- **metricName**

- **type**

Property	Type	Description
<b>buckets</b>	<b>array (string)</b>	A list of buckets to use when <b>type</b> is "Histogram". The list must be parsable as floats. When not set, Prometheus default buckets are used.
<b>charts</b>	<b>array</b>	Charts configuration, for the OpenShift Container Platform Console in the administrator view, Dashboards menu.
<b>direction</b>	<b>string</b>	Filter for ingress, egress or any direction flows. When set to <b>Ingress</b> , it is equivalent to adding the regular expression filter on <b>FlowDirection: 0 2</b> . When set to <b>Egress</b> , it is equivalent to adding the regular expression filter on <b>FlowDirection: 1 2</b> .
<b>divider</b>	<b>string</b>	When nonzero, scale factor (divider) of the value. Metric value = Flow value / Divider.
<b>filters</b>	<b>array</b>	<b>filters</b> is a list of fields and values used to restrict which flows are taken into account. Oftentimes, these filters must be used to eliminate duplicates: <b>Duplicate != "true"</b> and <b>FlowDirection = "0"</b> . Refer to the documentation for the list of available fields: <a href="https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html">https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html</a> .

Property	Type	Description
<b>labels</b>	<b>array (string)</b>	<p><b>labels</b> is a list of fields that should be used as Prometheus labels, also known as dimensions. From choosing labels results the level of granularity of this metric, and the available aggregations at query time. It must be done carefully as it impacts the metric cardinality (cf <a href="https://rhobs-handbook.netlify.app/products/openshiftmonitoring/telemetry.md/#what-is-the-cardinality-of-a-metric">https://rhobs-handbook.netlify.app/products/openshiftmonitoring/telemetry.md/#what-is-the-cardinality-of-a-metric</a>). In general, avoid setting very high cardinality labels such as IP or MAC addresses. "SrcK8S_OwnerName" or "DstK8S_OwnerName" should be preferred over "SrcK8S_Name" or "DstK8S_Name" as much as possible. Refer to the documentation for the list of available fields: <a href="https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html">https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html</a>.</p>
<b>metricName</b>	<b>string</b>	Name of the metric. In Prometheus, it is automatically prefixed with "netobserv_".
<b>type</b>	<b>string</b>	Metric type: "Counter" or "Histogram". Use "Counter" for any value that increases over time and on which you can compute a rate, such as Bytes or Packets. Use "Histogram" for any value that must be sampled independently, such as latencies.

Property	Type	Description
<b>valueField</b>	<b>string</b>	<b>valueField</b> is the flow field that must be used as a value for this metric. This field must hold numeric values. Leave empty to count flows rather than a specific value per flow. Refer to the documentation for the list of available fields: <a href="https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html">https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html</a> .

### 13.1.3. .spec.charts

#### Description

Charts configuration, for the OpenShift Container Platform Console in the administrator view, Dashboards menu.

#### Type

**array**

### 13.1.4. .spec.charts[]

#### Description

Configures charts / dashboard generation associated to a metric

#### Type

**object**

#### Required

- **dashboardName**
- **queries**
- **title**
- **type**

Property	Type	Description
<b>dashboardName</b>	<b>string</b>	Name of the containing dashboard. If this name does not refer to an existing dashboard, a new dashboard is created.

Property	Type	Description
<b>queries</b>	<b>array</b>	List of queries to be displayed on this chart. If <b>type</b> is <b>SingleStat</b> and multiple queries are provided, this chart is automatically expanded in several panels (one per query).
<b>sectionName</b>	<b>string</b>	Name of the containing dashboard section. If this name does not refer to an existing section, a new section is created. If <b>sectionName</b> is omitted or empty, the chart is placed in the global top section.
<b>title</b>	<b>string</b>	Title of the chart.
<b>type</b>	<b>string</b>	Type of the chart.
<b>unit</b>	<b>string</b>	Unit of this chart. Only a few units are currently supported. Leave empty to use generic number.

### 13.1.5. .spec.charts[].queries

#### Description

List of queries to be displayed on this chart. If **type** is **SingleStat** and multiple queries are provided, this chart is automatically expanded in several panels (one per query).

#### Type

**array**

### 13.1.6. .spec.charts[].queries[]

#### Description

Configures PromQL queries

#### Type

**object**

#### Required

- **legend**
- **promQL**
- **top**

Property	Type	Description
<b>legend</b>	<b>string</b>	The query legend that applies to each timeseries represented in this chart. When multiple timeseries are displayed, you should set a legend that distinguishes each of them. It can be done with the following format: <b>{{ Label }}</b> . For example, if the <b>promQL</b> groups timeseries per label such as: <b>sum(rate(\$METRIC[2m])) by (Label1, Label2)</b> , you might write as the legend: <b>Label1={{ Label1 }}</b> , <b>Label2={{ Label2 }}</b> .
<b>promQL</b>	<b>string</b>	The <b>promQL</b> query to be run against Prometheus. If the chart <b>type</b> is <b>SingleStat</b> , this query should only return a single timeseries. For other types, a top 7 is displayed. You can use <b>\$METRIC</b> to refer to the metric defined in this resource. For example: <b>sum(rate(\$METRIC[2m]))</b> . To learn more about <b>promQL</b> , refer to the Prometheus documentation: <a href="https://prometheus.io/docs/prometheus/latest/querying/basics/">https://prometheus.io/docs/prometheus/latest/querying/basics/</a>
<b>top</b>	<b>integer</b>	Top N series to display per timestamp. Does not apply to <b>SingleStat</b> chart type.

### 13.1.7. .spec.filters

#### Description

**filters** is a list of fields and values used to restrict which flows are taken into account. Oftentimes, these filters must be used to eliminate duplicates: **Duplicate != "true"** and **FlowDirection = "0"**. Refer to the documentation for the list of available fields: [https://docs.openshift.com/container-platform/latest/observability/network\\_observability/json-flows-format-reference.html](https://docs.openshift.com/container-platform/latest/observability/network_observability/json-flows-format-reference.html).

#### Type

**array**

### 13.1.8. .spec.filters[]

#### Description



Type

**object**

Required

- **field**
- **matchType**

Property	Type	Description
<b>field</b>	<b>string</b>	Name of the field to filter on
<b>matchType</b>	<b>string</b>	Type of matching to apply
<b>value</b>	<b>string</b>	Value to filter on. When <b>matchType</b> is <b>Equal</b> or <b>NotEqual</b> , you can use field injection with <b>\$(SomeField)</b> to refer to any other field of the flow.

## CHAPTER 14. NETWORK FLOWS FORMAT REFERENCE

These are the specifications for network flows format, used both internally and when exporting flows to Kafka.

### 14.1. NETWORK FLOWS FORMAT REFERENCE

This is the specification of the network flows format. That format is used when a Kafka exporter is configured, for Prometheus metrics labels as well as internally for the Loki store.

The "Filter ID" column shows which related name to use when defining Quick Filters (see [spec.consolePlugin.quickFilters](#) in the **FlowCollector** specification).

The "Loki label" column is useful when querying Loki directly: label fields need to be selected using [stream selectors](#).

The "Cardinality" column gives information about the implied metric cardinality if this field was to be used as a Prometheus label with the **FlowMetric** API. For more information, see the "FlowMetric API reference".

Name	Type	Description	Filter ID	Loki label	Cardinality
<b>Bytes</b>	number	Number of bytes	n/a	no	avoid
<b>DnsErrno</b>	number	Error number returned from DNS tracker ebpif hook function	<b>dns_errno</b>	no	fine
<b>DnsFlags</b>	number	DNS flags for DNS record	n/a	no	fine
<b>DnsFlagsResponseCode</b>	string	Parsed DNS header RCODEs name	<b>dns_flags_response_code</b>	no	fine
<b>DnsId</b>	number	DNS record id	<b>dns_id</b>	no	avoid
<b>DnsLatencyMs</b>	number	Time between a DNS request and response, in milliseconds	<b>dns_latency</b>	no	avoid
<b>Dscp</b>	number	Differentiated Services Code Point (DSCP) value	<b>dscp</b>	no	fine
<b>DstAddr</b>	string	Destination IP address (ipv4 or ipv6)	<b>dst_address</b>	no	avoid
<b>DstK8S_HostIP</b>	string	Destination node IP	<b>dst_host_address</b>	no	fine

Name	Type	Description	Filter ID	Loki label	Cardinality
<b>DstK8S_HostName</b>	string	Destination node name	<b>dst_host_name</b>	no	fine
<b>DstK8S_Name</b>	string	Name of the destination Kubernetes object, such as Pod name, Service name or Node name.	<b>dst_name</b>	no	careful
<b>DstK8S_Namespace</b>	string	Destination namespace	<b>dst_namespace</b>	yes	fine
<b>DstK8S_OwnerName</b>	string	Name of the destination owner, such as Deployment name, StatefulSet name, etc.	<b>dst_owner_name</b>	yes	fine
<b>DstK8S_OwnerType</b>	string	Kind of the destination owner, such as Deployment, StatefulSet, etc.	<b>dst_kind</b>	no	fine
<b>DstK8S_Type</b>	string	Kind of the destination Kubernetes object, such as Pod, Service or Node.	<b>dst_kind</b>	yes	fine
<b>DstK8S_Zone</b>	string	Destination availability zone	<b>dst_zone</b>	yes	fine
<b>DstMac</b>	string	Destination MAC address	<b>dst_mac</b>	no	avoid
<b>DstPort</b>	number	Destination port	<b>dst_port</b>	no	careful
<b>DstSubnetLabel</b>	string	Destination subnet label	<b>dst_subnet_label</b>	no	fine
<b>Duplicate</b>	boolean	Indicates if this flow was also captured from another interface on the same host	n/a	yes	fine
<b>Flags</b>	number	Logical OR combination of unique TCP flags comprised in the flow, as per RFC-9293, with additional custom flags to represent the following per-packet combinations: - SYN+ACK (0x100) - FIN+ACK (0x200) - RST+ACK (0x400)	n/a	no	fine

Name	Type	Description	Filter ID	Loki label	Cardinality
<b>FlowDirection</b>	number	Flow interpreted direction from the node observation point. Can be one of: - 0: Ingress (incoming traffic, from the node observation point) - 1: Egress (outgoing traffic, from the node observation point) - 2: Inner (with the same source and destination node)	<b>node_direction</b>	yes	fine
<b>IcmpCode</b>	number	ICMP code	<b>icmp_code</b>	no	fine
<b>IcmpType</b>	number	ICMP type	<b>icmp_type</b>	no	fine
<b>IfDirections</b>	number	Flow directions from the network interface observation point. Can be one of: - 0: Ingress (interface incoming traffic) - 1: Egress (interface outgoing traffic)	<b>ifdirections</b>	no	fine
<b>Interfaces</b>	string	Network interfaces	<b>interfaces</b>	no	careful
<b>K8S_ClusterName</b>	string	Cluster name or identifier	<b>cluster_name</b>	yes	fine
<b>K8S_FlowLayer</b>	string	Flow layer: 'app' or 'infra'	<b>flow_layer</b>	no	fine
<b>Packets</b>	number	Number of packets	n/a	no	avoid
<b>PktDropBytes</b>	number	Number of bytes dropped by the kernel	n/a	no	avoid
<b>PktDropLatestDropCause</b>	string	Latest drop cause	<b>pkt_drop_cause</b>	no	fine
<b>PktDropLatestFlags</b>	number	TCP flags on last dropped packet	n/a	no	fine

Name	Type	Description	Filter ID	Loki label	Cardinality
<b>PktDrop LatestState</b>	string	TCP state on last dropped packet	<b>pkt_drop_state</b>	no	fine
<b>PktDrop Packets</b>	number	Number of packets dropped by the kernel	n/a	no	avoid
<b>Proto</b>	number	L4 protocol	<b>protocol</b>	no	fine
<b>SrcAddr</b>	string	Source IP address (ipv4 or ipv6)	<b>src_address</b>	no	avoid
<b>SrcK8S_HostIP</b>	string	Source node IP	<b>src_host_address</b>	no	fine
<b>SrcK8S_HostName</b>	string	Source node name	<b>src_host_name</b>	no	fine
<b>SrcK8S_Name</b>	string	Name of the source Kubernetes object, such as Pod name, Service name or Node name.	<b>src_name</b>	no	careful
<b>SrcK8S_Namespace</b>	string	Source namespace	<b>src_namespace</b>	yes	fine
<b>SrcK8S_OwnerName</b>	string	Name of the source owner, such as Deployment name, StatefulSet name, etc.	<b>src_owner_name</b>	yes	fine
<b>SrcK8S_OwnerType</b>	string	Kind of the source owner, such as Deployment, StatefulSet, etc.	<b>src_kind</b>	no	fine
<b>SrcK8S_Type</b>	string	Kind of the source Kubernetes object, such as Pod, Service or Node.	<b>src_kind</b>	yes	fine
<b>SrcK8S_Zone</b>	string	Source availability zone	<b>src_zone</b>	yes	fine
<b>SrcMac</b>	string	Source MAC address	<b>src_mac</b>	no	avoid
<b>SrcPort</b>	number	Source port	<b>src_port</b>	no	careful
<b>SrcSubnetLabel</b>	string	Source subnet label	<b>src_subnet_label</b>	no	fine

Name	Type	Description	Filter ID	Loki label	Cardinality
<b>TimeFlowEndMs</b>	number	End timestamp of this flow, in milliseconds	n/a	no	avoid
<b>TimeFlowRttNs</b>	number	TCP Smoothed Round Trip Time (SRTT), in nanoseconds	<b>time_flow_rtt</b>	no	avoid
<b>TimeFlowStartMs</b>	number	Start timestamp of this flow, in milliseconds	n/a	no	avoid
<b>TimeReceived</b>	number	Timestamp when this flow was received and processed by the flow collector, in seconds	n/a	no	avoid
<b>_HashId</b>	string	In conversation tracking, the conversation identifier	<b>id</b>	no	avoid
<b>_RecordType</b>	string	Type of record: 'flowLog' for regular flow logs, or 'newConnection', 'heartbeat', 'endConnection' for conversation tracking	<b>type</b>	yes	fine

## CHAPTER 15. TROUBLESHOOTING NETWORK OBSERVABILITY

To assist in troubleshooting Network Observability issues, you can perform some troubleshooting actions.

### 15.1. USING THE MUST-GATHER TOOL

You can use the must-gather tool to collect information about the Network Observability Operator resources and cluster-wide resources, such as pod logs, **FlowCollector**, and **webhook** configurations.

#### Procedure

1. Navigate to the directory where you want to store the must-gather data.
2. Run the following command to collect cluster-wide must-gather resources:

```
$ oc adm must-gather
--image-stream=openshift/must-gather \
--image=quay.io/netobserv/must-gather
```

### 15.2. CONFIGURING NETWORK TRAFFIC MENU ENTRY IN THE OPENSIFT CONTAINER PLATFORM CONSOLE

Manually configure the network traffic menu entry in the OpenShift Container Platform console when the network traffic menu entry is not listed in **Observe** menu in the OpenShift Container Platform console.

#### Prerequisites

- You have installed OpenShift Container Platform version 4.10 or newer.

#### Procedure

1. Check if the **spec.consolePlugin.register** field is set to **true** by running the following command:

```
$ oc -n netobserv get flowcollector cluster -o yaml
```

#### Example output

```
apiVersion: flows.netobserv.io/v1alpha1
kind: FlowCollector
metadata:
  name: cluster
spec:
  consolePlugin:
    register: false
```

2. Optional: Add the **netobserv-plugin** plugin by manually editing the Console Operator config:

```
$ oc edit console.operator.openshift.io cluster
```

### Example output

```
...  
spec:  
  plugins:  
  - netobserv-plugin  
...
```

- Optional: Set the **spec.consolePlugin.register** field to **true** by running the following command:

```
$ oc -n netobserv edit flowcollector cluster -o yaml
```

### Example output

```
apiVersion: flows.netobserv.io/v1alpha1  
kind: FlowCollector  
metadata:  
  name: cluster  
spec:  
  consolePlugin:  
    register: true
```

- Ensure the status of console pods is **running** by running the following command:

```
$ oc get pods -n openshift-console -l app=console
```

- Restart the console pods by running the following command:

```
$ oc delete pods -n openshift-console -l app=console
```

- Clear your browser cache and history.

- Check the status of Network Observability plugin pods by running the following command:

```
$ oc get pods -n netobserv -l app=netobserv-plugin
```

### Example output

```
NAME                                READY STATUS RESTARTS AGE  
netobserv-plugin-68c7bbb9bb-b69q6  1/1   Running 0      21s
```

- Check the logs of the Network Observability plugin pods by running the following command:

```
$ oc logs -n netobserv -l app=netobserv-plugin
```

### Example output



```
time="2022-12-13T12:06:49Z" level=info msg="Starting netobserv-console-plugin [build
version: , build date: 2022-10-21 15:15] at log level info" module=main
time="2022-12-13T12:06:49Z" level=info msg="listening on https://:9001" module=server
```

### 15.3. FLOWLOGS-PIPELINE DOES NOT CONSUME NETWORK FLOWS AFTER INSTALLING KAFKA

If you deployed the flow collector first with **deploymentModel: KAFKA** and then deployed Kafka, the flow collector might not connect correctly to Kafka. Manually restart the flow-pipeline pods where Flowlogs-pipeline does not consume network flows from Kafka.

#### Procedure

1. Delete the flow-pipeline pods to restart them by running the following command:

```
$ oc delete pods -n netobserv -l app=flowlogs-pipeline-transformer
```

### 15.4. FAILING TO SEE NETWORK FLOWS FROM BOTH BR-INT AND BR-EX INTERFACES

`br-ex`` and **br-int** are virtual bridge devices operated at OSI layer 2. The eBPF agent works at the IP and TCP levels, layers 3 and 4 respectively. You can expect that the eBPF agent captures the network traffic passing through **br-ex** and **br-int**, when the network traffic is processed by other interfaces such as physical host or virtual pod interfaces. If you restrict the eBPF agent network interfaces to attach only to **br-ex** and **br-int**, you do not see any network flow.

Manually remove the part in the **interfaces** or **excludeInterfaces** that restricts the network interfaces to **br-int** and **br-ex**.

#### Procedure

1. Remove the **interfaces: [ 'br-int', 'br-ex' ]** field. This allows the agent to fetch information from all the interfaces. Alternatively, you can specify the Layer-3 interface for example, **eth0**. Run the following command:

```
$ oc edit -n netobserv flowcollector.yaml -o yaml
```

#### Example output

```
apiVersion: flows.netobserv.io/v1alpha1
kind: FlowCollector
metadata:
  name: cluster
spec:
  agent:
    type: EBPF
    ebpf:
      interfaces: [ 'br-int', 'br-ex' ] 1
```

- 1 Specifies the network interfaces.

## 15.5. NETWORK OBSERVABILITY CONTROLLER MANAGER POD RUNS OUT OF MEMORY

You can increase memory limits for the Network Observability operator by editing the `spec.config.resources.limits.memory` specification in the **Subscription** object.

### Procedure

1. In the web console, navigate to **Operators** → **Installed Operators**
2. Click **Network Observability** and then select **Subscription**.
3. From the **Actions** menu, click **Edit Subscription**.
  - a. Alternatively, you can use the CLI to open the YAML configuration for the **Subscription** object by running the following command:

```
$ oc edit subscription netobserv-operator -n openshift-netobserv-operator
```

4. Edit the **Subscription** object to add the `config.resources.limits.memory` specification and set the value to account for your memory requirements. See the Additional resources for more information about resource considerations:

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: netobserv-operator
  namespace: openshift-netobserv-operator
spec:
  channel: stable
  config:
    resources:
      limits:
        memory: 800Mi 1
      requests:
        cpu: 100m
        memory: 100Mi
  installPlanApproval: Automatic
  name: netobserv-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace
  startingCSV: <network_observability_operator_latest_version> 2
```

- 1** For example, you can increase the memory limit to **800Mi**.
- 2** This value should not be edited, but note that it changes depending on the most current release of the Operator.

## 15.6. RUNNING CUSTOM QUERIES TO LOKI

For troubleshooting, can run custom queries to Loki. There are two examples of ways to do this, which you can adapt according to your needs by replacing the `<api_token>` with your own.



## NOTE

These examples use the **netobserv** namespace for the Network Observability Operator and Loki deployments. Additionally, the examples assume that the LokiStack is named **loki**. You can optionally use a different namespace and naming by adapting the examples, specifically the **-n netobserv** or the **loki-gateway** URL.

## Prerequisites

- Installed Loki Operator for use with Network Observability Operator

## Procedure

- To get all available labels, run the following:

```
$ oc exec deployment/netobserv-plugin -n netobserv -- curl -G -s -H 'X-Scope-
OrgID:network' -H 'Authorization: Bearer <api_token>' -k https://loki-gateway-
http.netobserv.svc:8080/api/logs/v1/network/loki/api/v1/labels | jq
```

- To get all flows from the source namespace, **my-namespace**, run the following:

```
$ oc exec deployment/netobserv-plugin -n netobserv -- curl -G -s -H 'X-Scope-
OrgID:network' -H 'Authorization: Bearer <api_token>' -k https://loki-gateway-
http.netobserv.svc:8080/api/logs/v1/network/loki/api/v1/query --data-urlencode 'query=
{SrcK8S_Namespace="my-namespace"}' | jq
```

## Additional resources

- [Resource considerations](#)

## 15.7. TROUBLESHOOTING LOKI RESOURCEEXHAUSTED ERROR

Loki may return a **ResourceExhausted** error when network flow data sent by Network Observability exceeds the configured maximum message size. If you are using the Red Hat Loki Operator, this maximum message size is configured to 100 MiB.

## Procedure

1. Navigate to **Operators** → **Installed Operators**, viewing **All projects** from the **Project** dropdown menu.
2. In the **Provided APIs** list, select the Network Observability Operator.
3. Click the **Flow Collector** then the **YAML view** tab.
  - a. If you are using the Loki Operator, check that the **spec.loki.batchSize** value does not exceed 98 MiB.
  - b. If you are using a Loki installation method that is different from the Red Hat Loki Operator, such as Grafana Loki, verify that the **grpc\_server\_max\_recv\_msg\_size** [Grafana Loki server setting](#) is higher than the **FlowCollector** resource **spec.loki.batchSize** value. If it is not, you must either increase the **grpc\_server\_max\_recv\_msg\_size** value, or decrease the **spec.loki.batchSize** value so that it is lower than the limit.

4. Click **Save** if you edited the **FlowCollector**.

## 15.8. LOKI EMPTY RING ERROR

The Loki "empty ring" error results in flows not being stored in Loki and not showing up in the web console. This error might happen in various situations. A single workaround to address them all does not exist. There are some actions you can take to investigate the logs in your Loki pods, and verify that the **LokiStack** is healthy and ready.

Some of the situations where this error is observed are as follows:

- After a **LokiStack** is uninstalled and reinstalled in the same namespace, old PVCs are not removed, which can cause this error.
  - **Action:** You can try removing the **LokiStack** again, removing the PVC, then reinstalling the **LokiStack**.
- After a certificate rotation, this error can prevent communication with the **flowlogs-pipeline** and **console-plugin** pods.
  - **Action:** You can restart the pods to restore the connectivity.

## 15.9. RESOURCE TROUBLESHOOTING

### 15.10. LOKISTACK RATE LIMIT ERRORS

A rate-limit placed on the Loki tenant can result in potential temporary loss of data and a 429 error: **Per stream rate limit exceeded (limit:xMB/sec) while attempting to ingest for stream**. You might consider having an alert set to notify you of this error. For more information, see "Creating Loki rate limit alerts for the NetObserv dashboard" in the Additional resources of this section.

You can update the LokiStack CRD with the **perStreamRateLimit** and **perStreamRateLimitBurst** specifications, as shown in the following procedure.

#### Procedure

1. Navigate to **Operators** → **Installed Operators**, viewing **All projects** from the **Project** dropdown.
2. Look for **Loki Operator**, and select the **LokiStack** tab.
3. Create or edit an existing **LokiStack** instance using the **YAML view** to add the **perStreamRateLimit** and **perStreamRateLimitBurst** specifications:

```
apiVersion: loki.grafana.com/v1
kind: LokiStack
metadata:
  name: loki
  namespace: netobserv
spec:
  limits:
    global:
      ingestion:
        perStreamRateLimit: 6
```

1

```
perStreamRateLimitBurst: 30 2  
tenants:  
  mode: openshift-network  
  managementState: Managed
```

- 1** The default value for **perStreamRateLimit** is **3**.
- 2** The default value for **perStreamRateLimitBurst** is **15**.

4. Click **Save**.

## Verification

Once you update the **perStreamRateLimit** and **perStreamRateLimitBurst** specifications, the pods in your cluster restart and the 429 rate-limit error no longer occurs.