



OpenShift Container Platform 4.12

Red Hat build of OpenTelemetry

Configuring and using the Red Hat build of OpenTelemetry in OpenShift Container Platform

OpenShift Container Platform 4.12 Red Hat build of OpenTelemetry

Configuring and using the Red Hat build of OpenTelemetry 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[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] 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[®] 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 Red Hat build of the open source OpenTelemetry project to collect unified, standardized, and vendor-neutral telemetry data for cloud-native software in OpenShift Container Platform.

Table of Contents

CHAPTER 1. RELEASE NOTES	5
1.1. RELEASE NOTES FOR RED HAT BUILD OF OPENTELEMETRY 3.2	5
1.1.1. Red Hat build of OpenTelemetry overview	5
1.1.2. Technology Preview features	5
1.1.3. New features and enhancements	6
1.1.4. Deprecated functionality	6
1.1.5. Bug fixes	6
1.1.6. Getting support	7
1.1.7. Making open source more inclusive	7
1.2. RELEASE NOTES FOR PAST RELEASES OF RED HAT BUILD OF OPENTELEMETRY	7
1.2.1. Red Hat build of OpenTelemetry overview	7
1.2.2. Release notes for Red Hat build of OpenTelemetry 3.1.1	8
1.2.2.1. CVEs	8
1.2.3. Release notes for Red Hat build of OpenTelemetry 3.1	8
1.2.3.1. Technology Preview features	8
1.2.3.2. New features and enhancements	8
1.2.4. Release notes for Red Hat build of OpenTelemetry 3.0	8
1.2.4.1. New features and enhancements	9
1.2.4.2. Removal notice	9
1.2.4.3. Bug fixes	9
1.2.4.4. Known issues	9
1.2.5. Release notes for Red Hat build of OpenTelemetry 2.9.2	10
1.2.5.1. CVEs	11
1.2.5.2. Known issues	11
1.2.6. Release notes for Red Hat build of OpenTelemetry 2.9.1	11
1.2.6.1. CVEs	11
1.2.6.2. Known issues	11
1.2.7. Release notes for Red Hat build of OpenTelemetry 2.9	11
1.2.7.1. New features and enhancements	12
1.2.7.2. Known issues	12
1.2.8. Release notes for Red Hat build of OpenTelemetry 2.8	12
1.2.8.1. Bug fixes	13
1.2.9. Release notes for Red Hat build of OpenTelemetry 2.7	13
1.2.9.1. Bug fixes	13
1.2.10. Release notes for Red Hat build of OpenTelemetry 2.6	13
1.2.10.1. Bug fixes	13
1.2.11. Release notes for Red Hat build of OpenTelemetry 2.5	13
1.2.11.1. New features and enhancements	14
1.2.11.2. Bug fixes	14
1.2.12. Release notes for Red Hat build of OpenTelemetry 2.4	14
1.2.12.1. Bug fixes	14
1.2.13. Release notes for Red Hat build of OpenTelemetry 2.3	14
1.2.13.1. Bug fixes	15
1.2.14. Release notes for Red Hat build of OpenTelemetry 2.2	15
1.2.14.1. Technology Preview features	15
1.2.14.2. Bug fixes	15
1.2.15. Release notes for Red Hat build of OpenTelemetry 2.1	15
1.2.15.1. Technology Preview features	16
1.2.15.2. Bug fixes	16
1.2.16. Release notes for Red Hat build of OpenTelemetry 2.0	16
1.2.17. Getting support	17

1.2.18. Making open source more inclusive	17
CHAPTER 2. INSTALLING	18
2.1. INSTALLING THE RED HAT BUILD OF OPENTELEMETRY FROM THE WEB CONSOLE	18
2.2. INSTALLING THE RED HAT BUILD OF OPENTELEMETRY BY USING THE CLI	19
2.3. ADDITIONAL RESOURCES	22
CHAPTER 3. CONFIGURING THE COLLECTOR	23
3.1. OPENTELEMETRY COLLECTOR CONFIGURATION OPTIONS	23
3.2. OPENTELEMETRY COLLECTOR COMPONENTS	26
3.2.1. Receivers	26
3.2.1.1. OTLP Receiver	26
3.2.1.2. Jaeger Receiver	27
3.2.1.3. Host Metrics Receiver	28
3.2.1.4. Kubernetes Objects Receiver	30
3.2.1.5. Kubelet Stats Receiver	32
3.2.1.6. Prometheus Receiver	33
3.2.1.7. Zipkin Receiver	33
3.2.1.8. Kafka Receiver	34
3.2.1.9. Kubernetes Cluster Receiver	35
3.2.1.10. OpenCensus Receiver	37
3.2.1.11. Filelog Receiver	38
3.2.1.12. Journald Receiver	38
3.2.1.13. Kubernetes Events Receiver	41
3.2.2. Processors	43
3.2.2.1. Batch Processor	43
3.2.2.2. Memory Limiter Processor	44
3.2.2.3. Resource Detection Processor	45
3.2.2.4. Attributes Processor	46
3.2.2.5. Resource Processor	47
3.2.2.6. Span Processor	47
3.2.2.7. Kubernetes Attributes Processor	48
3.2.2.8. Filter Processor	49
3.2.2.9. Routing Processor	50
3.2.2.10. Cumulative to Delta Processor	50
3.2.3. Exporters	51
3.2.3.1. OTLP Exporter	51
3.2.3.2. OTLP HTTP Exporter	52
3.2.3.3. Debug Exporter	53
3.2.3.4. Load Balancing Exporter	53
3.2.3.5. Prometheus Exporter	54
3.2.3.6. Kafka Exporter	55
3.2.4. Connectors	56
3.2.4.1. Forward Connector	56
3.2.4.2. Spanmetrics Connector	57
3.2.5. Extensions	58
3.2.5.1. BearerTokenAuth Extension	58
3.2.5.2. OAuth2Client Extension	58
3.2.5.3. File Storage Extension	60
3.2.5.4. OIDC Auth Extension	61
3.2.5.5. Jaeger Remote Sampling Extension	62
3.2.5.6. Performance Profiler Extension	64
3.2.5.7. Health Check Extension	64

3.2.5.8. Memory Ballast Extension	65
3.2.5.9. zPages Extension	66
3.3. CREATING THE REQUIRED RBAC RESOURCES AUTOMATICALLY	67
3.4. TARGET ALLOCATOR	68
CHAPTER 4. CONFIGURING THE INSTRUMENTATION	71
4.1. OPENTELEMETRY INSTRUMENTATION CONFIGURATION OPTIONS	71
4.1.1. Instrumentation options	71
4.1.2. Using the instrumentation CR with Service Mesh	73
4.1.2.1. Configuration of the Apache HTTP Server auto-instrumentation	73
4.1.2.2. Configuration of the .NET auto-instrumentation	73
4.1.2.3. Configuration of the Go auto-instrumentation	74
4.1.2.4. Configuration of the Java auto-instrumentation	75
4.1.2.5. Configuration of the Node.js auto-instrumentation	75
4.1.2.6. Configuration of the Python auto-instrumentation	76
4.1.2.7. Configuration of the OpenTelemetry SDK variables	76
4.1.2.8. Multi-container pods	76
CHAPTER 5. SENDING TRACES AND METRICS TO THE OPENTELEMETRY COLLECTOR	78
5.1. SENDING TRACES AND METRICS TO THE OPENTELEMETRY COLLECTOR WITH SIDECAR INJECTION	78
5.2. SENDING TRACES AND METRICS TO THE OPENTELEMETRY COLLECTOR WITHOUT SIDECAR INJECTION	80
CHAPTER 6. CONFIGURING METRICS FOR THE MONITORING STACK	83
6.1. CONFIGURATION FOR SENDING METRICS TO THE MONITORING STACK	83
6.2. CONFIGURATION FOR RECEIVING METRICS FROM THE MONITORING STACK	84
6.3. ADDITIONAL RESOURCES	86
CHAPTER 7. FORWARDING TRACES TO A TEMPOSTACK INSTANCE	87
CHAPTER 8. CONFIGURING THE OPENTELEMETRY COLLECTOR METRICS	90
CHAPTER 9. GATHERING THE OBSERVABILITY DATA FROM MULTIPLE CLUSTERS	91
CHAPTER 10. TROUBLESHOOTING	96
10.1. GETTING THE OPENTELEMETRY COLLECTOR LOGS	96
10.2. EXPOSING THE METRICS	96
10.3. DEBUG EXPORTER	97
CHAPTER 11. MIGRATING	98
11.1. MIGRATING WITH SIDECARS	98
11.2. MIGRATING WITHOUT SIDECARS	100
CHAPTER 12. UPGRADING	103
12.1. ADDITIONAL RESOURCES	103
CHAPTER 13. REMOVING	104
13.1. REMOVING AN OPENTELEMETRY COLLECTOR INSTANCE BY USING THE WEB CONSOLE	104
13.2. REMOVING AN OPENTELEMETRY COLLECTOR INSTANCE BY USING THE CLI	104
13.3. ADDITIONAL RESOURCES	105

CHAPTER 1. RELEASE NOTES

1.1. RELEASE NOTES FOR RED HAT BUILD OF OPENTELEMETRY 3.2

1.1.1. Red Hat build of OpenTelemetry overview

Red Hat build of OpenTelemetry is based on the open source [OpenTelemetry project](#), which aims to provide unified, standardized, and vendor-neutral telemetry data collection for cloud-native software. Red Hat build of OpenTelemetry product provides support for deploying and managing the OpenTelemetry Collector and simplifying the workload instrumentation.

The [OpenTelemetry Collector](#) can receive, process, and forward telemetry data in multiple formats, making it the ideal component for telemetry processing and interoperability between telemetry systems. The Collector provides a unified solution for collecting and processing metrics, traces, and logs.

The OpenTelemetry Collector has a number of features including the following:

Data Collection and Processing Hub

It acts as a central component that gathers telemetry data like metrics and traces from various sources. This data can be created from instrumented applications and infrastructure.

Customizable telemetry data pipeline

The OpenTelemetry Collector is designed to be customizable. It supports various processors, exporters, and receivers.

Auto-instrumentation features

Automatic instrumentation simplifies the process of adding observability to applications. Developers don't need to manually instrument their code for basic telemetry data.

Here are some of the use cases for the OpenTelemetry Collector:

Centralized data collection

In a microservices architecture, the Collector can be deployed to aggregate data from multiple services.

Data enrichment and processing

Before forwarding data to analysis tools, the Collector can enrich, filter, and process this data.

Multi-backend receiving and exporting

The Collector can receive and send data to multiple monitoring and analysis platforms simultaneously.

The Red Hat build of OpenTelemetry is provided through the Red Hat build of OpenTelemetry Operator.

1.1.2. Technology Preview features

This update introduces the following Technology Preview features:

- Host Metrics Receiver
- OIDC Auth Extension
- Kubernetes Cluster Receiver

- Kubernetes Events Receiver
- Kubernetes Objects Receiver
- Load-Balancing Exporter
- Kubelet Stats Receiver
- Cumulative to Delta Processor
- Forward Connector
- Journald Receiver
- Filelog Receiver
- File Storage Extension



IMPORTANT

Each of these features 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](#).

1.1.3. New features and enhancements

This update introduces the following enhancement:

- Red Hat build of OpenTelemetry 3.2 is based on the open source [OpenTelemetry](#) release 0.99.0.

1.1.4. Deprecated functionality

In Red Hat build of OpenTelemetry 3.2, use of empty values and **null** keywords in the OpenTelemetry Collector custom resource is deprecated and planned to be unsupported in a future release. Red Hat will provide bug fixes and support for this syntax during the current release lifecycle, but this syntax will become unsupported. As an alternative to empty values and **null** keywords, you can update the OpenTelemetry Collector custom resource to contain empty JSON objects as open-closed braces **{}** instead.

1.1.5. Bug fixes

This update introduces the following bug fix:

- Before this update, the checkbox to enable Operator monitoring was not available in the web console when installing the Red Hat build of OpenTelemetry Operator. As a result, a **ServiceMonitor** resource was not created in the **openshift-opentelemetry-operator** namespace. With this update, the checkbox appears for the Red Hat build of OpenTelemetry Operator in the web console so that Operator monitoring can be enabled during installation. ([TRACING-3761](#))

1.1.6. Getting support

If you experience difficulty with a procedure described in this documentation, or with OpenShift Container Platform in general, visit the [Red Hat Customer Portal](#). From the Customer Portal, you can:

- Search or browse through the Red Hat Knowledgebase of articles and solutions relating to Red Hat products.
- Submit a support case to Red Hat Support.
- Access other product documentation.

To identify issues with your cluster, you can use Insights in [OpenShift Cluster Manager Hybrid Cloud Console](#). Insights provides details about issues and, if available, information on how to solve a problem.

If you have a suggestion for improving this documentation or have found an error, submit a [Jira issue](#) for the most relevant documentation component. Please provide specific details, such as the section name and OpenShift Container Platform version.

1.1.7. Making open source more inclusive

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

1.2. RELEASE NOTES FOR PAST RELEASES OF RED HAT BUILD OF OPENTELEMETRY

1.2.1. Red Hat build of OpenTelemetry overview

Red Hat build of OpenTelemetry is based on the open source [OpenTelemetry project](#), which aims to provide unified, standardized, and vendor-neutral telemetry data collection for cloud-native software. Red Hat build of OpenTelemetry product provides support for deploying and managing the OpenTelemetry Collector and simplifying the workload instrumentation.

The [OpenTelemetry Collector](#) can receive, process, and forward telemetry data in multiple formats, making it the ideal component for telemetry processing and interoperability between telemetry systems. The Collector provides a unified solution for collecting and processing metrics, traces, and logs.

The OpenTelemetry Collector has a number of features including the following:

Data Collection and Processing Hub

It acts as a central component that gathers telemetry data like metrics and traces from various sources. This data can be created from instrumented applications and infrastructure.

Customizable telemetry data pipeline

The OpenTelemetry Collector is designed to be customizable. It supports various processors, exporters, and receivers.

Auto-instrumentation features

Automatic instrumentation simplifies the process of adding observability to applications. Developers don't need to manually instrument their code for basic telemetry data.

Here are some of the use cases for the OpenTelemetry Collector:

Centralized data collection

In a microservices architecture, the Collector can be deployed to aggregate data from multiple services.

Data enrichment and processing

Before forwarding data to analysis tools, the Collector can enrich, filter, and process this data.

Multi-backend receiving and exporting

The Collector can receive and send data to multiple monitoring and analysis platforms simultaneously.

1.2.2. Release notes for Red Hat build of OpenTelemetry 3.1.1

The Red Hat build of OpenTelemetry is provided through the Red Hat build of OpenTelemetry Operator.

1.2.2.1. CVEs

This release fixes [CVE-2023-39326](#).

1.2.3. Release notes for Red Hat build of OpenTelemetry 3.1

The Red Hat build of OpenTelemetry is provided through the Red Hat build of OpenTelemetry Operator.

1.2.3.1. Technology Preview features

This update introduces the following Technology Preview feature:

- The target allocator is an optional component of the OpenTelemetry Operator that shards Prometheus receiver scrape targets across the deployed fleet of OpenTelemetry Collector instances. The target allocator provides integration with the Prometheus **PodMonitor** and **ServiceMonitor** custom resources.



IMPORTANT

The target allocator 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](#).

1.2.3.2. New features and enhancements

This update introduces the following enhancement:

- Red Hat build of OpenTelemetry 3.1 is based on the open source [OpenTelemetry](#) release 0.93.0.

1.2.4. Release notes for Red Hat build of OpenTelemetry 3.0

1.2.4.1. New features and enhancements

This update introduces the following enhancements:

- Red Hat build of OpenTelemetry 3.0 is based on the open source [OpenTelemetry](#) release 0.89.0.
- The **OpenShift distributed tracing data collection Operator** is renamed as the **Red Hat build of OpenTelemetry Operator**.
- Support for the ARM architecture.
- Support for the Prometheus receiver for metrics collection.
- Support for the Kafka receiver and exporter for sending traces and metrics to Kafka.
- Support for cluster-wide proxy environments.
- The Red Hat build of OpenTelemetry Operator creates the Prometheus **ServiceMonitor** custom resource if the Prometheus exporter is enabled.
- The Operator enables the **Instrumentation** custom resource that allows injecting upstream OpenTelemetry auto-instrumentation libraries.

1.2.4.2. Removal notice

In Red Hat build of OpenTelemetry 3.0, the Jaeger exporter has been removed. Bug fixes and support are provided only through the end of the 2.9 lifecycle. As an alternative to the Jaeger exporter for sending data to the Jaeger collector, you can use the OTLP exporter instead.

1.2.4.3. Bug fixes

This update introduces the following bug fixes:

- Fixed support for disconnected environments when using the **oc adm catalog mirror** CLI command.

1.2.4.4. Known issues

There is currently a known issue:

- Currently, the cluster monitoring of the Red Hat build of OpenTelemetry Operator is disabled due to a bug ([TRACING-3761](#)). The bug is preventing the cluster monitoring from scraping metrics from the Red Hat build of OpenTelemetry Operator due to a missing label **openshift.io/cluster-monitoring=true** that is required for the cluster monitoring and service monitor object.

Workaround

You can enable the cluster monitoring as follows:

1. Add the following label in the Operator namespace: **oc label namespace openshift-opentelemetry-operator openshift.io/cluster-monitoring=true**
2. Create a service monitor, role, and role binding:

```
apiVersion: monitoring.coreos.com/v1
```

```

kind: ServiceMonitor
metadata:
  name: opentelemetry-operator-controller-manager-metrics-service
  namespace: openshift-opentelemetry-operator
spec:
  endpoints:
    - bearerTokenFile: /var/run/secrets/kubernetes.io/serviceaccount/token
      path: /metrics
      port: https
      scheme: https
      tlsConfig:
        insecureSkipVerify: true
  selector:
    matchLabels:
      app.kubernetes.io/name: opentelemetry-operator
      control-plane: controller-manager
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: otel-operator-prometheus
  namespace: openshift-opentelemetry-operator
  annotations:
    include.release.openshift.io/self-managed-high-availability: "true"
    include.release.openshift.io/single-node-developer: "true"
rules:
- apiGroups:
  - ""
  resources:
  - services
  - endpoints
  - pods
  verbs:
  - get
  - list
  - watch
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: otel-operator-prometheus
  namespace: openshift-opentelemetry-operator
  annotations:
    include.release.openshift.io/self-managed-high-availability: "true"
    include.release.openshift.io/single-node-developer: "true"
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: otel-operator-prometheus
subjects:
- kind: ServiceAccount
  name: prometheus-k8s
  namespace: openshift-monitoring

```

1.2.5. Release notes for Red Hat build of OpenTelemetry 2.9.2



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.9.2 is based on the open source [OpenTelemetry](#) release 0.81.0.

1.2.5.1. CVEs

- This release fixes [CVE-2023-46234](#).

1.2.5.2. Known issues

There is currently a known issue:

- Currently, you must manually set [Operator maturity](#) to Level IV, Deep Insights. ([TRACING-3431](#))

1.2.6. Release notes for Red Hat build of OpenTelemetry 2.9.1



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.9.1 is based on the open source [OpenTelemetry](#) release 0.81.0.

1.2.6.1. CVEs

- This release fixes [CVE-2023-44487](#).

1.2.6.2. Known issues

There is currently a known issue:

- Currently, you must manually set [Operator maturity](#) to Level IV, Deep Insights. ([TRACING-3431](#))

1.2.7. Release notes for Red Hat build of OpenTelemetry 2.9



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.9 is based on the open source [OpenTelemetry](#) release 0.81.0.

1.2.7.1. New features and enhancements

This release introduces the following enhancements for the Red Hat build of OpenTelemetry:


- Support OTLP metrics ingestion. The metrics can be forwarded and stored in the **user-workload-monitoring** via the Prometheus exporter.
- Support the [Operator maturity](#) Level IV, Deep Insights, which enables upgrading and monitoring of **OpenTelemetry Collector** instances and the Red Hat build of OpenTelemetry Operator.
- Report traces and metrics from remote clusters using OTLP or HTTP and HTTPS.
- Collect OpenShift Container Platform resource attributes via the **resourcedetection** processor.
- Support the **managed** and **unmanaged** states in the **OpenTelemetryCollector** custom resource.

1.2.7.2. Known issues

There is currently a known issue:

- Currently, you must manually set [Operator maturity](#) to Level IV, Deep Insights. ([TRACING-3431](#))

1.2.8. Release notes for Red Hat build of OpenTelemetry 2.8



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.8 is based on the open source [OpenTelemetry](#) release 0.74.0.

1.2.8.1. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.9. Release notes for Red Hat build of OpenTelemetry 2.7



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.7 is based on the open source [OpenTelemetry](#) release 0.63.1.

1.2.9.1. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.10. Release notes for Red Hat build of OpenTelemetry 2.6



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.6 is based on the open source [OpenTelemetry](#) release 0.60.

1.2.10.1. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.11. Release notes for Red Hat build of OpenTelemetry 2.5



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.5 is based on the open source [OpenTelemetry](#) release 0.56.

1.2.11.1. New features and enhancements


This update introduces the following enhancement:

- Support for collecting Kubernetes resource attributes to the Red Hat build of OpenTelemetry Operator.

1.2.11.2. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.12. Release notes for Red Hat build of OpenTelemetry 2.4



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.4 is based on the open source [OpenTelemetry](#) release 0.49.

1.2.12.1. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.13. Release notes for Red Hat build of OpenTelemetry 2.3



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.3.1 is based on the open source [OpenTelemetry](#) release 0.44.1.

Red Hat build of OpenTelemetry 2.3.0 is based on the open source [OpenTelemetry](#) release 0.44.0.

1.2.13.1. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.14. Release notes for Red Hat build of OpenTelemetry 2.2



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.2 is based on the open source [OpenTelemetry](#) release 0.42.0.

1.2.14.1. Technology Preview features

The unsupported OpenTelemetry Collector components included in the 2.1 release are removed.

1.2.14.2. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.15. Release notes for Red Hat build of OpenTelemetry 2.1



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.1 is based on the open source [OpenTelemetry](#) release 0.41.1.

1.2.15.1. Technology Preview features

This release introduces a breaking change to how to configure certificates in the OpenTelemetry custom resource file. With this update, the `ca_file` moves under `tls` in the custom resource, as shown in the following examples.

CA file configuration for OpenTelemetry version 0.33

```
spec:
  mode: deployment
  config: |
    exporters:
      jaeger:
        endpoint: jaeger-production-collector-headless.tracing-system.svc:14250
        ca_file: "/var/run/secrets/kubernetes.io/serviceaccount/service-ca.crt"
```


CA file configuration for OpenTelemetry version 0.41.1

```
spec:
  mode: deployment
  config: |
    exporters:
      jaeger:
        endpoint: jaeger-production-collector-headless.tracing-system.svc:14250
        tls:
          ca_file: "/var/run/secrets/kubernetes.io/serviceaccount/service-ca.crt"
```

1.2.15.2. Bug fixes

This release addresses Common Vulnerabilities and Exposures (CVEs) and bug fixes.

1.2.16. Release notes for Red Hat build of OpenTelemetry 2.0



IMPORTANT

The Red Hat build of OpenTelemetry 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](#).

Red Hat build of OpenTelemetry 2.0 is based on the open source [OpenTelemetry](#) release 0.33.0.

This release adds the Red Hat build of OpenTelemetry as a [Technology Preview](#), which you install using the Red Hat build of OpenTelemetry Operator. Red Hat build of OpenTelemetry is based on the [OpenTelemetry](#) APIs and instrumentation. The Red Hat build of OpenTelemetry includes the OpenTelemetry Operator and Collector. You can use the Collector to receive traces in the OpenTelemetry or Jaeger protocol and send the trace data to the Red Hat build of OpenTelemetry. Other capabilities of the Collector are not supported at this time. The OpenTelemetry Collector allows developers to instrument their code with vendor agnostic APIs, avoiding vendor lock-in and enabling a growing ecosystem of observability tooling.

1.2.17. Getting support

If you experience difficulty with a procedure described in this documentation, or with OpenShift Container Platform in general, visit the [Red Hat Customer Portal](#). From the Customer Portal, you can:

- Search or browse through the Red Hat Knowledgebase of articles and solutions relating to Red Hat products.
- Submit a support case to Red Hat Support.
- Access other product documentation.

To identify issues with your cluster, you can use Insights in [OpenShift Cluster Manager Hybrid Cloud Console](#). Insights provides details about issues and, if available, information on how to solve a problem.

If you have a suggestion for improving this documentation or have found an error, submit a [Jira issue](#) for the most relevant documentation component. Please provide specific details, such as the section name and OpenShift Container Platform version.

1.2.18. Making open source more inclusive

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

CHAPTER 2. INSTALLING

Installing the Red Hat build of OpenTelemetry involves the following steps:

1. Installing the Red Hat build of OpenTelemetry Operator.
2. Creating a namespace for an OpenTelemetry Collector instance.
3. Creating an **OpenTelemetryCollector** custom resource to deploy the OpenTelemetry Collector instance.

2.1. INSTALLING THE RED HAT BUILD OF OPENTELEMETRY FROM THE WEB CONSOLE

You can install the Red Hat build of OpenTelemetry from the **Administrator** view of the web console.

Prerequisites

- You are logged in to the web console as a cluster administrator with the **cluster-admin** role.
- For Red Hat OpenShift Dedicated, you must be logged in using an account with the **dedicated-admin** role.

Procedure

1. Install the Red Hat build of OpenTelemetry Operator:
 - a. Go to **Operators** → **OperatorHub** and search for **Red Hat build of OpenTelemetry Operator**.
 - b. Select the **Red Hat build of OpenTelemetry Operator** that is provided by Red Hat → **Install** → **Install** → **View Operator**.



IMPORTANT

This installs the Operator with the default presets:

- **Update channel** → **stable**
- **Installation mode** → **All namespaces on the cluster**
- **Installed Namespace** → **openshift-operators**
- **Update approval** → **Automatic**

- c. In the **Details** tab of the installed Operator page, under **ClusterServiceVersion details**, verify that the installation **Status** is **Succeeded**.
2. Create a project of your choice for the **OpenTelemetry Collector** instance that you will create in the next step by going to **Home** → **Projects** → **Create Project**.
3. Create an **OpenTelemetry Collector** instance.
 - a. Go to **Operators** → **Installed Operators**.

- b. Select **OpenTelemetry Collector** → **Create OpenTelemetry Collector** → **YAML view**.
- c. In the **YAML view**, customize the **OpenTelemetryCollector** custom resource (CR) with the OTLP, Jaeger, Zipkin receivers and the debug exporter.

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: <project_of_opentelemetry_collector_instance>
spec:
  mode: deployment
  config: |
    receivers:
      otlp:
        protocols:
          grpc:
          http:
      jaeger:
        protocols:
          grpc: {}
          thrift_binary: {}
          thrift_compact: {}
          thrift_http: {}
      zipkin: {}
    processors:
      batch: {}
      memory_limiter:
        check_interval: 1s
        limit_percentage: 50
        spike_limit_percentage: 30
    exporters:
      debug: {}
  service:
    pipelines:
      traces:
        receivers: [otlp,jaeger,zipkin]
        processors: [memory_limiter,batch]
        exporters: [debug]

```

- d. Select **Create**.

Verification

1. Use the **Project**: dropdown list to select the project of the **OpenTelemetry Collector** instance.
2. Go to **Operators** → **Installed Operators** to verify that the **Status** of the **OpenTelemetry Collector** instance is **Condition: Ready**.
3. Go to **Workloads** → **Pods** to verify that all the component pods of the **OpenTelemetry Collector** instance are running.

2.2. INSTALLING THE RED HAT BUILD OF OPENTELEMETRY BY USING THE CLI

You can install the Red Hat build of OpenTelemetry from the command line.

Prerequisites

- An active OpenShift CLI (**oc**) session by a cluster administrator with the **cluster-admin** role.

TIP

- Ensure that your OpenShift CLI (**oc**) version is up to date and matches your OpenShift Container Platform version.
- Run **oc login**:

```
$ oc login --username=<your_username>
```

Procedure

1. Install the Red Hat build of OpenTelemetry Operator:
 - a. Create a project for the Red Hat build of OpenTelemetry Operator by running the following command:

```
$ oc apply -f - << EOF
apiVersion: project.openshift.io/v1
kind: Project
metadata:
  labels:
    kubernetes.io/metadata.name: openshift-opentelemetry-operator
    openshift.io/cluster-monitoring: "true"
  name: openshift-opentelemetry-operator
EOF
```

- b. Create an Operator group by running the following command:

```
$ oc apply -f - << EOF
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: openshift-opentelemetry-operator
  namespace: openshift-opentelemetry-operator
spec:
  upgradeStrategy: Default
EOF
```

- c. Create a subscription by running the following command:

```
$ oc apply -f - << EOF
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: opentelemetry-product
  namespace: openshift-opentelemetry-operator
spec:
  channel: stable
```



```
installPlanApproval: Automatic
name: opentelemetry-product
source: redhat-operators
sourceNamespace: openshift-marketplace
EOF
```

- d. Check the Operator status by running the following command:

```
$ oc get csv -n openshift-opentelemetry-operator
```

2. Create a project of your choice for the OpenTelemetry Collector instance that you will create in a subsequent step:


- To create a project without metadata, run the following command:

```
$ oc new-project <project_of_opentelemetry_collector_instance>
```

- To create a project with metadata, run the following command:

```
$ oc apply -f - << EOF
apiVersion: project.openshift.io/v1
kind: Project
metadata:
  name: <project_of_opentelemetry_collector_instance>
EOF
```

3. Create an OpenTelemetry Collector instance in the project that you created for it.



NOTE

You can create multiple OpenTelemetry Collector instances in separate projects on the same cluster.

- a. Customize the **OpenTelemetry Collector** custom resource (CR) with the OTLP, Jaeger, and Zipkin receivers and the debug exporter:

```
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: <project_of_opentelemetry_collector_instance>
spec:
  mode: deployment
  config: |
    receivers:
      otlp:
        protocols:
          grpc:
          http:
      jaeger:
        protocols:
          grpc: {}
          thrift_binary: {}
          thrift_compact: {}
```

```

    thrift_http: {}
  zipkin:
  processors:
    batch: {}
    memory_limiter:
      check_interval: 1s
      limit_percentage: 50
      spike_limit_percentage: 30
  exporters:
    debug: {}
  service:
    pipelines:
      traces:
        receivers: [otlp,jaeger,zipkin]
        processors: [memory_limiter,batch]
        exporters: [debug]

```

- b. Apply the customized CR by running the following command:

```

$ oc apply -f - << EOF
<OpenTelemetryCollector_custom_resource>
EOF

```

Verification

1. Verify that the **status.phase** of the OpenTelemetry Collector pod is **Running** and the **conditions** are **type: Ready** by running the following command:

```

$ oc get pod -l app.kubernetes.io/managed-by=opentelemetry-
operator,app.kubernetes.io/instance=<namespace>.<instance_name> -o yaml

```

2. Get the OpenTelemetry Collector service by running the following command:

```

$ oc get service -l app.kubernetes.io/managed-by=opentelemetry-
operator,app.kubernetes.io/instance=<namespace>.<instance_name>

```

2.3. ADDITIONAL RESOURCES

- [Creating a cluster admin](#)
- [OperatorHub.io](#)
- [Accessing the web console](#)
- [Installing from OperatorHub using the web console](#)
- [Creating applications from installed Operators](#)
- [Getting started with the OpenShift CLI](#)

CHAPTER 3. CONFIGURING THE COLLECTOR

The Red Hat build of OpenTelemetry Operator uses a custom resource definition (CRD) file that defines the architecture and configuration settings to be used when creating and deploying the Red Hat build of OpenTelemetry resources. You can install the default configuration or modify the file.

3.1. OPENTELEMETRY COLLECTOR CONFIGURATION OPTIONS

The OpenTelemetry Collector consists of five types of components that access telemetry data:

Receivers

A receiver, which can be push or pull based, is how data gets into the Collector. Generally, a receiver accepts data in a specified format, translates it into the internal format, and passes it to processors and exporters defined in the applicable pipelines. By default, no receivers are configured. One or more receivers must be configured. Receivers may support one or more data sources.

Processors

Optional. Processors process the data between it is received and exported. By default, no processors are enabled. Processors must be enabled for every data source. Not all processors support all data sources. Depending on the data source, multiple processors might be enabled. Note that the order of processors matters.

Exporters

An exporter, which can be push or pull based, is how you send data to one or more back ends or destinations. By default, no exporters are configured. One or more exporters must be configured. Exporters can support one or more data sources. Exporters might be used with their default settings, but many exporters require configuration to specify at least the destination and security settings.

Connectors

A connector connects two pipelines. It consumes data as an exporter at the end of one pipeline and emits data as a receiver at the start of another pipeline. It can consume and emit data of the same or different data type. It can generate and emit data to summarize the consumed data, or it can merely replicate or route data.

Extensions

An extension adds capabilities to the Collector. For example, authentication can be added to the receivers and exporters automatically.

You can define multiple instances of components in a custom resource YAML file. When configured, these components must be enabled through pipelines defined in the **spec.config.service** section of the YAML file. As a best practice, only enable the components that you need.

Example of the OpenTelemetry Collector custom resource file

```
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: cluster-collector
  namespace: tracing-system
spec:
  mode: deployment
  observability:
    metrics:
      enableMetrics: true
  config: |
    receivers:
```

```

otlp:
  protocols:
    grpc: {}
    http: {}
processors: {}
exporters:
  otlp:
    endpoint: jaeger-production-collector-headless.tracing-system.svc:4317
  tls:
    ca_file: "/var/run/secrets/kubernetes.io/serviceaccount/service-ca.crt"
prometheus:
  endpoint: 0.0.0.0:8889
  resource_to_telemetry_conversion:
    enabled: true # by default resource attributes are dropped
service: ❶
pipelines:
  traces:
    receivers: [otlp]
    processors: []
    exporters: [jaeger]
  metrics:
    receivers: [otlp]
    processors: []
    exporters: [prometheus]

```

- ❶ If a component is configured but not defined in the **service** section, the component is not enabled.

Table 3.1. Parameters used by the Operator to define the OpenTelemetry Collector

Parameter	Description	Values	Default
receivers:	A receiver is how data gets into the Collector. By default, no receivers are configured. There must be at least one enabled receiver for a configuration to be considered valid. Receivers are enabled by being added to a pipeline.	otlp, jaeger, prometheus, zipkin, kafka, opencensus	None
processors:	Processors run through the received data before it is exported. By default, no processors are enabled.	batch, memory_limiter, resourcedetection, attributes, span, k8sattributes, filter, routing	None

Parameter	Description	Values	Default
exporters:	An exporter sends data to one or more back ends or destinations. By default, no exporters are configured. There must be at least one enabled exporter for a configuration to be considered valid. Exporters are enabled by being added to a pipeline. Exporters might be used with their default settings, but many require configuration to specify at least the destination and security settings.	otlp, otlphttp, debug, prometheus, kafka	None
connectors:	Connectors join pairs of pipelines by consuming data as end-of-pipeline exporters and emitting data as start-of-pipeline receivers. Connectors can be used to summarize, replicate, or route consumed data.	spanmetrics	None
extensions:	Optional components for tasks that do not involve processing telemetry data.	bearertokenauth, oauth2client, jaegerremotesampling, pprof, health_check, memory_ballast, zpages	None
service: pipelines:	Components are enabled by adding them to a pipeline under services.pipeline .		
service: pipelines: traces: receivers:	You enable receivers for tracing by adding them under service.pipelines.traces .		None

Parameter	Description	Values	Default
<code>service: pipelines: traces: processors:</code>	You enable processors for tracing by adding them under <code>service.pipelines.traces</code> .		None
<code>service: pipelines: traces: exporters:</code>	You enable exporters for tracing by adding them under <code>service.pipelines.traces</code> .		None
<code>service: pipelines: metrics: receivers:</code>	You enable receivers for metrics by adding them under <code>service.pipelines.metrics</code> .		None
<code>service: pipelines: metrics: processors:</code>	You enable processors for metrics by adding them under <code>service.pipelines.metrics</code> .		None
<code>service: pipelines: metrics: exporters:</code>	You enable exporters for metrics by adding them under <code>service.pipelines.metrics</code> .		None

3.2. OPENTELEMETRY COLLECTOR COMPONENTS

3.2.1. Receivers

Receivers get data into the Collector.

3.2.1.1. OTLP Receiver

The OTLP Receiver ingests traces, metrics, and logs by using the OpenTelemetry Protocol (OTLP). The OTLP Receiver ingests traces and metrics using the OpenTelemetry protocol (OTLP).

OpenTelemetry Collector custom resource with an enabled OTLP Receiver

`config: |`

```

receivers:
  otlp:
    protocols:
      grpc:
        endpoint: 0.0.0.0:4317 1
        tls: 2
          ca_file: ca.pem
          cert_file: cert.pem
          key_file: key.pem
          client_ca_file: client.pem 3
          reload_interval: 1h 4
      http:
        endpoint: 0.0.0.0:4318 5
        tls: 6

service:
  pipelines:
    traces:
      receivers: [otlp]
    metrics:
      receivers: [otlp]

```

- 1 The OTLP gRPC endpoint. If omitted, the default **0.0.0.0:4317** is used.
- 2 The server-side TLS configuration. Defines paths to TLS certificates. If omitted, the TLS is disabled.
- 3 The path to the TLS certificate at which the server verifies a client certificate. This sets the value of **ClientCAs** and **ClientAuth** to **RequireAndVerifyClientCert** in the **TLSConfig**. For more information, see the [Config of the Golang TLS package](#).
- 4 Specifies the time interval at which the certificate is reloaded. If the value is not set, the certificate is never reloaded. The **reload_interval** field accepts a string containing valid units of time such as **ns**, **us** (or **µs**), **ms**, **s**, **m**, **h**.
- 5 The OTLP HTTP endpoint. The default value is **0.0.0.0:4318**.
- 6 The server-side TLS configuration. For more information, see the **grpc** protocol configuration section.

3.2.1.2. Jaeger Receiver

The Jaeger Receiver ingests traces in the Jaeger formats.

OpenTelemetry Collector custom resource with an enabled Jaeger Receiver

```

config: |
  receivers:
    jaeger:
      protocols:
        grpc:
          endpoint: 0.0.0.0:14250 1
        thrift_http:
          endpoint: 0.0.0.0:14268 2

```

```

thrift_compact:
  endpoint: 0.0.0.0:6831 3
thrift_binary:
  endpoint: 0.0.0.0:6832 4
tls: 5

service:
  pipelines:
  traces:
    receivers: [jaeger]

```

- 1 The Jaeger gRPC endpoint. If omitted, the default **0.0.0.0:14250** is used.
- 2 The Jaeger Thrift HTTP endpoint. If omitted, the default **0.0.0.0:14268** is used.
- 3 The Jaeger Thrift Compact endpoint. If omitted, the default **0.0.0.0:6831** is used.
- 4 The Jaeger Thrift Binary endpoint. If omitted, the default **0.0.0.0:6832** is used.
- 5 The server-side TLS configuration. See the OTLP Receiver configuration section for more details.

3.2.1.3. Host Metrics Receiver



IMPORTANT

The Host Metrics Receiver 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](#).

The Host Metrics Receiver ingests metrics in the OTLP format.

OpenTelemetry Collector custom resource with an enabled Host Metrics Receiver

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-hostfs-daemonset
  namespace: <namespace>
---
apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
allowHostDirVolumePlugin: true
allowHostIPC: false
allowHostNetwork: false
allowHostPID: true
allowHostPorts: false
allowPrivilegeEscalation: true
allowPrivilegedContainer: true

```



```

allowedCapabilities: null
defaultAddCapabilities:
- SYS_ADMIN
fsGroup:
  type: RunAsAny
groups: []
metadata:
  name: otel-hostmetrics
readOnlyRootFilesystem: true
runAsUser:
  type: RunAsAny
seLinuxContext:
  type: RunAsAny
supplementalGroups:
  type: RunAsAny
users:
- system:serviceaccount:<namespace>:otel-hostfs-daemonset
volumes:
- configMap
- emptyDir
- hostPath
- projected
---
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: <namespace>
spec:
  serviceAccount: otel-hostfs-daemonset
  mode: daemonset
  volumeMounts:
  - mountPath: /hostfs
    name: host
    readOnly: true
  volumes:
  - hostPath:
    path: /
    name: host
  config: |
    receivers:
      hostmetrics:
        collection_interval: 10s 1
        initial_delay: 1s 2
        root_path: / 3
        scrapers: 4
          cpu:
          memory:
          disk:
    service:
      pipelines:
        metrics:
          receivers: [hostmetrics]


```

1

Sets the time interval for host metrics collection. If omitted, the default value is **1m**.

- 2 Sets the initial time delay for host metrics collection. If omitted, the default value is **1s**.
- 3 Configures the **root_path** so that the Host Metrics Receiver knows where the root filesystem is. If running multiple instances of the Host Metrics Receiver, set the same **root_path** value for each instance.
- 4 Lists the enabled host metrics scrapers. Available scrapers are **cpu**, **disk**, **load**, **filesystem**, **memory**, **network**, **paging**, **processes**, and **process**.

3.2.1.4. Kubernetes Objects Receiver



IMPORTANT

The Kubernetes Objects Receiver 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](#).

The Kubernetes Objects Receiver pulls or watches objects to be collected from the Kubernetes API server. This receiver watches primarily Kubernetes events, but it can collect any type of Kubernetes objects. This receiver gathers telemetry for the cluster as a whole, so only one instance of this receiver suffices for collecting all the data.

OpenTelemetry Collector custom resource with an enabled Kubernetes Objects Receiver

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-k8sobj
  namespace: <namespace>
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-k8sobj
  namespace: <namespace>
rules:
- apiGroups:
  - ""
  resources:
  - events
  - pods
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - "events.k8s.io"
  resources:

```

```

- events
verbs:
- watch
- list
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-k8sobj
subjects:
- kind: ServiceAccount
  name: otel-k8sobj
  namespace: <namespace>
roleRef:
  kind: ClusterRole
  name: otel-k8sobj
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel-k8s-obj
  namespace: <namespace>
spec:
  serviceAccount: otel-k8sobj
  image: ghcr.io/os-observability/redhat-opentelemetry-collector/redhat-opentelemetry-collector:main
  mode: deployment
  config: |
    receivers:
      k8sobjects:
        auth_type: serviceAccount
        objects:
          - name: pods 1
            mode: pull 2
            interval: 30s 3
            label_selector: 4
            field_selector: 5
            namespaces: [<namespace>,...] 6
          - name: events
            mode: watch
        exporters:
          debug:
            service:
              pipelines:
                logs:
                  receivers: [k8sobjects]
                  exporters: [debug]

```

- 1** The Resource name that this receiver observes: for example, **pods**, **deployments**, or **events**.
- 2** The observation mode that this receiver uses: **pull** or **watch**.
- 3** Only applicable to the pull mode. The request interval for pulling an object. If omitted, the default value is **1h**.

- 4 The label selector to define targets.
- 5 The field selector to filter targets.
- 6 The list of namespaces to collect events from. If omitted, the default value is **all**.

3.2.1.5. Kubelet Stats Receiver



IMPORTANT

The Kubelet Stats Receiver 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](#).

The Kubelet Stats Receiver extracts metrics related to nodes, pods, containers, and volumes from the kubelet's API server. These metrics are then channeled through the metrics-processing pipeline for additional analysis.

OpenTelemetry Collector custom resource with an enabled Kubelet Stats Receiver

```
# ...
config: |
  receivers:
    kubeletstats:
      collection_interval: 20s
      auth_type: "serviceAccount"
      endpoint: "https://${env:K8S_NODE_NAME}:10250"
      insecure_skip_verify: true
  service:
    pipelines:
      metrics:
        receivers: [kubeletstats]
env:
  - name: K8S_NODE_NAME 1
    valueFrom:
      fieldRef:
        fieldPath: spec.nodeName
# ...
```

- 1 Sets the **K8S_NODE_NAME** to authenticate to the API.

The Kubelet Stats Receiver requires additional permissions for the service account used for running the OpenTelemetry Collector.

Permissions required by the service account

```
apiVersion: rbac.authorization.k8s.io/v1
```

```

kind: ClusterRole
metadata:
  name: otel-collector
rules:
- apiGroups: []
  resources: ['nodes/stats']
  verbs: ['get', 'watch', 'list']
- apiGroups: [""]
  resources: ["nodes/proxy"] ❶
  verbs: ["get"]

```

- ❶ The permissions required when using the **extra_metadata_labels** or **request_utilization** or **limit_utilization** metrics.

3.2.1.6. Prometheus Receiver

The Prometheus Receiver is currently a [Technology Preview](#) feature only.

The Prometheus Receiver scrapes the metrics endpoints.

OpenTelemetry Collector custom resource with an enabled Prometheus Receiver

```

config: |
  receivers:
    prometheus:
      config:
        scrape_configs: ❶
        - job_name: 'my-app' ❷
          scrape_interval: 5s ❸
          static_configs:
            - targets: ['my-app.example.svc.cluster.local:8888'] ❹
      service:
        pipelines:
          metrics:
            receivers: [prometheus]

```

- ❶ Scrapes configurations using the Prometheus format.
- ❷ The Prometheus job name.
- ❸ The Interval for scraping the metrics data. Accepts time units. The default value is **1m**.
- ❹ The targets at which the metrics are exposed. This example scrapes the metrics from a **my-app** application in the **example** project.

3.2.1.7. Zipkin Receiver

The Zipkin Receiver ingests traces in the Zipkin v1 and v2 formats.

OpenTelemetry Collector custom resource with the enabled Zipkin Receiver

```

config: |

```

```

receivers:
  zipkin:
    endpoint: 0.0.0.0:9411 1
    tls: 2

service:
  pipelines:
    traces:
      receivers: [zipkin]

```

- 1 The Zipkin HTTP endpoint. If omitted, the default **0.0.0.0:9411** is used.
- 2 The server-side TLS configuration. See the OTLP Receiver configuration section for more details.

3.2.1.8. Kafka Receiver

The Kafka Receiver is currently a [Technology Preview](#) feature only.

The Kafka Receiver receives traces, metrics, and logs from Kafka in the OTLP format.

OpenTelemetry Collector custom resource with the enabled Kafka Receiver

```

config: |
  receivers:
    kafka:
      brokers: ["localhost:9092"] 1
      protocol_version: 2.0.0 2
      topic: otlp_spans 3
      auth:
        plain_text: 4
          username: example
          password: example
        tls: 5
          ca_file: ca.pem
          cert_file: cert.pem
          key_file: key.pem
          insecure: false 6
          server_name_override: kafka.example.corp 7
      service:
        pipelines:
          traces:
            receivers: [kafka]

```

- 1 The list of Kafka brokers. The default is **localhost:9092**.
- 2 The Kafka protocol version. For example, **2.0.0**. This is a required field.
- 3 The name of the Kafka topic to read from. The default is **otlp_spans**.
- 4 The plaintext authentication configuration. If omitted, plaintext authentication is disabled.
- 5 The client-side TLS configuration. Defines paths to the TLS certificates. If omitted, TLS authentication is disabled.

- 6 Disables verifying the server's certificate chain and host name. The default is **false**.
- 7 ServerName indicates the name of the server requested by the client to support virtual hosting.

3.2.1.9. Kubernetes Cluster Receiver



IMPORTANT

The Kubernetes Cluster Receiver 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](#).

The Kubernetes Cluster Receiver gathers cluster metrics and entity events from the Kubernetes API server. It uses the Kubernetes API to receive information about updates. Authentication for this receiver is only supported through service accounts.

OpenTelemetry Collector custom resource with the enabled Kubernetes Cluster Receiver

```
# ...
receivers:
  k8s_cluster:
    distribution: openshift
    collection_interval: 10s
exporters:
  debug:
service:
  pipelines:
    metrics:
      receivers: [k8s_cluster]
      exporters: [debug]
    logs/entity_events:
      receivers: [k8s_cluster]
      exporters: [debug]
# ...
```

This receiver requires a configured service account, RBAC rules for the cluster role, and the cluster role binding that binds the RBAC with the service account.

ServiceAccount object

```
apiVersion: v1
kind: ServiceAccount
metadata:
  labels:
    app: otelcontribcol
  name: otelcontribcol
```

RBAC rules for the ClusterRole object

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otelcontribcol
  labels:
    app: otelcontribcol
rules:
- apiGroups:
  - quota.openshift.io
  resources:
  - clusterresourcequotas
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - ""
  resources:
  - events
  - namespaces
  - namespaces/status
  - nodes
  - nodes/spec
  - pods
  - pods/status
  - replicationcontrollers
  - replicationcontrollers/status
  - resourcequotas
  - services
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - apps
  resources:
  - daemonsets
  - deployments
  - replicaset
  - statefulsets
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - extensions
  resources:
  - daemonsets
  - deployments
  - replicaset
  verbs:
  - get
  - list
  - watch
```



```

- apiGroups:
  - batch
  resources:
  - jobs
  - cronjobs
  verbs:
  - get
  - list
  - watch
- apiGroups:
  - autoscaling
  resources:
  - horizontalpodautoscalers
  verbs:
  - get
  - list
  - watch

```

ClusterRoleBinding object

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otelcontribcol
  labels:
    app: otelcontribcol
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: otelcontribcol
subjects:
- kind: ServiceAccount
  name: otelcontribcol
  namespace: default

```

3.2.1.10. OpenCensus Receiver

The OpenCensus Receiver provides backwards compatibility with the OpenCensus project for easier migration of instrumented codebases. It receives metrics and traces in the OpenCensus format via gRPC or HTTP and Json.

OpenTelemetry Collector custom resource with the enabled OpenCensus Receiver

```


config: |
  receivers:
    opencensus:
      endpoint: 0.0.0.0:9411 1
      tls: 2
      cors_allowed_origins: 3
        - https://*.<example>.com
  service:
    pipelines:

```

```
traces:
  receivers: [opencensus]
  ...
```

- 1 The OpenCensus endpoint. If omitted, the default is **0.0.0.0:55678**.
- 2 The server-side TLS configuration. See the OTLP Receiver configuration section for more details.
- 3 You can also use the HTTP JSON endpoint to optionally configure CORS, which is enabled by specifying a list of allowed CORS origins in this field. Wildcards with * are accepted under the **cors_allowed_origins**. To match any origin, enter only *.

3.2.1.11. Filelog Receiver



IMPORTANT

The Filelog Receiver 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](#).

The Filelog Receiver tails and parses logs from files.

OpenTelemetry Collector custom resource with the enabled Filelog Receiver that tails a text file

```
receivers:
  filelog:
    include: [ /simple.log ] 1
    operators: 2
    - type: regex_parser
      regex: '^(?P<time>\d{4}-\d{2}-\d{2} \d{2}:\d{2}:\d{2}) (?P<sev>[A-Z]*) (?P<msg>.*)*$'
      timestamp:
        parse_from: attributes.time
        layout: '%Y-%m-%d %H:%M:%S'
      severity:
        parse_from: attributes.sev
```

- 1 A list of file glob patterns that match the file paths to be read.
- 2 An array of Operators. Each Operator performs a simple task such as parsing a timestamp or JSON. To process logs into a desired format, chain the Operators together.

3.2.1.12. Journald Receiver



IMPORTANT

The Journald Receiver 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](#).

The Journald Receiver parses **journald** events from the **systemd** journal and sends them as logs.

OpenTelemetry Collector custom resource with the enabled Journald Receiver

```
kubectl apply -f - <<EOF
apiVersion: v1
kind: Namespace
metadata:
  name: otel-journald
  labels:
    security.openshift.io/scc.podSecurityLabelSync: "false"
    pod-security.kubernetes.io/enforce: "privileged"
    pod-security.kubernetes.io/audit: "privileged"
    pod-security.kubernetes.io/warn: "privileged"
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: privileged-sa
  namespace: otel-journald
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-journald-binding
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: system:openshift:scc:privileged
subjects:
- kind: ServiceAccount
  name: privileged-sa
  namespace: otel-journald
---
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel-journald-logs
  namespace: otel-journald
spec:
  mode: daemonset
  serviceAccount: privileged-sa
  securityContext:
    allowPrivilegeEscalation: false
  capabilities:
```

```

drop:
- CHOWN
- DAC_OVERRIDE
- FOWNER
- FSETID
- KILL
- NET_BIND_SERVICE
- SETGID
- SETPCAP
- SETUID
readOnlyRootFilesystem: true
seLinuxOptions:
  type: spc_t
seccompProfile:
  type: RuntimeDefault
config: |
receivers:
  journald:
    files: /var/log/journal/*/*
    priority: info 1
    units: 2
      - kubelet
      - crio
      - init.scope
      - dnsmasq
    all: true 3
    retry_on_failure:
      enabled: true 4
      initial_interval: 1s 5
      max_interval: 30s 6
      max_elapsed_time: 5m 7
processors:
exporters:
  debug:
    verbosity: detailed
service:
  pipelines:
    logs:
      receivers: [journald]
      exporters: [debug]
volumeMounts:
- name: journal-logs
  mountPath: /var/log/journal/
  readOnly: true
volumes:
- name: journal-logs
  hostPath:
    path: /var/log/journal
tolerations:
- key: node-role.kubernetes.io/master
  operator: Exists
  effect: NoSchedule
EOF

```

1 Filters output by message priorities or priority ranges. The default value is **info**.

- 2 Lists the units to read entries from. If empty, entries are read from all units.
- 3 Includes very long logs and logs with unprintable characters. The default value is **false**.
- 4 If set to **true**, the receiver pauses reading a file and attempts to resend the current batch of logs when encountering an error from downstream components. The default value is **false**.
- 5 The time interval to wait after the first failure before retrying. The default value is **1s**. The units are **ms, s, m, h**.
- 6 The upper bound for the retry backoff interval. When this value is reached, the time interval between consecutive retry attempts remains constant at this value. The default value is **30s**. The supported units are **ms, s, m, h**.
- 7 The maximum time interval, including retry attempts, for attempting to send a logs batch to a downstream consumer. When this value is reached, the data are discarded. If the set value is **0**, retrying never stops. The default value is **5m**. The supported units are **ms, s, m, h**.

3.2.1.13. Kubernetes Events Receiver



IMPORTANT

The Kubernetes Events Receiver 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](#).

The Kubernetes Events Receiver collects events from the Kubernetes API server. The collected events are converted into logs.

OpenShift Container Platform permissions required for the Kubernetes Events Receiver

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-collector
  labels:
    app: otel-collector
rules:
- apiGroups:
  - ""
  resources:
  - events
  - namespaces
  - namespaces/status
  - nodes
  - nodes/spec
  - pods
  - pods/status
```

```
- replicationcontrollers
- replicationcontrollers/status
- resourcequotas
- services
verbs:
- get
- list
- watch
- apiGroups:
- apps
resources:
- daemonsets
- deployments
- replicaset
- statefulsets
verbs:
- get
- list
- watch
- apiGroups:
- extensions
resources:
- daemonsets
- deployments
- replicaset
verbs:
- get
- list
- watch
- apiGroups:
- batch
resources:
- jobs
- cronjobs
verbs:
- get
- list
- watch
- apiGroups:
- autoscaling
resources:
- horizontalpodautoscalers
verbs:
- get
- list
- watch
```

OpenTelemetry Collector custom resource with the enabled Kubernetes Event Receiver

```
serviceAccount: otel-collector 1
config: |
  receivers:
    k8s_events:
      namespaces: [project1, project2] 2
  service:
```

```

pipelines:
  logs:
    receivers: [k8s_events]

```

- 1 The service account of the Collector that has the required ClusterRole **otel-collector** RBAC.
- 2 The list of namespaces to collect events from. The default value is empty, which means that all namespaces are collected.

3.2.2. Processors

Processors run through the data between it is received and exported.

3.2.2.1. Batch Processor

The Batch Processor batches traces and metrics to reduce the number of outgoing connections needed to transfer the telemetry information.

Example of the OpenTelemetry Collector custom resource when using the Batch Processor

```

config: |
  processor:
    batch:
      timeout: 5s
      send_batch_max_size: 10000
  service:
    pipelines:
      traces:
        processors: [batch]
      metrics:
        processors: [batch]

```

Table 3.2. Parameters used by the Batch Processor

Parameter	Description	Default
timeout	Sends the batch after a specific time duration and irrespective of the batch size.	200ms
send_batch_size	Sends the batch of telemetry data after the specified number of spans or metrics.	8192
send_batch_max_size	The maximum allowable size of the batch. Must be equal or greater than the send_batch_size .	0

Parameter	Description	Default
<code>metadata_keys</code>	When activated, a batcher instance is created for each unique set of values found in the client.Metadata .	<code>[]</code>
<code>metadata_cardinality_limit</code>	When the metadata_keys are populated, this configuration restricts the number of distinct metadata key-value combinations processed throughout the duration of the process.	1000

3.2.2.2. Memory Limiter Processor

The Memory Limiter Processor periodically checks the Collector's memory usage and pauses data processing when the soft memory limit is reached. This processor supports traces, metrics, and logs. The preceding component, which is typically a receiver, is expected to retry sending the same data and may apply a backpressure to the incoming data. When memory usage exceeds the hard limit, the Memory Limiter Processor forces garbage collection to run.

Example of the OpenTelemetry Collector custom resource when using the Memory Limiter Processor

```

config: |
  processor:
    memory_limiter:
      check_interval: 1s
      limit_mib: 4000
      spike_limit_mib: 800
  service:
    pipelines:
      traces:
        processors: [batch]
      metrics:
        processors: [batch]

```

Table 3.3. Parameters used by the Memory Limiter Processor

Parameter	Description	Default
<code>check_interval</code>	Time between memory usage measurements. The optimal value is 1s . For spiky traffic patterns, you can decrease the check_interval or increase the spike_limit_mib .	0s

Parameter	Description	Default
<code>limit_mib</code>	The hard limit, which is the maximum amount of memory in MiB allocated on the heap. Typically, the total memory usage of the OpenTelemetry Collector is about 50 MiB greater than this value.	0
<code>spike_limit_mib</code>	Spike limit, which is the maximum expected spike of memory usage in MiB. The optimal value is approximately 20% of limit_mib . To calculate the soft limit, subtract the spike_limit_mib from the limit_mib .	20% of limit_mib
<code>limit_percentage</code>	Same as the limit_mib but expressed as a percentage of the total available memory. The limit_mib setting takes precedence over this setting.	0
<code>spike_limit_percentage</code>	Same as the spike_limit_mib but expressed as a percentage of the total available memory. Intended to be used with the limit_percentage setting.	0

3.2.2.3. Resource Detection Processor

The Resource Detection Processor is currently a [Technology Preview](#) feature only.

The Resource Detection Processor identifies host resource details in alignment with OpenTelemetry's resource semantic standards. Using the detected information, this processor can add or replace the resource values in telemetry data. This processor supports traces and metrics. You can use this processor with multiple detectors such as the Docket metadata detector or the **OTEL_RESOURCE_ATTRIBUTES** environment variable detector.

OpenShift Container Platform permissions required for the Resource Detection Processor

```
kind: ClusterRole
metadata:
  name: otel-collector
rules:
- apiGroups: ["config.openshift.io"]
  resources: ["infrastructures", "infrastructures/status"]
  verbs: ["get", "watch", "list"]
```

OpenTelemetry Collector using the Resource Detection Processor

-

```

config: |
  processor:
    resourcedetection:
      detectors: [openshift]
      override: true
  service:
    pipelines:
      traces:
        processors: [resourcedetection]
      metrics:
        processors: [resourcedetection]

```

OpenTelemetry Collector using the Resource Detection Processor with an environment variable detector

```

config: |
  processors:
    resourcedetection/env:
      detectors: [env] ①
      timeout: 2s
      override: false

```

① Specifies which detector to use. In this example, the environment detector is specified.

3.2.2.4. Attributes Processor

The Attributes Processor is currently a [Technology Preview](#) feature only.

The Attributes Processor can modify attributes of a span, log, or metric. You can configure this processor to filter and match input data and include or exclude such data for specific actions.

This processor operates on a list of actions, executing them in the order specified in the configuration. The following actions are supported:

Insert

Inserts a new attribute into the input data when the specified key does not already exist.

Update

Updates an attribute in the input data if the key already exists.

Upsert

Combines the insert and update actions: Inserts a new attribute if the key does not exist yet. Updates the attribute if the key already exists.

Delete

Removes an attribute from the input data.

Hash

Hashes an existing attribute value as SHA1.

Extract

Extracts values by using a regular expression rule from the input key to the target keys defined in the rule. If a target key already exists, it is overridden similarly to the Span Processor's **to_attributes** setting with the existing attribute as the source.

Convert

Converts an existing attribute to a specified type.

OpenTelemetry Collector using the Attributes Processor

```

config: |
  processors:
    attributes/example:
      actions:
        - key: db.table
          action: delete
        - key: redacted_span
          value: true
          action: upsert
        - key: copy_key
          from_attribute: key_original
          action: update
        - key: account_id
          value: 2245
          action: insert
        - key: account_password
          action: delete
        - key: account_email
          action: hash
        - key: http.status_code
          action: convert
          converted_type: int

```

3.2.2.5. Resource Processor

The Resource Processor is currently a [Technology Preview](#) feature only.

The Resource Processor applies changes to the resource attributes. This processor supports traces, metrics, and logs.

OpenTelemetry Collector using the Resource Detection Processor

```

config: |
  processor:
    attributes:
      - key: cloud.availability_zone
        value: "zone-1"
        action: upsert
      - key: k8s.cluster.name
        from_attribute: k8s-cluster
        action: insert
      - key: redundant-attribute
        action: delete

```

Attributes represent the actions that are applied to the resource attributes, such as delete the attribute, insert the attribute, or upsert the attribute.

3.2.2.6. Span Processor

The Span Processor is currently a [Technology Preview](#) feature only.

The Span Processor modifies the span name based on its attributes or extracts the span attributes from the span name. This processor can also change the span status and include or exclude spans. This processor supports traces.

Span renaming requires specifying attributes for the new name by using the **from_attributes** configuration.

OpenTelemetry Collector using the Span Processor for renaming a span

```
config: |
  processor:
    span:
      name:
        from_attributes: [<key1>, <key2>, ...] 1
        separator: <value> 2
```

1 Defines the keys to form the new span name.

2 An optional separator.

You can use this processor to extract attributes from the span name.

OpenTelemetry Collector using the Span Processor for extracting attributes from a span name

```
config: |
  processor:
    span/to_attributes:
      name:
        to_attributes:
          rules:
            - ^\api\v1\document\(?P<documentId>.*\)update$ 1
```

1 This rule defines how the extraction is to be executed. You can define more rules: for example, in this case, if the regular expression matches the name, a **documentID** attribute is created. In this example, if the input span name is **/api/v1/document/12345678/update**, this results in the **/api/v1/document/{documentId}/update** output span name, and a new **"documentId"="12345678"** attribute is added to the span.

You can have the span status modified.

OpenTelemetry Collector using the Span Processor for status change

```
config: |
  processor:
    span/set_status:
      status:
        code: Error
        description: "<error_description>"
```

3.2.2.7. Kubernetes Attributes Processor

The Kubernetes Attributes Processor is currently a [Technology Preview](#) feature only.

The Kubernetes Attributes Processor enables automatic configuration of spans, metrics, and log resource attributes by using the Kubernetes metadata. This processor supports traces, metrics, and logs. This processor automatically identifies the Kubernetes resources, extracts the metadata from them, and incorporates this extracted metadata as resource attributes into relevant spans, metrics, and logs. It utilizes the Kubernetes API to discover all pods operating within a cluster, maintaining records of their IP addresses, pod UIDs, and other relevant metadata.

Minimum OpenShift Container Platform permissions required for the Kubernetes Attributes Processor

```
kind: ClusterRole
metadata:
  name: otel-collector
rules:
- apiGroups: []
  resources: ['pods', 'namespaces']
  verbs: ['get', 'watch', 'list']
```

OpenTelemetry Collector using the Kubernetes Attributes Processor

```
config: |
  processors:
    k8sattributes:
      filter:
        node_from_env_var: KUBE_NODE_NAME
```

3.2.2.8. Filter Processor

The Filter Processor is currently a [Technology Preview](#) feature only.

The Filter Processor leverages the OpenTelemetry Transformation Language to establish criteria for discarding telemetry data. If any of these conditions are satisfied, the telemetry data are discarded. You can combine the conditions by using the logical OR operator. This processor supports traces, metrics, and logs.

OpenTelemetry Collector custom resource with an enabled OTLP Exporter

```
config: |
  processors:
    filter/otl:
      error_mode: ignore 1
      traces:
        span:
          - 'attributes["container.name"] == "app_container_1"' 2
          - 'resource.attributes["host.name"] == "localhost"' 3
```

1 Defines the error mode. When set to **ignore**, ignores errors returned by conditions. When set to **propagate**, returns the error up the pipeline. An error causes the payload to be dropped from the Collector.

2 Filters the spans that have the **container.name == app_container_1** attribute.

- Filters the spans that have the **host.name == localhost** resource attribute.

3.2.2.9. Routing Processor

The Routing Processor is currently a [Technology Preview](#) feature only.

The Routing Processor routes logs, metrics, or traces to specific exporters. This processor can read a header from an incoming gRPC or plain HTTP request or read a resource attribute, and then direct the trace information to relevant exporters according to the read value.

OpenTelemetry Collector custom resource with an enabled OTLP Exporter

```
config: |
  processors:
    routing:
      from_attribute: X-Tenant 1
      default_exporters: 2
      - jaeger
      table: 3
      - value: acme
        exporters: [jaeger/acme]
  exporters:
    jaeger:
      endpoint: localhost:14250
    jaeger/acme:
      endpoint: localhost:24250
```

- The HTTP header name for the lookup value when performing the route.
- The default exporter when the attribute value is not present in the table in the next section.
- The table that defines which values are to be routed to which exporters.

You can optionally create an **attribute_source** configuration, which defines where to look for the attribute in **from_attribute**. The allowed value is **context** to search the context, which includes the HTTP headers, or **resource** to search the resource attributes.

3.2.2.10. Cumulative to Delta Processor

This processor converts monotonic, cumulative-sum, and histogram metrics to monotonic delta metrics.

You can filter metrics by using the **include:** or **exclude:** fields and specifying the **strict** or **regexp** metric name matching.

This processor does not convert non-monotonic sums and exponential histograms.



IMPORTANT

The Cumulative to Delta Processor 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](#).

Example of an OpenTelemetry Collector custom resource with an enabled Cumulative to Delta Processor

```
config: |
  processors:
    cumulativetodelta:
      include: 1
      match_type: strict 2
      metrics: 3
      - <metric_1_name>
      - <metric_2_name>
      exclude: 4
      match_type: regexp
      metrics:
      - "<regular_expression_for_metric_names>"
```

- 1 Optional: Configures which metrics to include. When omitted, all metrics, except for those listed in the **exclude** field, are converted to delta metrics.
- 2 Defines a value provided in the **metrics** field as a **strict** exact match or **regexp** regular expression.
- 3 Lists the metric names, which are exact matches or matches for regular expressions, of the metrics to be converted to delta metrics. If a metric matches both the **include** and **exclude** filters, the **exclude** filter takes precedence.
- 4 Optional: Configures which metrics to exclude. When omitted, no metrics are excluded from conversion to delta metrics.

3.2.3. Exporters

Exporters send data to one or more back ends or destinations.

3.2.3.1. OTLP Exporter

The OTLP gRPC Exporter exports traces and metrics by using the OpenTelemetry protocol (OTLP).

OpenTelemetry Collector custom resource with an enabled OTLP Exporter

```
config: |
  exporters:
    otlp:
```

```

endpoint: tempo-ingester:4317 1
tls: 2
  ca_file: ca.pem
  cert_file: cert.pem
  key_file: key.pem
  insecure: false 3
  insecure_skip_verify: false # 4
  reload_interval: 1h 5
  server_name_override: <name> 6
headers: 7
  X-Scope-OrgID: "dev"
service:
  pipelines:
  traces:
    exporters: [otlp]
  metrics:
    exporters: [otlp]

```

- 1 The OTLP gRPC endpoint. If the **https://** scheme is used, then client transport security is enabled and overrides the **insecure** setting in the **tls**.
- 2 The client-side TLS configuration. Defines paths to TLS certificates.
- 3 Disables client transport security when set to **true**. The default value is **false** by default.
- 4 Skips verifying the certificate when set to **true**. The default value is **false**.
- 5 Specifies the time interval at which the certificate is reloaded. If the value is not set, the certificate is never reloaded. The **reload_interval** accepts a string containing valid units of time such as **ns, us** (or **µs**), **ms, s, m, h**.
- 6 Overrides the virtual host name of authority such as the authority header field in requests. You can use this for testing.
- 7 Headers are sent for every request performed during an established connection.

3.2.3.2. OTLP HTTP Exporter

The OTLP HTTP Exporter exports traces and metrics by using the OpenTelemetry protocol (OTLP).

OpenTelemetry Collector custom resource with an enabled OTLP Exporter

```

config: |
  exporters:
    otlphttp:
      endpoint: http://tempo-ingester:4318 1
      tls: 2
      headers: 3
        X-Scope-OrgID: "dev"
      disable_keep_alives: false 4
  service:
    pipelines:

```



```
traces:
  exporters: [otlphttp]
metrics:
  exporters: [otlphttp]
```

- 1 The OTLP HTTP endpoint. If the **https://** scheme is used, then client transport security is enabled and overrides the **insecure** setting in the **tls**.
- 2 The client side TLS configuration. Defines paths to TLS certificates.
- 3 Headers are sent in every HTTP request.
- 4 If true, disables HTTP keep-alives. It will only use the connection to the server for a single HTTP request.

3.2.3.3. Debug Exporter

The Debug Exporter prints traces and metrics to the standard output.


OpenTelemetry Collector custom resource with an enabled Debug Exporter

```
config: |
  exporters:
    debug:
      verbosity: detailed 1
  service:
    pipelines:
      traces:
        exporters: [logging]
      metrics:
        exporters: [logging]
```

- 1 Verbosity of the debug export: **detailed** or **normal** or **basic**. When set to **detailed**, pipeline data is verbosely logged. Defaults to **normal**.

3.2.3.4. Load Balancing Exporter

The Load Balancing Exporter consistently exports spans, metrics, and logs according to the **routing_key** configuration.



IMPORTANT

The Load Balancing Exporter 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](#).

OpenTelemetry Collector custom resource with an enabled Load Balancing Exporter

```
# ...
config: |
  exporters:
    loadbalancing:
      routing_key: "service" 1
      protocol:
        otlp: 2
        timeout: 1s
      resolver: 3
      static: 4
        hostnames:
          - backend-1:4317
          - backend-2:4317
      dns: 5
        hostname: otelcol-headless.observability.svc.cluster.local
      k8s: 6
        service: lb-svc.kube-public
      ports:
        - 15317
        - 16317
# ...
```

- 1 The **routing_key: service** exports spans for the same service name to the same Collector instance to provide accurate aggregation. The **routing_key: traceID** exports spans based on their **traceID**. The implicit default is **traceID** based routing.
- 2 The OTLP is the only supported load-balancing protocol. All options of the OTLP exporter are supported.
- 3 You can configure only one resolver.
- 4 The static resolver distributes the load across the listed endpoints.
- 5 You can use the DNS resolver only with a Kubernetes headless service.
- 6 The Kubernetes resolver is recommended.

3.2.3.5. Prometheus Exporter

The Prometheus Exporter is currently a [Technology Preview](#) feature only.

The Prometheus Exporter exports metrics in the Prometheus or OpenMetrics formats.

OpenTelemetry Collector custom resource with an enabled Prometheus Exporter

```
ports:
- name: promexporter 1
  port: 8889
  protocol: TCP
config: |
  exporters:
    prometheus:
      endpoint: 0.0.0.0:8889 2
```

```

tls: ❸
  ca_file: ca.pem
  cert_file: cert.pem
  key_file: key.pem
  namespace: prefix ❹
  const_labels: ❺
    label1: value1
  enable_open_metrics: true ❻
  resource_to_telemetry_conversion: ❼
    enabled: true
  metric_expiration: 180m ❽
  add_metric_suffixes: false ❾
service:
  pipelines:
  metrics:
    exporters: [prometheus]

```

- ❶ Exposes the Prometheus port from the Collector pod and service. You can enable scraping of metrics by Prometheus by using the port name in **ServiceMonitor** or **PodMonitor** custom resource.
- ❷ The network endpoint where the metrics are exposed.
- ❸ The server-side TLS configuration. Defines paths to TLS certificates.
- ❹ If set, exports metrics under the provided value. No default.
- ❺ Key-value pair labels that are applied for every exported metric. No default.
- ❻ If **true**, metrics are exported using the OpenMetrics format. Exemplars are only exported in the OpenMetrics format and only for histogram and monotonic sum metrics such as **counter**. Disabled by default.
- ❼ If **enabled** is **true**, all the resource attributes are converted to metric labels by default. Disabled by default.
- ❽ Defines how long metrics are exposed without updates. The default is **5m**.
- ❾ Adds the metrics types and units suffixes. Must be disabled if the monitor tab in Jaeger console is enabled. The default is **true**.

3.2.3.6. Kafka Exporter

The Kafka Exporter is currently a [Technology Preview](#) feature only.

The Kafka Exporter exports logs, metrics, and traces to Kafka. This exporter uses a synchronous producer that blocks and does not batch messages. You must use it with batch and queued retry processors for higher throughput and resiliency.

OpenTelemetry Collector custom resource with an enabled Kafka Exporter

```

config: |
  exporters:
    kafka:
      brokers: ["localhost:9092"] ❶

```

```

protocol_version: 2.0.0 2
topic: otlp_spans 3
auth:
  plain_text: 4
    username: example
    password: example
  tls: 5
    ca_file: ca.pem
    cert_file: cert.pem
    key_file: key.pem
    insecure: false 6
    server_name_override: kafka.example.corp 7
service:
  pipelines:
  traces:
    exporters: [kafka]

```

- 1 The list of Kafka brokers. The default is **localhost:9092**.
- 2 The Kafka protocol version. For example, **2.0.0**. This is a required field.
- 3 The name of the Kafka topic to read from. The following are the defaults: **otlp_spans** for traces, **otlp_metrics** for metrics, **otlp_logs** for logs.
- 4 The plaintext authentication configuration. If omitted, plaintext authentication is disabled.
- 5 The client-side TLS configuration. Defines paths to the TLS certificates. If omitted, TLS authentication is disabled.
- 6 Disables verifying the server's certificate chain and host name. The default is **false**.
- 7 ServerName indicates the name of the server requested by the client to support virtual hosting.

3.2.4. Connectors

Connectors connect two pipelines.

3.2.4.1. Forward Connector



IMPORTANT

The Forward Connector 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](#).

The Forward Connector merges two pipelines of the same type.

OpenTelemetry Collector custom resource with an enabled Forward Connector

```

receivers:
  otlp:
    protocols:
      grpc:
  jaeger:
    protocols:
      grpc:
processors:
  batch:
exporters:
  otlp:
    endpoint: tempo-simplest-distributor:4317
    tls:
      insecure: true
connectors:
  forward:
service:
  pipelines:
    traces/regiona:
      receivers: [otlp]
      processors: []
      exporters: [forward]
    traces/regionb:
      receivers: [jaeger]
      processors: []
      exporters: [forward]
    traces:
      receivers: [forward]
      processors: [batch]
      exporters: [otlp]

```

3.2.4.2. Spanmetrics Connector

The Spanmetrics Connector is currently a [Technology Preview](#) feature only.

The Spanmetrics Connector aggregates Request, Error, and Duration (R.E.D) OpenTelemetry metrics from span data.

OpenTelemetry Collector custom resource with an enabled Spanmetrics Connector

```

config: |
  connectors:
    spanmetrics:
      metrics_flush_interval: 15s 1
  service:
    pipelines:
      traces:
        exporters: [spanmetrics]
      metrics:
        receivers: [spanmetrics]

```

1 Defines the flush interval of the generated metrics. Defaults to **15s**.

3.2.5. Extensions

Extensions add capabilities to the Collector.

3.2.5.1. BearerTokenAuth Extension

The BearerTokenAuth Extension is currently a [Technology Preview](#) feature only.

The BearerTokenAuth Extension is an authenticator for receivers and exporters that are based on the HTTP and the gRPC protocol. You can use the OpenTelemetry Collector custom resource to configure client authentication and server authentication for the BearerTokenAuth Extension on the receiver and exporter side. This extension supports traces, metrics, and logs.

OpenTelemetry Collector custom resource with client and server authentication configured for the BearerTokenAuth Extension

```
config: |
  extensions:
    bearertokenauth:
      scheme: "Bearer" 1
      token: "<token>" 2
      filename: "<token_file>" 3

  receivers:
    otlp:
      protocols:
        http:
          auth:
            authenticator: bearertokenauth 4

  exporters:
    otlp:
      auth:
        authenticator: bearertokenauth 5

  service:
    extensions: [bearertokenauth]
    pipelines:
      traces:
        receivers: [otlp]
        exporters: [otlp]
```

- 1 You can configure the BearerTokenAuth Extension to send a custom **scheme**. The default is **Bearer**.
- 2 You can add the BearerTokenAuth Extension token as metadata to identify a message.
- 3 Path to a file that contains an authorization token that is transmitted with every message.
- 4 You can assign the authenticator configuration to an OTLP Receiver.
- 5 You can assign the authenticator configuration to an OTLP Exporter.

3.2.5.2. OAuth2Client Extension

The OAuth2Client Extension is currently a [Technology Preview](#) feature only.

The OAuth2Client Extension is an authenticator for exporters that are based on the HTTP and the gRPC protocol. Client authentication for the OAuth2Client Extension is configured in a separate section in the OpenTelemetry Collector custom resource. This extension supports traces, metrics, and logs.

OpenTelemetry Collector custom resource with client authentication configured for the OAuth2Client Extension

```

config: |
  extensions:
    oauth2client:
      client_id: <client_id> 1
      client_secret: <client_secret> 2
      endpoint_params: 3
        audience: <audience>
      token_url: https://example.com/oauth2/default/v1/token 4
      scopes: ["api.metrics"] 5
      # tls settings for the token client
      tls: 6
        insecure: true 7
        ca_file: /var/lib/mycert.pem 8
        cert_file: <cert_file> 9
        key_file: <key_file> 10
      timeout: 2s 11

  receivers:
    otlp:
      protocols:
        http: {}

  exporters:
    otlp:
      auth:
        authenticator: oauth2client 12


  service:
    extensions: [oauth2client]
    pipelines:
      traces:
        receivers: [otlp]
        exporters: [otlp]

```

- 1 Client identifier, which is provided by the identity provider.
- 2 Confidential key used to authenticate the client to the identity provider.
- 3 Further metadata, in the key-value pair format, which is transferred during authentication. For example, **audience** specifies the intended audience for the access token, indicating the recipient of the token.
- 4 The URL of the OAuth2 token endpoint, where the Collector requests access tokens.
- 5 The scopes define the specific permissions or access levels requested by the client.

- 6 The Transport Layer Security (TLS) settings for the token client, which is used to establish a secure connection when requesting tokens.
- 7 When set to **true**, configures the Collector to use an insecure or non-verified TLS connection to call the configured token endpoint.
- 8 The path to a Certificate Authority (CA) file that is used to verify the server's certificate during the TLS handshake.
- 9 The path to the client certificate file that the client must use to authenticate itself to the OAuth2 server if required.
- 10 The path to the client's private key file that is used with the client certificate if needed for authentication.
- 11 Sets a timeout for the token client's request.
- 12 You can assign the authenticator configuration to an OTLP exporter.

3.2.5.3. File Storage Extension



IMPORTANT

The File Storage Extension 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](#).

The File Storage Extension supports traces, metrics, and logs. This extension can persist the state to the local file system. This extension persists the sending queue for the OTLP exporters that are based on the HTTP and the gRPC protocols. This extension requires the read and write access to a directory. This extension can use a default directory, but the default directory must already exist.

OpenTelemetry Collector custom resource with a configured File Storage Extension that persists an OTLP sending queue

```

config: |
  extensions:
    file_storage/all_settings:
      directory: /var/lib/otelcol/mydir 1
      timeout: 1s 2
      compaction:
        on_start: true 3
        directory: /tmp/ 4
        max_transaction_size: 65_536 5
      fsync: false 6

  exporters:
    otlp:

```



```

sending_queue:
  storage: file_storage/all_settings

service:
  extensions: [file_storage/all_settings]
  pipelines:
  traces:
    receivers: [otlp]
    exporters: [otlp]

```

- 1 Specifies the directory in which the telemetry data is stored.
- 2 Specifies the timeout time interval for opening the stored files.
- 3 Starts compaction when the Collector starts. If omitted, the default is **false**.
- 4 Specifies the directory in which the compactor stores the telemetry data.
- 5 Defines the maximum size of the compaction transaction. To ignore the transaction size, set to zero. If omitted, the default is **65536** bytes.
- 6 When set, forces the database to perform an **fsync** call after each write operation. This helps to ensure database integrity if there is an interruption to the database process, but at the cost of performance.

3.2.5.4. OIDC Auth Extension



IMPORTANT

The OIDC Auth Extension 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](#).

The OIDC Auth Extension authenticates incoming requests to receivers by using the OpenID Connect (OIDC) protocol. It validates the ID token in the authorization header against the issuer and updates the authentication context of the incoming request.

OpenTelemetry Collector custom resource with the configured OIDC Auth Extension

```

config: |
  extensions:
    oidc:
      attribute: authorization 1
      issuer_url: https://example.com/auth/realms/opentelemetry 2
      issuer_ca_path: /var/run/tls/issuer.pem 3
      audience: otel-collector 4
      username_claim: email 5
  receivers:

```

```

otlp:
  protocols:
    grpc:
      auth:
        authenticator: oidc
exporters:
  otlp:
    endpoint: <endpoint>
service:
  extensions: [oidc]
  pipelines:
    traces:
      receivers: [otlp]
      exporters: [otlp]

```

- 1 The name of the header that contains the ID token. The default name is **authorization**.
- 2 The base URL of the OIDC provider.
- 3 Optional: The path to the issuer's CA certificate.
- 4 The audience for the token.
- 5 The name of the claim that contains the username. The default name is **sub**.

3.2.5.5. Jaeger Remote Sampling Extension

The Jaeger Remote Sampling Extension is currently a [Technology Preview](#) feature only.

The Jaeger Remote Sampling Extension enables serving sampling strategies after Jaeger's remote sampling API. You can configure this extension to proxy requests to a backing remote sampling server such as a Jaeger collector down the pipeline or to a static JSON file from the local file system.

OpenTelemetry Collector custom resource with a configured Jaeger Remote Sampling Extension

```

config: |
  extensions:
    jaegerremotesampling:
      source:
        reload_interval: 30s 1
      remote:
        endpoint: jaeger-collector:14250 2
        file: /etc/otelcol/sampling_strategies.json 3

  receivers:
    otlp:
      protocols:
        http: {}

  exporters:
    otlp:

  service:

```

```

extensions: [jaegerremotesampling]
pipelines:
  traces:
    receivers: [otlp]
    exporters: [otlp]

```

- 1 The time interval at which the sampling configuration is updated.
- 2 The endpoint for reaching the Jaeger remote sampling strategy provider.
- 3 The path to a local file that contains a sampling strategy configuration in the JSON format.

Example of a Jaeger Remote Sampling strategy file

```

{
  "service_strategies": [
    {
      "service": "foo",
      "type": "probabilistic",
      "param": 0.8,
      "operation_strategies": [
        {
          "operation": "op1",
          "type": "probabilistic",
          "param": 0.2
        },
        {
          "operation": "op2",
          "type": "probabilistic",
          "param": 0.4
        }
      ]
    },
    {
      "service": "bar",
      "type": "ratelimiting",
      "param": 5
    }
  ],
  "default_strategy": {
    "type": "probabilistic",
    "param": 0.5,
    "operation_strategies": [
      {
        "operation": "/health",
        "type": "probabilistic",
        "param": 0.0
      },
      {
        "operation": "/metrics",
        "type": "probabilistic",
        "param": 0.0
      }
    ]
  }
}

```

```

]
}
}

```

3.2.5.6. Performance Profiler Extension

The Performance Profiler Extension is currently a [Technology Preview](#) feature only.

The Performance Profiler Extension enables the Go **net/http/pprof** endpoint. Developers use this extension to collect performance profiles and investigate issues with the service.

OpenTelemetry Collector custom resource with the configured Performance Profiler Extension

```

config: |
  extensions:
    pprof:
      endpoint: localhost:1777 1
      block_profile_fraction: 0 2
      mutex_profile_fraction: 0 3
      save_to_file: test.pprof 4

  receivers:
    otlp:
      protocols:
        http: {}

  exporters:
    otlp:

  service:
    extensions: [pprof]
    pipelines:
      traces:
        receivers: [otlp]
        exporters: [otlp]

```

- 1** The endpoint at which this extension listens. Use **localhost**: to make it available only locally or **""** to make it available on all network interfaces. The default value is **localhost:1777**.
- 2** Sets a fraction of blocking events to be profiled. To disable profiling, set this to **0** or a negative integer. See the [documentation](#) for the **runtime** package. The default value is **0**.
- 3** Set a fraction of mutex contention events to be profiled. To disable profiling, set this to **0** or a negative integer. See the [documentation](#) for the **runtime** package. The default value is **0**.
- 4** The name of the file in which the CPU profile is to be saved. Profiling starts when the Collector starts. Profiling is saved to the file when the Collector is terminated.

3.2.5.7. Health Check Extension

The Health Check Extension is currently a [Technology Preview](#) feature only.

The Health Check Extension provides an HTTP URL for checking the status of the OpenTelemetry Collector. You can use this extension as a liveness and readiness probe on OpenShift.

OpenTelemetry Collector custom resource with the configured Health Check Extension

```

config: |
  extensions:
    health_check:
      endpoint: "0.0.0.0:13133" 1
      tls: 2
        ca_file: "/path/to/ca.crt"
        cert_file: "/path/to/cert.crt"
        key_file: "/path/to/key.key"
      path: "/health/status" 3
      check_collector_pipeline: 4
        enabled: true 5
        interval: "5m" 6
        exporter_failure_threshold: 5 7

  receivers:
    otlp:
      protocols:
        http: {}

  exporters:
    otlp:

  service:
    extensions: [health_check]
    pipelines:
      traces:
        receivers: [otlp]
        exporters: [otlp]

```

- 1 The target IP address for publishing the health check status. The default is **0.0.0.0:13133**.
- 2 The TLS server-side configuration. Defines paths to TLS certificates. If omitted, the TLS is disabled.
- 3 The path for the health check server. The default is `/`.
- 4 Settings for the Collector pipeline health check.
- 5 Enables the Collector pipeline health check. The default is **false**.
- 6 The time interval for checking the number of failures. The default is **5m**.
- 7 The threshold of a number of failures until which a container is still marked as healthy. The default is **5**.

3.2.5.8. Memory Ballast Extension

The Memory Ballast Extension is currently a [Technology Preview](#) feature only.

The Memory Ballast Extension enables applications to configure memory ballast for the process.

OpenTelemetry Collector custom resource with the configured Memory Ballast Extension

```

config: |
  extensions:
    memory_ballast:
      size_mib: 64 1
      size_in_percentage: 20 2

  receivers:
    otlp:
      protocols:
        http: {}

  exporters:
    otlp:

  service:
    extensions: [memory_ballast]
    pipelines:
      traces:
        receivers: [otlp]
        exporters: [otlp]

```

- 1** Sets the memory ballast size in MiB. Takes priority over the **size_in_percentage** if both are specified.
- 2** Sets the memory ballast as a percentage, **1-100**, of the total memory. Supports containerized and physical host environments.

3.2.5.9. zPages Extension

The zPages Extension is currently a [Technology Preview](#) feature only.

The zPages Extension provides an HTTP endpoint for extensions that serve zPages. At the endpoint, this extension serves live data for debugging instrumented components. All core exporters and receivers provide some zPages instrumentation.

zPages are useful for in-process diagnostics without having to depend on a back end to examine traces or metrics.

OpenTelemetry Collector custom resource with the configured zPages Extension

```

config: |
  extensions:
    zpages:
      endpoint: "localhost:55679" 1

  receivers:
    otlp:
      protocols:
        http: {}

  exporters:

```

```

otlp:

service:
  extensions: [zpages]
  pipelines:
    traces:
      receivers: [otlp]
      exporters: [otlp]

```

- 1 Specifies the HTTP endpoint that serves zPages. Use **localhost:** to make it available only locally, or **":"** to make it available on all network interfaces. The default is **localhost:55679**.

3.3. CREATING THE REQUIRED RBAC RESOURCES AUTOMATICALLY

Some Collector components require configuring the RBAC resources.

Procedure


- Add the following permissions to the **opentelemetry-operator-controller-manage** service account so that the Red Hat build of OpenTelemetry Operator can create them automatically:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: generate-processors-rbac
rules:
- apiGroups:
  - rbac.authorization.k8s.io
  resources:
  - clusterrolebindings
  - clusterroles
  verbs:
  - create
  - delete
  - get
  - list
  - patch
  - update
  - watch
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: generate-processors-rbac
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: generate-processors-rbac
subjects:
- kind: ServiceAccount
  name: opentelemetry-operator-controller-manager
  namespace: openshift-opentelemetry-operator

```

3.4. TARGET ALLOCATOR



IMPORTANT

The target allocator 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](#).

The target allocator is an optional component of the OpenTelemetry Operator that shards scrape targets across the deployed fleet of OpenTelemetry Collector instances. The target allocator integrates with the Prometheus **PodMonitor** and **ServiceMonitor** custom resources (CR). When the target allocator is enabled, the OpenTelemetry Operator adds the **http_sd_config** field to the enabled **prometheus** receiver that connects to the target allocator service.

Example OpenTelemetryCollector CR with the enabled target allocator

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: observability
spec:
  mode: statefulset 1
  targetAllocator:
    enabled: true 2
    serviceAccount: 3
  prometheusCR:
    enabled: true 4
    scrapeInterval: 10s
    serviceMonitorSelector: 5
      name: app1
    podMonitorSelector: 6
      name: app2
  config: |
    receivers:
      prometheus: 7
        config:
          scrape_configs: []
    processors:
    exporters:
      debug: {}
  service:
    pipelines:
      metrics:
        receivers: [prometheus]
        processors: []
        exporters: [debug]

```


- 1 When the target allocator is enabled, the deployment mode must be set to **statefulset**.
- 2 Enables the target allocator. Defaults to **false**.
- 3 The service account name of the target allocator deployment. The service account needs to have RBAC to get the **ServiceMonitor**, **PodMonitor** custom resources, and other objects from the cluster to properly set labels on scraped metrics. The default service name is **<collector_name>-targetallocator**.
- 4 Enables integration with the Prometheus **PodMonitor** and **ServiceMonitor** custom resources.
- 5 Label selector for the Prometheus **ServiceMonitor** custom resources. When left empty, enables all service monitors.
- 6 Label selector for the Prometheus **PodMonitor** custom resources. When left empty, enables all pod monitors.
- 7 Prometheus receiver with the minimal, empty **scrape_config: []** configuration option.

The target allocator deployment uses the Kubernetes API to get relevant objects from the cluster, so it requires a custom RBAC configuration.

RBAC configuration for the target allocator service account

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-targetallocator
rules:
  - apiGroups: [""]
    resources:
      - services
      - pods
    verbs: ["get", "list", "watch"]
  - apiGroups: ["monitoring.coreos.com"]
    resources:
      - servicemonitors
      - podmonitors
    verbs: ["get", "list", "watch"]
  - apiGroups: ["discovery.k8s.io"]
    resources:
      - endpointslices
    verbs: ["get", "list", "watch"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-targetallocator
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: otel-targetallocator
subjects:

```

```
- kind: ServiceAccount  
  name: otel-targetallocator 1  
  namespace: observability 2
```

1 The name of the target allocator service account name.

2 The namespace of the target allocator service account.

CHAPTER 4. CONFIGURING THE INSTRUMENTATION



IMPORTANT

OpenTelemetry instrumentation injection 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](#).

The Red Hat build of OpenTelemetry Operator uses a custom resource definition (CRD) file that defines the configuration of the instrumentation.

4.1. OPENTELEMETRY INSTRUMENTATION CONFIGURATION OPTIONS

The Red Hat build of OpenTelemetry can inject and configure the OpenTelemetry auto-instrumentation libraries into your workloads. Currently, the project supports injection of the instrumentation libraries from Go, Java, Node.js, Python, .NET, and the Apache HTTP Server (**httpd**).

Auto-instrumentation in OpenTelemetry refers to the capability where the framework automatically instruments an application without manual code changes. This enables developers and administrators to get observability into their applications with minimal effort and changes to the existing codebase.



IMPORTANT

The Red Hat build of OpenTelemetry Operator only supports the injection mechanism of the instrumentation libraries but does not support instrumentation libraries or upstream images. Customers can build their own instrumentation images or use community images.

4.1.1. Instrumentation options

Instrumentation options are specified in the **OpenTelemetryCollector** custom resource.

Sample OpenTelemetryCollector custom resource file

```
apiVersion: opentelemetry.io/v1alpha1
kind: Instrumentation
metadata:
  name: java-instrumentation
spec:
  env:
    - name: OTEL_EXPORTER_OTLP_TIMEOUT
      value: "20"
  exporter:
    endpoint: http://production-collector.observability.svc.cluster.local:4317
  propagators:
    - w3c
  sampler:
```

```

type: parentbased_traceidratio
argument: "0.25"
java:
  env:
  - name: OTEL_JAVAAGENT_DEBUG
    value: "true"

```

Table 4.1. Parameters used by the Operator to define the Instrumentation

Parameter	Description	Values
env	Common environment variables to define across all the instrumentations.	
exporter	Exporter configuration.	
propagators	Propagators defines inter-process context propagation configuration.	tracecontext, baggage, b3, b3multi, jaeger, ottrace, none
resource	Resource attributes configuration.	
sampler	Sampling configuration.	
apacheHttpd	Configuration for the Apache HTTP Server instrumentation.	
dotnet	Configuration for the .NET instrumentation.	
go	Configuration for the Go instrumentation.	
java	Configuration for the Java instrumentation.	
nodejs	Configuration for the Node.js instrumentation.	
python	Configuration for the Python instrumentation.	

4.1.2. Using the instrumentation CR with Service Mesh

When using the instrumentation custom resource (CR) with Red Hat OpenShift Service Mesh, you must use the **b3multi** propagator.

4.1.2.1. Configuration of the Apache HTTP Server auto-instrumentation

Table 4.2. Parameters for the `.spec.apacheHttpd` field

Name	Description	Default
<code>attrs</code>	Attributes specific to the Apache HTTP Server.	
<code>configPath</code>	Location of the Apache HTTP Server configuration.	<code>/usr/local/apache2/conf</code>
<code>env</code>	Environment variables specific to the Apache HTTP Server.	
<code>image</code>	Container image with the Apache SDK and auto-instrumentation.	
<code>resourceRequirements</code>	The compute resource requirements.	
<code>version</code>	Apache HTTP Server version.	2.4

The PodSpec annotation to enable injection

```
instrumentation.opentelemetry.io/inject-apache-httpd: "true"
```

4.1.2.2. Configuration of the .NET auto-instrumentation

Name	Description
<code>env</code>	Environment variables specific to .NET.
<code>image</code>	Container image with the .NET SDK and auto-instrumentation.
<code>resourceRequirements</code>	The compute resource requirements.

For the .NET auto-instrumentation, the required **OTEL_EXPORTER_OTLP_ENDPOINT** environment variable must be set if the endpoint of the exporters is set to **4317**. The .NET autoinstrumentation uses **http/proto** by default, and the telemetry data must be set to the **4318** port.

The PodSpec annotation to enable injection

```
instrumentation.opentelemetry.io/inject-dotnet: "true"
```

4.1.2.3. Configuration of the Go auto-instrumentation

Name	Description
<code>env</code>	Environment variables specific to Go.
<code>image</code>	Container image with the Go SDK and auto-instrumentation.
<code>resourceRequirements</code>	The compute resource requirements.

The PodSpec annotation to enable injection

```
instrumentation.opentelemetry.io/inject-go: "true"
```

Additional permissions required for the Go auto-instrumentation in the OpenShift cluster

```
apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
  name: otel-go-instrumentation-scc
allowHostDirVolumePlugin: true
allowPrivilegeEscalation: true
allowPrivilegedContainer: true
allowedCapabilities:
- "SYS_PTRACE"
fsGroup:
  type: RunAsAny
runAsUser:
  type: RunAsAny
seLinuxContext:
  type: RunAsAny
seccompProfiles:
- "*"
supplementalGroups:
  type: RunAsAny
```

TIP

The CLI command for applying the permissions for the Go auto-instrumentation in the OpenShift cluster is as follows:

```
$ oc adm policy add-scc-to-user otel-go-instrumentation-scc -z <service_account>
```

4.1.2.4. Configuration of the Java auto-instrumentation

Name	Description
env	Environment variables specific to Java.
image	Container image with the Java SDK and auto-instrumentation.
resourceRequirements	The compute resource requirements.

The PodSpec annotation to enable injection

```
instrumentation.opentelemetry.io/inject-java: "true"
```

4.1.2.5. Configuration of the Node.js auto-instrumentation

Name	Description
env	Environment variables specific to Node.js.
image	Container image with the Node.js SDK and auto-instrumentation.
resourceRequirements	The compute resource requirements.

The PodSpec annotations to enable injection

```
instrumentation.opentelemetry.io/inject-nodejs: "true"
instrumentation.opentelemetry.io/otel-go-auto-target-exe: "/path/to/container/executable"
```

The **instrumentation.opentelemetry.io/otel-go-auto-target-exe** annotation sets the value for the required **OTEL_GO_AUTO_TARGET_EXE** environment variable.

4.1.2.6. Configuration of the Python auto-instrumentation

Name	Description
<code>env</code>	Environment variables specific to Python.
<code>image</code>	Container image with the Python SDK and auto-instrumentation.
<code>resourceRequirements</code>	The compute resource requirements.

For Python auto-instrumentation, the **OTEL_EXPORTER_OTLP_ENDPOINT** environment variable must be set if the endpoint of the exporters is set to **4317**. Python auto-instrumentation uses **http/proto** by default, and the telemetry data must be set to the **4318** port.

The PodSpec annotation to enable injection

```
instrumentation.opentelemetry.io/inject-python: "true"
```

4.1.2.7. Configuration of the OpenTelemetry SDK variables

The OpenTelemetry SDK variables in your pod are configurable by using the following annotation:

```
instrumentation.opentelemetry.io/inject-sdk: "true"
```

Note that all the annotations accept the following values:

true

Injects the **Instrumentation** resource from the namespace.

false

Does not inject any instrumentation.

instrumentation-name

The name of the instrumentation resource to inject from the current namespace.

other-namespace/instrumentation-name

The name of the instrumentation resource to inject from another namespace.

4.1.2.8. Multi-container pods

The instrumentation is run on the first container that is available by default according to the pod specification. In some cases, you can also specify target containers for injection.

Pod annotation

```
instrumentation.opentelemetry.io/container-names: "<container_1>,<container_2>"
```


**NOTE**

The Go auto-instrumentation does not support multi-container auto-instrumentation injection.

CHAPTER 5. SENDING TRACES AND METRICS TO THE OPENTELEMETRY COLLECTOR

You can set up and use the Red Hat build of OpenTelemetry to send traces to the OpenTelemetry Collector or the TempoStack instance.

Sending traces and metrics to the OpenTelemetry Collector is possible with or without sidecar injection.

5.1. SENDING TRACES AND METRICS TO THE OPENTELEMETRY COLLECTOR WITH SIDECAR INJECTION

You can set up sending telemetry data to an OpenTelemetry Collector instance with sidecar injection.

The Red Hat build of OpenTelemetry Operator allows sidecar injection into deployment workloads and automatic configuration of your instrumentation to send telemetry data to the OpenTelemetry Collector.

Prerequisites

- The Red Hat OpenShift distributed tracing platform (Tempo) is installed, and a TempoStack instance is deployed.
- You have access to the cluster through the web console or the OpenShift CLI (**oc**):
 - You are logged in to the web console as a cluster administrator with the **cluster-admin** role.
 - An active OpenShift CLI (**oc**) session by a cluster administrator with the **cluster-admin** role.
 - For Red Hat OpenShift Dedicated, you must have an account with the **dedicated-admin** role.

Procedure

1. Create a project for an OpenTelemetry Collector instance.

```
apiVersion: project.openshift.io/v1
kind: Project
metadata:
  name: observability
```

2. Create a service account.

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-collector-sidecar
  namespace: observability
```

3. Grant the permissions to the service account for the **k8sattributes** and **resourcedetection** processors.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
```

```

metadata:
  name: otel-collector
rules:
- apiGroups: ["", "config.openshift.io"]
  resources: ["pods", "namespaces", "infrastructures", "infrastructures/status"]
  verbs: ["get", "watch", "list"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-collector
subjects:
- kind: ServiceAccount
  name: otel-collector-sidecar
  namespace: observability
roleRef:
  kind: ClusterRole
  name: otel-collector
  apiGroup: rbac.authorization.k8s.io

```

4. Deploy the OpenTelemetry Collector as a sidecar.

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: observability
spec:
  serviceAccount: otel-collector-sidecar
  mode: sidecar
  config: |
    serviceAccount: otel-collector-sidecar
    receivers:
      otlp:
        protocols:
          grpc: {}
          http: {}
    processors:
      batch: {}
      memory_limiter:
        check_interval: 1s
        limit_percentage: 50
        spike_limit_percentage: 30
      resourcedetection:
        detectors: [openshift]
        timeout: 2s
    exporters:
      otlp:
        endpoint: "tempo-<example>-gateway:8090" 1
        tls:
          insecure: true
    service:
      pipelines:
        traces:

```

```

receivers: [jaeger]
processors: [memory_limiter, resourcedetection, batch]
exporters: [otlp]

```

- 1 This points to the Gateway of the TempoStack instance deployed by using the **<example>** Tempo Operator.

5. Create your deployment using the **otel-collector-sidecar** service account.
6. Add the **sidecar.opentelemetry.io/inject: "true"** annotation to your **Deployment** object. This will inject all the needed environment variables to send data from your workloads to the OpenTelemetry Collector instance.

5.2. SENDING TRACES AND METRICS TO THE OPENTELEMETRY COLLECTOR WITHOUT SIDECAR INJECTION

You can set up sending telemetry data to an OpenTelemetry Collector instance without sidecar injection, which involves manually setting several environment variables.

Prerequisites

- The Red Hat OpenShift distributed tracing platform (Tempo) is installed, and a TempoStack instance is deployed.
- You have access to the cluster through the web console or the OpenShift CLI (**oc**):
 - You are logged in to the web console as a cluster administrator with the **cluster-admin** role.
 - An active OpenShift CLI (**oc**) session by a cluster administrator with the **cluster-admin** role.
 - For Red Hat OpenShift Dedicated, you must have an account with the **dedicated-admin** role.

Procedure

1. Create a project for an OpenTelemetry Collector instance.

```

apiVersion: project.openshift.io/v1
kind: Project
metadata:
  name: observability

```

2. Create a service account.

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-collector-deployment
namespace: observability

```

3. Grant the permissions to the service account for the **k8sattributes** and **resourcedetection** processors.

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-collector
rules:
- apiGroups: ["", "config.openshift.io"]
  resources: ["pods", "namespaces", "infrastructures", "infrastructures/status"]
  verbs: ["get", "watch", "list"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-collector
subjects:
- kind: ServiceAccount
  name: otel-collector-deployment
  namespace: observability
roleRef:
  kind: ClusterRole
  name: otel-collector
  apiGroup: rbac.authorization.k8s.io

```

4. Deploy the OpenTelemetry Collector instance with the **OpenTelemetryCollector** custom resource.

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: observability
spec:
  mode: deployment
  serviceAccount: otel-collector-deployment
  config: |
    receivers:
      jaeger:
        protocols:
          grpc: {}
          thrift_binary: {}
          thrift_compact: {}
          thrift_http: {}
      opencensus: {}
      otlp:
        protocols:
          grpc: {}
          http: {}
      zipkin: {}
    processors:
      batch: {}
      k8sattributes: {}
      memory_limiter:
        check_interval: 1s
        limit_percentage: 50
        spike_limit_percentage: 30
      resourcedetection:

```

```

detectors: [openshift]
exporters:
  otlp:
    endpoint: "tempo-<example>-distributor:4317" 1
    tls:
      insecure: true
service:
  pipelines:
    traces:
      receivers: [jaeger, opencensus, otlp, zipkin]
      processors: [memory_limiter, k8sattributes, resourcedetection, batch]
      exporters: [otlp]

```

- 1** This points to the Gateway of the TempoStack instance deployed by using the **<example>** Tempo Operator.

5. Set the environment variables in the container with your instrumented application.

Name	Description	Default value
OTEL_SERVICE_NAME	Sets the value of the service.name resource attribute.	""
OTEL_EXPORTER_OTLP_ENDPOINT	Base endpoint URL for any signal type with an optionally specified port number.	https://localhost:4317
OTEL_EXPORTER_OTLP_CERTIFICATE	Path to the certificate file for the TLS credentials of the gRPC client.	https://localhost:4317
OTEL_TRACES_SAMPLER	Sampler to be used for traces.	parentbased_always_on
OTEL_EXPORTER_OTLP_PROTOCOL	Transport protocol for the OTLP exporter.	grpc
OTEL_EXPORTER_OTLP_TIMEOUT	Maximum time interval for the OTLP exporter to wait for each batch export.	10s
OTEL_EXPORTER_OTLP_INSECURE	Disables client transport security for gRPC requests. An HTTPS schema overrides it.	False

CHAPTER 6. CONFIGURING METRICS FOR THE MONITORING STACK

As a cluster administrator, you can configure the OpenTelemetry Collector custom resource (CR) to perform the following tasks:

- Create a Prometheus **ServiceMonitor** CR for scraping the Collector's pipeline metrics and the enabled Prometheus exporters.
- Configure the Prometheus receiver to scrape metrics from the in-cluster monitoring stack.

6.1. CONFIGURATION FOR SENDING METRICS TO THE MONITORING STACK

One of two following custom resources (CR) configures the sending of metrics to the monitoring stack:

- OpenTelemetry Collector CR
- Prometheus **PodMonitor** CR

A configured OpenTelemetry Collector CR can create a Prometheus **ServiceMonitor** CR for scraping the Collector's pipeline metrics and the enabled Prometheus exporters.

Example of the OpenTelemetry Collector CR with the Prometheus exporter

```
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
spec:
  mode: deployment
  observability:
    metrics:
      enableMetrics: true 1
  config: |
    exporters:
      prometheus:
        endpoint: 0.0.0.0:8889
        resource_to_telemetry_conversion:
          enabled: true # by default resource attributes are dropped
    service:
      telemetry:
        metrics:
          address: ":8888"
      pipelines:
        metrics:
          receivers: [otlp]
          exporters: [prometheus]
```

- 1** Configures the Operator to create the Prometheus **ServiceMonitor** CR to scrape the Collector's internal metrics endpoint and Prometheus exporter metric endpoints. The metrics will be stored in the OpenShift monitoring stack.

Alternatively, a manually created Prometheus **PodMonitor** CR can provide fine control, for example removing duplicated labels added during Prometheus scraping.

Example of the PodMonitor CR that configures the monitoring stack to scrape the Collector metrics

```

apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
  name: otel-collector
spec:
  selector:
    matchLabels:
      app.kubernetes.io/name: <cr_name>-collector 1
  podMetricsEndpoints:
    - port: metrics 2
    - port: promexporter 3
  relabelings:
    - action: labeldrop
      regex: pod
    - action: labeldrop
      regex: container
    - action: labeldrop
      regex: endpoint
  metricRelabelings:
    - action: labeldrop
      regex: instance
    - action: labeldrop
      regex: job

```

- 1** The name of the OpenTelemetry Collector CR.
- 2** The name of the internal metrics port for the OpenTelemetry Collector. This port name is always **metrics**.
- 3** The name of the Prometheus exporter port for the OpenTelemetry Collector.

6.2. CONFIGURATION FOR RECEIVING METRICS FROM THE MONITORING STACK

A configured OpenTelemetry Collector custom resource (CR) can set up the Prometheus receiver to scrape metrics from the in-cluster monitoring stack.

Example of the OpenTelemetry Collector CR for scraping metrics from the in-cluster monitoring stack

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-collector
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-monitoring-view 1
subjects:
  - kind: ServiceAccount

```



```

name: otel-collector
namespace: observability
---
kind: ConfigMap
apiVersion: v1
metadata:
  name: cabundle
  namespace: observability
  annotations:
    service.beta.openshift.io/inject-cabundle: "true" 2
---
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: observability
spec:
  volumeMounts:
    - name: cabundle-volume
      mountPath: /etc/pki/ca-trust/source/service-ca
      readOnly: true
  volumes:
    - name: cabundle-volume
      configMap:
        name: cabundle
  mode: deployment
  config: |
    receivers:
      prometheus: 3
      config:
        scrape_configs:
          - job_name: 'federate'
            scrape_interval: 15s
            scheme: https
            tls_config:
              ca_file: /etc/pki/ca-trust/source/service-ca/service-ca.crt
              bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
            honor_labels: false
            params:
              'match[]':
                - '{__name__="<metric_name>"}' 4
            metrics_path: '/federate'
            static_configs:
              - targets:
                  - "prometheus-k8s.openshift-monitoring.svc.cluster.local:9091"
    exporters:
      debug: 5
      verbosity: detailed
  service:
    pipelines:
      metrics:
        receivers: [prometheus]
        processors: []
        exporters: [debug]

```

- 1 Assigns the **cluster-monitoring-view** cluster role to the service account of the OpenTelemetry Collector so that it can access the metrics data.
- 2 Injects the OpenShift service CA for configuring the TLS in the Prometheus receiver.
- 3 Configures the Prometheus receiver to scrape the federate endpoint from the in-cluster monitoring stack.
- 4 Uses the Prometheus query language to select the metrics to be scraped. See the in-cluster monitoring documentation for more details and limitations of the federate endpoint.
- 5 Configures the debug exporter to print the metrics to the standard output.

6.3. ADDITIONAL RESOURCES

- [Querying metrics by using the federation endpoint for Prometheus](#)

CHAPTER 7. FORWARDING TRACES TO A TEMPOSTACK INSTANCE

To configure forwarding traces to a TempoStack instance, you can deploy and configure the OpenTelemetry Collector. You can deploy the OpenTelemetry Collector in the deployment mode by using the specified processors, receivers, and exporters. For other modes, see the OpenTelemetry Collector documentation linked in *Additional resources*.

Prerequisites

- The Red Hat build of OpenTelemetry Operator is installed.
- The Tempo Operator is installed.
- A TempoStack instance is deployed on the cluster.

Procedure

1. Create a service account for the OpenTelemetry Collector.

Example ServiceAccount

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-collector-deployment
```

2. Create a cluster role for the service account.

Example ClusterRole

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-collector
rules:
  1
  2
- apiGroups: ["", "config.openshift.io"]
  resources: ["pods", "namespaces", "infrastructures", "infrastructures/status"]
  verbs: ["get", "watch", "list"]
```

1 The **k8sattributesprocessor** requires permissions for pods and namespaces resources.

2 The **resourcedetectionprocessor** requires permissions for infrastructures and status.

3. Bind the cluster role to the service account.

Example ClusterRoleBinding

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
```

```

metadata:
  name: otel-collector
subjects:
- kind: ServiceAccount
  name: otel-collector-deployment
  namespace: otel-collector-example
roleRef:
  kind: ClusterRole
  name: otel-collector
  apiGroup: rbac.authorization.k8s.io

```

4. Create the YAML file to define the **OpenTelemetryCollector** custom resource (CR).

Example OpenTelemetryCollector

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
spec:
  mode: deployment
  serviceAccount: otel-collector-deployment
  config: |
    receivers:
      jaeger:
        protocols:
          grpc: {}
          thrift_binary: {}
          thrift_compact: {}
          thrift_http: {}
      opencensus: {}
      otlp:
        protocols:
          grpc: {}
          http: {}
      zipkin: {}
    processors:
      batch: {}
      k8sattributes: {}
      memory_limiter:
        check_interval: 1s
        limit_percentage: 50
        spike_limit_percentage: 30
      resourcedetection:
        detectors: [openshift]
    exporters:
      otlp:
        endpoint: "tempo-simplest-distributor:4317" 1
        tls:
          insecure: true
  service:
    pipelines:
      traces:

```

```

receivers: [jaeger, opencensus, otlp, zipkin] 2
processors: [memory_limiter, k8sattributes, resourcedetection, batch]
exporters: [otlp]

```

- 1** The Collector exporter is configured to export OTLP and points to the Tempo distributor endpoint, "**tempo-simplest-distributor:4317**" in this example, which is already created.
- 2** The Collector is configured with a receiver for Jaeger traces, OpenCensus traces over the OpenCensus protocol, Zipkin traces over the Zipkin protocol, and OTLP traces over the GRPC protocol.

TIP

You can deploy **telemetrygen** as a test:

```

apiVersion: batch/v1
kind: Job
metadata:
  name: telemetrygen
spec:
  template:
    spec:
      containers:
        - name: telemetrygen
          image: ghcr.io/open-telemetry/opentelemetry-collector-contrib/telemetrygen:latest
          args:
            - traces
            - --otlp-endpoint=otel-collector:4317
            - --otlp-insecure
            - --duration=30s
            - --workers=1
      restartPolicy: Never
      backoffLimit: 4

```

Additional resources

- [OpenTelemetry Collector documentation](#)
- [Deployment examples on GitHub](#)

CHAPTER 8. CONFIGURING THE OPENTELEMETRY COLLECTOR METRICS

You can enable metrics and alerts of OpenTelemetry Collector instances.

Prerequisites

- Monitoring for user-defined projects is enabled in the cluster.

Procedure

- To enable metrics of an OpenTelemetry Collector instance, set the **spec.observability.metrics.enableMetrics** field to **true**:

```
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: <name>
spec:
  observability:
    metrics:
      enableMetrics: true
```

Verification

You can use the **Administrator** view of the web console to verify successful configuration:

- Go to **Observe** → **Targets**, filter by **Source: User**, and check that the **ServiceMonitors** in the **opentelemetry-collector-<instance_name>** format have the **Up** status.

Additional resources

- [Enabling monitoring for user-defined projects](#)

CHAPTER 9. GATHERING THE OBSERVABILITY DATA FROM MULTIPLE CLUSTERS

For a multicluster configuration, you can create one OpenTelemetry Collector instance in each one of the remote clusters and then forward all the telemetry data to one OpenTelemetry Collector instance.

Prerequisites

- The Red Hat build of OpenTelemetry Operator is installed.
- The Tempo Operator is installed.
- A TempoStack instance is deployed on the cluster.
- The following mounted certificates: Issuer, self-signed certificate, CA issuer, client and server certificates. To create any of these certificates, see step 1.

Procedure

1. Mount the following certificates in the OpenTelemetry Collector instance, skipping already mounted certificates.
 - a. An Issuer to generate the certificates by using the cert-manager Operator for Red Hat OpenShift.

```
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: selfsigned-issuer
spec:
  selfSigned: {}
```

- b. A self-signed certificate.

```
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: ca
spec:
  isCA: true
  commonName: ca
  subject:
    organizations:
      - Organization # <your_organization_name>
  organizationalUnits:
    - Widgets
  secretName: ca-secret
  privateKey:
    algorithm: ECDSA
    size: 256
  issuerRef:
    name: selfsigned-issuer
    kind: Issuer
    group: cert-manager.io
```

c. A CA issuer.

```
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
  name: test-ca-issuer
spec:
  ca:
    secretName: ca-secret
```

d. The client and server certificates.

```
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: server
spec:
  secretName: server-tls
  isCA: false
  usages:
    - server auth
    - client auth
  dnsNames:
    - "otel.observability.svc.cluster.local" 1
  issuerRef:
    name: ca-issuer
---
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: client
spec:
  secretName: client-tls
  isCA: false
  usages:
    - server auth
    - client auth
  dnsNames:
    - "otel.observability.svc.cluster.local" 2
  issuerRef:
    name: ca-issuer
```

1 List of exact DNS names to be mapped to a solver in the server OpenTelemetry Collector instance.

2 List of exact DNS names to be mapped to a solver in the client OpenTelemetry Collector instance.

2. Create a service account for the OpenTelemetry Collector instance.

Example ServiceAccount

```
apiVersion: v1
kind: ServiceAccount
```



```

metadata:
  name: otel-collector-deployment

```

3. Create a cluster role for the service account.

Example ClusterRole

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-collector
rules:
  1
  2
  - apiGroups: ["", "config.openshift.io"]
    resources: ["pods", "namespaces", "infrastructures", "infrastructures/status"]
    verbs: ["get", "watch", "list"]

```

- 1 The **k8sattributesprocessor** requires permissions for pods and namespace resources.
- 2 The **resourcedetectionprocessor** requires permissions for infrastructures and status.

4. Bind the cluster role to the service account.

Example ClusterRoleBinding

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-collector
subjects:
  - kind: ServiceAccount
    name: otel-collector-deployment
    namespace: otel-collector-<example>
roleRef:
  kind: ClusterRole
  name: otel-collector
  apiGroup: rbac.authorization.k8s.io

```

5. Create the YAML file to define the **OpenTelemetryCollector** custom resource (CR) in the edge clusters.

Example OpenTelemetryCollector custom resource for the edge clusters

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: otel-collector-<example>
spec:
  mode: daemonset
  serviceAccount: otel-collector-deployment
  config: |

```

```

receivers:
  jaeger:
    protocols:
      grpc: {}
      thrift_binary: {}
      thrift_compact: {}
      thrift_http: {}
  opencensus:
  otlp:
    protocols:
      grpc: {}
      http: {}
  zipkin: {}
processors:
  batch: {}
  k8sattributes: {}
  memory_limiter:
    check_interval: 1s
    limit_percentage: 50
    spike_limit_percentage: 30
  resourcedetection:
    detectors: [openshift]
exporters:
  otlphttp:
    endpoint: https://observability-cluster.com:443 1
    tls:
      insecure: false
      cert_file: /certs/server.crt
      key_file: /certs/server.key
      ca_file: /certs/ca.crt
  service:
    pipelines:
      traces:
        receivers: [jaeger, opencensus, otlp, zipkin]
        processors: [memory_limiter, k8sattributes, resourcedetection, batch]
        exporters: [otlp]
volumes:
- name: otel-certs
  secret:
    name: otel-certs
volumeMounts:
- name: otel-certs
  mountPath: /certs

```

- 1** The Collector exporter is configured to export OTLP HTTP and points to the OpenTelemetry Collector from the central cluster.

6. Create the YAML file to define the **OpenTelemetryCollector** custom resource (CR) in the central cluster.

Example OpenTelemetryCollector custom resource for the central cluster

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:

```

```

name: otlp-receiver
namespace: observability
spec:
  mode: "deployment"
  ingress:
    type: route
    route:
      termination: "passthrough"
  config: |
    receivers:
      otlp:
        protocols:
          http:
            tls: ❶
              cert_file: /certs/server.crt
              key_file: /certs/server.key
              client_ca_file: /certs/ca.crt
    exporters:
      logging: {}
      otlp:
        endpoint: "tempo-<simplest>-distributor:4317" ❷
        tls:
          insecure: true
    service:
      pipelines:
        traces:
          receivers: [otlp]
          processors: []
          exporters: [otlp]
  volumes:
    - name: otel-certs
      secret:
        name: otel-certs
  volumeMounts:
    - name: otel-certs
      mountPath: /certs

```

- ❶ The Collector receiver requires the certificates listed in the first step.
- ❷ The Collector exporter is configured to export OTLP and points to the Tempo distributor endpoint, which in this example is **"tempo-simplest-distributor:4317"** and already created.

CHAPTER 10. TROUBLESHOOTING

The OpenTelemetry Collector offers multiple ways to measure its health as well as investigate data ingestion issues.

10.1. GETTING THE OPENTELEMETRY COLLECTOR LOGS

You can get the logs for the OpenTelemetry Collector as follows.

Procedure

1. Set the relevant log level in the **OpenTelemetryCollector** custom resource (CR):

```
config: |
  service:
    telemetry:
      logs:
        level: debug 1
```

- 1** Collector's log level. Supported values include **info**, **warn**, **error**, or **debug**. Defaults to **info**.

2. Use the **oc logs** command or the web console to retrieve the logs.

10.2. EXPOSING THE METRICS

The OpenTelemetry Collector exposes the metrics about the data volumes it has processed. The following metrics are for spans, although similar metrics are exposed for metrics and logs signals:

otelcol_receiver_accepted_spans

The number of spans successfully pushed into the pipeline.

otelcol_receiver_refused_spans

The number of spans that could not be pushed into the pipeline.

otelcol_exporter_sent_spans

The number of spans successfully sent to the destination.

otelcol_exporter_enqueue_failed_spans

The number of spans failed to be added to the sending queue.

The Operator creates a **<cr_name>-collector-monitoring** telemetry service that you can use to scrape the metrics endpoint.

Procedure

1. Enable the telemetry service by adding the following lines in the **OpenTelemetryCollector** custom resource:

```
config: |
  service:
    telemetry:
```

```
metrics:  
  address: ":8888" 1
```

- 1** The address at which the internal collector metrics are exposed. Defaults to **:8888**.

1. Retrieve the metrics by running the following command, which uses the port-forwarding Collector pod:

```
$ oc port-forward <collector_pod>
```

2. Access the metrics endpoint at **http://localhost:8888/metrics**.

10.3. DEBUG EXPORTER

You can configure the debug exporter to export the collected data to the standard output.

Procedure

1. Configure the **OpenTelemetryCollector** custom resource as follows:

```
config: |  
  exporters:  
    debug:  
      verbosity: detailed  
  service:  
    pipelines:  
      traces:  
        exporters: [debug]  
      metrics:  
        exporters: [debug]  
      logs:  
        exporters: [debug]
```

2. Use the **oc logs** command or the web console to export the logs to the standard output.

CHAPTER 11. MIGRATING



IMPORTANT

The Red Hat OpenShift distributed tracing platform (Jaeger) is a deprecated feature. Deprecated functionality is still included in OpenShift Container Platform and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments.

For the most recent list of major functionality that has been deprecated or removed within OpenShift Container Platform, refer to the *Deprecated and removed features* section of the OpenShift Container Platform release notes.

If you are already using the Red Hat OpenShift distributed tracing platform (Jaeger) for your applications, you can migrate to the Red Hat build of OpenTelemetry, which is based on the [OpenTelemetry](#) open-source project.

The Red Hat build of OpenTelemetry provides a set of APIs, libraries, agents, and instrumentation to facilitate observability in distributed systems. The OpenTelemetry Collector in the Red Hat build of OpenTelemetry can ingest the Jaeger protocol, so you do not need to change the SDKs in your applications.

Migration from the distributed tracing platform (Jaeger) to the Red Hat build of OpenTelemetry requires configuring the OpenTelemetry Collector and your applications to report traces seamlessly. You can migrate sidecar and sidecarless deployments.

11.1. MIGRATING WITH SIDECARS

The Red Hat build of OpenTelemetry Operator supports sidecar injection into deployment workloads, so you can migrate from a distributed tracing platform (Jaeger) sidecar to a Red Hat build of OpenTelemetry sidecar.

Prerequisites

- The Red Hat OpenShift distributed tracing platform (Jaeger) is used on the cluster.
- The Red Hat build of OpenTelemetry is installed.

Procedure

1. Configure the OpenTelemetry Collector as a sidecar.

```
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: <otel-collector-namespace>
spec:
  mode: sidecar
  config: |
    receivers:
      jaeger:
        protocols:
          grpc: {}
```

```

    thrift_binary: {}
    thrift_compact: {}
    thrift_http: {}
processors:
  batch: {}
  memory_limiter:
    check_interval: 1s
    limit_percentage: 50
    spike_limit_percentage: 30
  resourcedetection:
    detectors: [openshift]
    timeout: 2s
exporters:
  otlp:
    endpoint: "tempo-<example>-gateway:8090" ❶
    tls:
      insecure: true
service:
  pipelines:
    traces:
      receivers: [jaeger]
      processors: [memory_limiter, resourcedetection, batch]
      exporters: [otlp]

```

❶ This endpoint points to the Gateway of a TempoStack instance deployed by using the **<example>** Tempo Operator.

2. Create a service account for running your application.

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-collector-sidecar

```

3. Create a cluster role for the permissions needed by some processors.

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-collector-sidecar
rules:
  ❶
- apiGroups: ["config.openshift.io"]
  resources: ["infrastructures", "infrastructures/status"]
  verbs: ["get", "watch", "list"]

```

❶ The **resourcedetectionprocessor** requires permissions for infrastructures and infrastructures/status.

4. Create a **ClusterRoleBinding** to set the permissions for the service account.

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding

```

```

metadata:
  name: otel-collector-sidecar
subjects:
- kind: ServiceAccount
  name: otel-collector-deployment
  namespace: otel-collector-example
roleRef:
  kind: ClusterRole
  name: otel-collector
  apiGroup: rbac.authorization.k8s.io

```

5. Deploy the OpenTelemetry Collector as a sidecar.
6. Remove the injected Jaeger Agent from your application by removing the **"sidecar.jaegertracing.io/inject": "true"** annotation from your **Deployment** object.
7. Enable automatic injection of the OpenTelemetry sidecar by adding the **sidecar.opentelemetry.io/inject: "true"** annotation to the **.spec.template.metadata.annotations** field of your **Deployment** object.
8. Use the created service account for the deployment of your application to allow the processors to get the correct information and add it to your traces.

11.2. MIGRATING WITHOUT SIDECARS

You can migrate from the distributed tracing platform (Jaeger) to the Red Hat build of OpenTelemetry without sidecar deployment.

Prerequisites

- The Red Hat OpenShift distributed tracing platform (Jaeger) is used on the cluster.
- The Red Hat build of OpenTelemetry is installed.

Procedure

1. Configure OpenTelemetry Collector deployment.
2. Create the project where the OpenTelemetry Collector will be deployed.

```

apiVersion: project.openshift.io/v1
kind: Project
metadata:
  name: observability

```

3. Create a service account for running the OpenTelemetry Collector instance.

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: otel-collector-deployment
  namespace: observability

```

4. Create a cluster role for setting the required permissions for the processors.


```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: otel-collector
rules:
  1
  2
- apiGroups: [ "", "config.openshift.io" ]
  resources: [ "pods", "namespaces", "infrastructures", "infrastructures/status" ]
  verbs: [ "get", "watch", "list" ]

```

- 1 Permissions for the **pods** and **namespaces** resources are required for the **k8sattributesprocessor**.
- 2 Permissions for **infrastructures** and **infrastructures/status** are required for **resourcedetectionprocessor**.


5. Create a ClusterRoleBinding to set the permissions for the service account.

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: otel-collector
subjects:
- kind: ServiceAccount
  name: otel-collector-deployment
  namespace: observability
roleRef:
  kind: ClusterRole
  name: otel-collector
  apiGroup: rbac.authorization.k8s.io

```

6. Create the OpenTelemetry Collector instance.



NOTE

This collector will export traces to a TempoStack instance. You must create your TempoStack instance by using the Red Hat Tempo Operator and place here the correct endpoint.

```

apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
  namespace: observability
spec:
  mode: deployment
  serviceAccount: otel-collector-deployment
  config: |
    receivers:
      jaeger:
        protocols:
          grpc: {}

```

```
    thrift_binary: {}
    thrift_compact: {}
    thrift_http: {}
processors:
  batch: {}
  k8sattributes:
  memory_limiter:
    check_interval: 1s
    limit_percentage: 50
    spike_limit_percentage: 30
  resourcedetection:
    detectors: [openshift]
exporters:
  otlp:
    endpoint: "tempo-example-gateway:8090"
    tls:
      insecure: true
service:
  pipelines:
    traces:
      receivers: [jaeger]
      processors: [memory_limiter, k8sattributes, resourcedetection, batch]
      exporters: [otlp]
```

7. Point your tracing endpoint to the OpenTelemetry Operator.
8. If you are exporting your traces directly from your application to Jaeger, change the API endpoint from the Jaeger endpoint to the OpenTelemetry Collector endpoint.

Example of exporting traces by using the `jaegerexporter` with Golang

```
exp, err := jaeger.New(jaeger.WithCollectorEndpoint(jaeger.WithEndpoint(url))) 1
```

- 1** The URL points to the OpenTelemetry Collector API endpoint.

CHAPTER 12. UPGRADING

For version upgrades, the Red Hat build of OpenTelemetry Operator uses the Operator Lifecycle Manager (OLM), which controls installation, upgrade, and role-based access control (RBAC) of Operators in a cluster.

The OLM runs in the OpenShift Container Platform by default. The OLM queries for available Operators as well as upgrades for installed Operators.

When the Red Hat build of OpenTelemetry Operator is upgraded to the new version, it scans for running OpenTelemetry Collector instances that it manages and upgrades them to the version corresponding to the Operator's new version.

12.1. ADDITIONAL RESOURCES

- [Operator Lifecycle Manager concepts and resources](#)
- [Updating installed Operators](#)

CHAPTER 13. REMOVING

The steps for removing the Red Hat build of OpenTelemetry from an OpenShift Container Platform cluster are as follows:

1. Shut down all Red Hat build of OpenTelemetry pods.
2. Remove any OpenTelemetryCollector instances.
3. Remove the Red Hat build of OpenTelemetry Operator.

13.1. REMOVING AN OPENTELEMETRY COLLECTOR INSTANCE BY USING THE WEB CONSOLE


You can remove an OpenTelemetry Collector instance in the **Administrator** view of the web console.

Prerequisites

- You are logged in to the web console as a cluster administrator with the **cluster-admin** role.
- For Red Hat OpenShift Dedicated, you must be logged in using an account with the **dedicated-admin** role.

Procedure

1. Go to **Operators** → **Installed Operators** → **Red Hat build of OpenTelemetry Operator** → **OpenTelemetryInstrumentation** or **OpenTelemetryCollector**.

2. To remove the relevant instance, select  → **Delete ...** → **Delete**.

3. Optional: Remove the Red Hat build of OpenTelemetry Operator.

13.2. REMOVING AN OPENTELEMETRY COLLECTOR INSTANCE BY USING THE CLI

You can remove an OpenTelemetry Collector instance on the command line.

Prerequisites

- An active OpenShift CLI (**oc**) session by a cluster administrator with the **cluster-admin** role.

TIP

- Ensure that your OpenShift CLI (**oc**) version is up to date and matches your OpenShift Container Platform version.
- Run **oc login**:

```
$ oc login --username=<your_username>
```

Procedure

1. Get the name of the OpenTelemetry Collector instance by running the following command:

```
$ oc get deployments -n <project_of_opentelemetry_instance>
```

2. Remove the OpenTelemetry Collector instance by running the following command:

```
$ oc delete opentelemetrycollectors <opentelemetry_instance_name> -n  
<project_of_opentelemetry_instance>
```

3. Optional: Remove the Red Hat build of OpenTelemetry Operator.

Verification

- To verify successful removal of the OpenTelemetry Collector instance, run **oc get deployments** again:

```
$ oc get deployments -n <project_of_opentelemetry_instance>
```

13.3. ADDITIONAL RESOURCES

- [Deleting Operators from a cluster](#)
- [Getting started with the OpenShift CLI](#)