

Configuring SRM

Date published: 2019-09-13

Date modified: 2019-09-13

CLOUdera

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Configuring Clusters and Replications.....	4
Configuring the Driver Role Target Clusters.....	5
Configuring the Service Role Target Cluster.....	6
Configuring Properties Not Exposed in Cloudera Manager.....	6
New Topic and Consumer Group Discovery.....	7
Configuration Examples.....	7
Bidirectional Replication of Active Clusters.....	8
Cross Data Center Replication of Multiple Clusters.....	9
Configuration Properties Reference.....	13
Configuration Properties Reference for Properties Available in Cloudera Manager.....	14

Configuring Clusters and Replications

Learn how to add and specify clusters and replications .

About this task

In order for SRM to replicate data to or from clusters, you need to:

1. Specify cluster aliases
2. Specify cluster connection information
3. Add and enable replications

Specifying your clusters and enabling replications does not start replication of data itself. When a clusters and replications are added with the following method to the configuration, SRM will connect and set up communication with them, but will not automatically replicate any data. To start replicating data you need to specify which topics to replicate with the `srm-control` command line tool.

Use the following steps as reference when you want to add new clusters or replications to your deployment.

Procedure

1. In Cloudera Manager select Streams Replication Manager.
2. Go to Configuration.
3. Specify cluster aliases:
 - a) Find the Streams Replication Manager Cluster alias. property.
 - b) Add a comma delimited list of cluster aliases. For example:

```
primary, secondary
```

Cluster aliases are arbitrary names defined by the user. Aliases specified here are used in other configuration properties and with the `srm-control` tool to refer to the clusters added for replication.

4. Specify cluster connection information:
 - a) Find the `streams.replication.manager's` replication configs property.
 - b) Click the add button and add new lines for each cluster alias you have specified in the Streams Replication Manager Cluster alias. property
 - c) Add connection information for your clusters. For example:

```
primary.bootstrap.servers=primary_host1:9092,primary_host2:9092,primary_
host3:9092
secondary.bootstrap.servers=secondary_host1:9092,secondary_host2:9092
,secondary_host3:9092
```

Each cluster has to be added to a new line. If a cluster has multiple hosts, add them to the same line but delimit them with commas.

5. Add and enable replications:
 - a) Find the `streams.replication.manager's` replication configs property.
 - b) Click the add button and add new lines for each unique replication you want to add and enable.
 - c) Add and enable your replications. For example:

```
primary->secondary.enabled=true
secondary->primary.enabled=true
```

6. Enter a Reason for change, and then click Save Changes to commit the changes.
7. Restart Streams Replication Manager.

Results

Replicating data to or from the specified clusters is now possible.

What to do next

Use the `srm-control` tool to kick off replication by adding topics or groups to the allowlist.

Related Information

[srm-control](#)

Configuring the Driver Role Target Clusters

Learn about Streams Replication Manager Driver role target clusters and how to configure them.

About this task

The Streams Replication Manager Driver Role is responsible for connecting to the specified clusters and performing replication between them. The driver can be installed on one or more hosts within a cluster.

The clusters the driver connects to are the clusters that you specify with the Streams Replication Manager Cluster alias. and `streams.replication.manager's` replication configs properties.

Target clusters of the driver are clusters that the driver writes data to. By default when the driver is started it will connect to all clusters, gather data from them, and write to all of them. In other words, by default a driver targets all clusters in your configuration. However, in a more complex deployment this setup is not ideal, as the driver role instances can become overburdened. In cases like these Cloudera recommends that you limit the number of clusters that each driver targets.

This can be done with the `streams.replication.manager.driver.target.cluster` property, which allows you to specify which cluster or clusters the driver targets.

When you specify a driver target, the driver still connects to all clusters and gathers data from them. However, it will only write to the clusters specified.

Configuring driver targets and limiting the number of clusters each instance of the driver writes to is essential in an advanced deployment as it allows you to distribute replication workloads.

By default the `streams.replication.manager.driver.target.cluster` property is left empty, meaning that all clusters are targeted. The property accepts any cluster alias that is specified in Streams Replication Manager Cluster alias.. When adding multiple cluster aliases, delimit them with a comma.

Procedure

1. In Cloudera Manager select Streams Replication Manager.
2. Go to Configuration.
3. Find the `streams.replication.manager.driver.target.cluster` property.
4. Add the cluster aliases that you want the driver role to target. For example:

```
primary, secondary
```

5. Enter a Reason for change, and then click Save Changes to commit the changes.
6. Restart Streams Replication Manager.

Results

Driver targets are configured. Drivers only write data to the targeted clusters.

Configuring the Service Role Target Cluster

Learn about the Streams Replication Manager Service role target cluster and how to configure it.

About this task

The Streams Replication Manager Service role consists of a REST API and a Kafka Streams application that aggregates and exposes cluster, topic, and consumer group metrics. With the help of these metrics, users can monitor replications. The service can only be installed on one host per cluster.

Each instance of the service is associated with a single target cluster. The target is the cluster that the service gathers and exposes metrics from. Because each instance of the service can only target and expose metrics from a single cluster, monitoring multiple clusters requires the deployment of multiple instances of the service.

The target cluster of the service is configured with the streams replication manager service target cluster property. The property accepts any cluster alias that is specified in Streams Replication Manager Cluster alias. as long as data is being replicated to that cluster. Configuring a service target is mandatory.

Procedure

1. In Cloudera Manager select Streams Replication Manager.
2. Go to Configuration.
3. Find the streams replication manager service target cluster property.
4. Add the target cluster alias. For example:

```
secondary
```

5. Enter a Reason for change, and then click Save Changes to commit the changes.
6. Restart Streams Replication Manager.

Results

The service target is set. The service gathers and exposes metrics from the specified cluster.

Configuring Properties Not Exposed in Cloudera Manager

Learn how to configure configuration properties not exposed in Cloudera Manager.

About this task

In addition to the configuration properties exposed directly for configuration through Cloudera Manager, SRM accepts a number of additional SRM specific properties as well as all Kafka client properties available in the version of Kafka that you are using. Properties not exposed directly can be set through the streams.replication.manager's replication configs property. For a comprehensive list of SRM specific properties, see the configuration properties reference. For a comprehensive list of Kafka client properties, see the upstream Apache Kafka documentation.

Configuration properties for SRM can be set on three different levels, which are the following:

- Top level: Top level or global configuration is achieved by adding the property on its own. For example:

```
replication.factor=3
```

- Cluster level: Cluster level configuration can be achieved by prepending the configuration property with a cluster alias specified in Streams Replication Manager Cluster alias.. For example:

```
primary.replication.factor=3
```

- Replication level: Replication level configuration can be achieved by prepending the configuration property with the name of the replication. For example:

```
primary->secondary.replication.factor=3
```

Before you begin

Make sure that cluster aliases and replications are configured. Otherwise cluster or replication level configuration is not possible.

Procedure

1. In Cloudera Manager select Streams Replication Manager.
2. Go to Configuration.
3. Configure properties not exposed in Cloudera Manager:
 - a) Find the streams.replication.manager's replication configs property.
 - b) Click the add button and add new lines for each additional property you want to configure.
 - c) Add configuration properties. For example:

```
replication.factor=3
```

4. Enter a Reason for change, and then click Save Changes to commit the changes.
5. Restart Streams Replication Manager.

Results

Configuration properties not directly exposed in Cloudera Manager are configured.

Related Information

[Configuration Properties Reference](#)

[Configuration Properties Reference for Properties Available in Cloudera Manager](#)

New Topic and Consumer Group Discovery

Understand how SRM discovers new topics and consumer groups on source clusters.

The discovery and replication of newly created topics or consumer groups is not instantaneous. SRM checks source clusters for new topics and consumer groups periodically, as controlled by the `refresh.topics.interval.seconds` and `refresh.groups.interval.seconds` properties. By default both properties are set to 10 minutes. As a result, the discovery and replication of new topics or groups can take up to 10 minutes.

Cloudera does not recommend using a refresh interval lower than the default value for production environments as it can lead to severe performance degradation.

Configuration Examples

Configuration examples for typical SRM architectures.


Bidirectional Replication of Active Clusters

Configuration example for two active Kafka clusters setup with bidirectional replication.

About this task

A typical scenario involves two active Kafka clusters within the same region but in separate availability zones. Clients can connect to either cluster in case one is temporarily unavailable. This example demonstrates the steps required to set up a deployment with two clusters configured with bidirectional replication. Additionally, it also provides example commands to start replication between clusters.

Figure 1: Bidirectional Replication of Active Clusters



The steps shown here have to be carried out on all clusters in a given deployment. Configuration properties presented in Steps 3-5 are configured identically on all clusters. The configuration property presented in Step 7 will differ for each cluster.



Note: The following list of steps assume that the Streams Replication Manager Service role is running on 1 host on each cluster and is targeting the cluster it is running on.

Procedure

1. In Cloudera Manager select Streams Replication Manager.
2. Go to Configuration.
3. Specify cluster aliases:
 - a) Find the Streams Replication Manager Cluster alias. property.
 - b) Add a comma delimited list of cluster aliases. For example:

```
primary, secondary
```

Cluster aliases are arbitrary names defined by the user. Aliases specified here are used in other configuration properties and with the `srm-control` tool to refer to the clusters added for replication.

4. Specify cluster connection information:

- a) Find the streams.replication.manager's replication configs property.
- b) Click the add button and add new lines for each cluster alias you have specified in the Streams Replication Manager Cluster alias. property
- c) Add connection information for your clusters. For example:

```
primary.bootstrap.servers=primary_host1:9092,primary_host2:9092,primary_
host3:9092
secondary.bootstrap.servers=secondary_host1:9092,secondary_host2:9092
,secondary_host3:9092
```

Each cluster has to be added to a new line. If a cluster has multiple hosts, add them to the same line but delimit them with commas.

5. Add and enable replications:

- a) Find the streams.replication.manager's replication configs property.
- b) Click the add button and add new lines for each unique replication you want to add and enable.
- c) Add and enable your replications. For example:

```
primary->secondary.enabled=true
secondary->primary.enabled=true
```

6. Enter a Reason for change, and then click Save Changes to commit the changes.**7. Add Streams Replication Manager Driver role instances to all Kafka broker hosts:**

- a) Go to Instances.
- b) Click Add Role Instances.
- c) Click Select Hosts.
- d) Select all Kafka broker hosts and click Ok.
- e) Click Continue.
- f) Find the streams replication manager driver target cluster property.
- g) Add the cluster aliases that you want the driver role to target. For example:

- On the primary cluster:

```
primary
```

- On the secondary cluster:

```
secondary
```

The streams replication manager driver target cluster property allows you to specify which clusters the driver should write to. In this example, the drivers read data from all clusters, but only write to the cluster they are running on. This allows you to distribute replication workloads.

- h) Click Continue.

8. Restart Streams Replication Manager.**9. Replicate data between clusters with the following commands:**

```
srm-control topics --source primary --target secondary --add ".*"
```

```
srm-control topics --source secondary --target primary --add ".*"
```

Cross Data Center Replication of Multiple Clusters


Configuration example with three data centers that have two Kafka clusters each.

About this task

In more advanced deployments, you may have multiple Kafka clusters in each of several data centers. To prevent creating a fully-connected mesh of all Kafka clusters, Cloudera recommends leveraging a single Kafka cluster in each data center for cross data center replication.

This example demonstrates the steps required to configure a deployment with three data centers that each have two Kafka clusters. Additionally, it also provides example commands to start bidirectional replication of all topics within each data center and an example on how to replicate a single topic across all data centers.

Figure 2: Cross Data Center Replication of Multiple Clusters



The steps shown here have to be carried out on all clusters in a given deployment. Configuration properties presented in Steps 3-5 are configured identically on all clusters. The configuration property presented in Step 7 will differ for each cluster.



Note: The following list of steps assume that the Streams Replication Manager Service role is running on 1 host on each cluster and is targeting the cluster it is running on.

Procedure

1. In Cloudera Manager select Streams Replication Manager.
2. Go to Configuration.
3. Specify cluster aliases:
 - a) Find the Streams Replication Manager Cluster alias. property.
 - b) Add a comma delimited list of cluster aliases. For example:

```
west1, west2, east1, east2, south1, south2
```

Cluster aliases are arbitrary names defined by the user. Aliases specified here are used in other configuration properties and with the `srn-control` tool to refer to the clusters added for replication.

4. Specify cluster connection information:
 - a) Find the `streams.replication.manager's` replication configs property.
 - b) Click the add button and add new lines for each cluster alias you have specified in the Streams Replication Manager Cluster alias. property
 - c) Add connection information for your clusters. For example:

```
west1.bootstrap.servers=west1_host1:9092,west1_host2:9092,west1_host3:9092
west2.bootstrap.servers=west2_host1:9092,west2_host2:9092,west2_host3:9092
```

```
east1.bootstrap.servers=east1_host1:9092, east1_host2:9092, east1_host3:
9092
east2.bootstrap.servers=east2_host1:9092, east2_host2:9092, east2_host3:90
92

south1.bootstrap.servers=south1_host1:9092, south1_host2:9092, south1_host3
:9092
south2.bootstrap.servers=south2_host1:9092, south2_host2:9092, south2_hos
t3:9092
```

Each cluster has to be added to a new line. If a cluster has multiple hosts, add them to the same line but delimit them with commas.

5. Add and enable replications:

- a) Find the streams.replication.manager's replication configs property.
- b) Click the add button and add new lines for each unique replication you want to add and enable.
- c) Add and enable your replications. For example:

Enable cross data center replication by adding the following replications:

```
west1->east1.enabled=true
west1->south1.enabled=true
east1->west1.enabled=true
east1->south1.enabled=true
south1->west1.enabled=true
south1->east1.enabled=true
```

Enable bidirectional replication within each data center by adding the following replications:

```
west1->west2.enabled=true
west2->west1.enabled=true
east1->east2.enabled=true
east2->east1.enabled=true
south1->south2.enabled=true
south2->south1.enabled=true
```

6. Enter a Reason for change, and then click Save Changes to commit the changes.

7. Add Streams Replication Manager Driver role instances to all Kafka broker hosts:

- a) Go to Instances.
- b) Click Add Role Instances.
- c) Click Select Hosts.
- d) Select all Kafka broker hosts and click Ok.
- e) Click Continue.
- f) Find the streams replication manager driver target cluster property.
- g) Add the cluster aliases that you want the driver role to target. For example:

- In the west data center:

```
west1, west2
```

- In the east data center:

```
east1, east2
```

- In the south data center:

```
south1, south2
```

The streams replication manager driver target cluster property allows you to specify which clusters the driver should write to. In this example, the drivers read data from all clusters, but only write to the cluster they are running on. This allows you to distribute replication workloads.

- h) Click Continue.

8. Restart Streams Replication Manager.**9. Replicate topics between hosts within each data center:**

```
srm-control topics --source west1 --target west2 --add ".*"
```

```
srm-control topics --source west2 --target west1 --add ".*"
```

```
srm-control topics --source east1 --target east2 --add ".*"
```

```
srm-control topics --source east2 --target east1 --add ".*"
```

```
srm-control topics --source south1 --target south2 --add ".*"
```

```
srm-control topics --source south2 --target south1 --add ".*"
```

10. Replicate topic1 across all data centers:

```
srm-control topics --source west1 --target east1 --add topic1,west2.topic1
```

```
srm-control topics --source west1 --target south1 --add topic1,west2.topic1
```

```
srm-control topics --source east1 --target west1 --add topic1,east2.topic1
```

```
srm-control topics --source east1 --target south1 --add topic1,east2.topic1
```

```
srm-control topics --source south1 --target west1 --add topic1,south2.topic1
```

```
srm-control topics --source south1 --target east1 --add topic1,south2.topic1
```

Configuration Properties Reference

A collection of SRM configuration properties not available in Cloudera Manager

The following table lists all SRM specific configuration properties that are not available directly for configuration via Cloudera Manager.

Property	Default Value	Description
sync.topic.configs.enabled	true	Enables the monitoring of the source cluster for configuration changes.
emit.heartbeats.enabled	true	Enables periodic emission of heartbeats.
emit.heartbeats.interval.seconds	5 (seconds)	The interval at which SRM emits heartbeats.
emit.checkpoints.enabled	true	Enables periodic emission of consumer offset information.
refresh.topics.enabled	true	Enables a periodical check for new topics on source clusters.
refresh.groups.enabled	true	Enables a periodical check for new consumer groups on source clusters.
replication.policy.class	org.apache.kafka.connect.mirror.DefaultReplicationPolicy	Replication policies to use. Use LegacyReplicationPolicy to mimic legacy MirrorMaker behavior.
replication.factor	2	Replication factor used for remote topics.
heartbeats.topic.replication.factor	3	Replication factor used for internal heartbeat topics.
checkpoints.topic.replication.factor	3	Replication factor used for internal checkpoints topics.
offset-syncs.topic.replication.factor	3	Replication factor used for internal offset-syncs topics.

Configuration Properties Reference for Properties Available in Cloudera Manager

Service-Wide

Table 1: Advanced

Display Name	Description	Related Name	Default Value	API Name	Required
System Group	The group that this service's processes should run as.		streamsrepmgr	process_groupname	true
System User	The user that this service's processes should run as.		streamsrepmgr	process_username	true
Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/globals.xml	For advanced use only, a string to be inserted into staging/globals.xml. Applies to configurations of all roles in this service except client configuration.			staging/globals.xml_service_safety_valve	false
Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml	For advanced use only, a string to be inserted into staging/srm-common.xml. Applies to configurations of all roles in this service except client configuration.			staging/srm-common.xml_service_safety_valve	false
Streams Replication Manager Service Environment Advanced Configuration Snippet (Safety Valve)	For advanced use only, key-value pairs (one on each line) to be inserted into a role's environment. Applies to configurations of all roles in this service except client configuration.			STREAMS_REPLICATION_MANAGER_service_env_safety_valve	false

Table 2: Monitoring

Display Name	Description	Related Name	Default Value	API Name	Required
Enable Service Level Health Alerts	When set, Cloudera Manager will send alerts when the health of this service reaches the threshold specified by the EventServer setting eventserver_health_events_alert_threshold		true	enable_alerts	false
Enable Configuration Change Alerts	When set, Cloudera Manager will send alerts when this entity's configuration changes.		false	enable_config_alerts	false

Display Name	Description	Related Name	Default Value	API Name	Required
Service Triggers	<p>The configured triggers for this service. This is a JSON-formatted list of triggers. These triggers are evaluated as part of the health system. Every trigger expression is parsed, and if the trigger condition is met, the list of actions provided in the trigger expression is executed. Each trigger has the following fields:</p> <ul style="list-style-type: none"> triggerName (mandatory) - The name of the trigger. This value must be unique for the specific service. triggerExpression (mandatory) - A tsquery expression representing the trigger. streamThreshold (optional) - The maximum number of streams that can satisfy a condition of a trigger before the condition fires. By default set to 0, and any stream returned causes the condition to fire. enabled (optional) - By default set to 'true'. If set to 'false', the trigger is not evaluated. expressionEditorConfig (optional) - Metadata for the trigger editor. If present, the trigger should only be edited from the Edit Trigger page; editing the trigger here can lead to inconsistencies. <p>For example, the following JSON formatted trigger fires if there are more than 10 DataNodes with more than 500 file descriptors opened: <code>{ "triggerName": "sample-trigger", "triggerExpression": "IF (SELECT fd_open WHERE roleType = DataNode and last(fd_open) > 500) DO health:bad", "streamThreshold": 10, "enabled": "true" }</code> See the trigger rules documentation for more details on how to write triggers using tsquery. The JSON format is evolving and may change and, as a result, backward compatibility is not guaranteed between releases.</p>		[]	service_triggers	true
Service Monitor Derived Configs Advanced Configuration Snippet (Safety Valve)	For advanced use only, a list of derived configuration properties that will be used by the Service Monitor instead of the default ones.			smon_derived_configs_safety_valve	false
Streams Replication Manager Driver Role Health Test	When computing the overall STREAMS_REPLICATION_MANAGER health, consider Streams Replication Manager Driver's health		true	STREAMS_REPLICATION_MANAGER_STREAMS_REPLICATION_MANAGER_DRIVER_health_enabled	false

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Service Role Health Test	When computing the overall STREAMS_REPLICATION_MANAGER health, consider Streams Replication Manager Service's health		true	STREAMS_REPLICATION_MANAGER_STREAMS_REPLICATION_MANAGER_SERVICE_health_enabled	false

Table 3: Other

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Cluster alias.	Specifies cluster aliases for the clusters SRM is connecting to. Cluster aliases are comma delimited. For example, 'primary, backup'.	clusters		clusters	true
emit checkpoints interval seconds	The interval at which SRM emits checkpoint information.	emit.checkpoints.interval.seconds	5 second(s)	emit.checkpoints.interval.seconds	true
KAFKA Service	Name of the KAFKA service that this Streams Replication Manager service instance depends on			kafka_service	false
Enable Kerberos Authentication	Enables Kerberos authentication for this Streams Replication Manager.	kerberos.auth.enabled	false	kerberos.auth.enabled	false
refresh groups interval seconds	The interval at which SRM looks for new consumer groups on source clusters.	refresh.groups.interval.seconds	10 minute(s)	refresh.groups.interval.seconds	true
refresh topics interval seconds	The interval at which SRM looks for new topics on source clusters.	refresh.topics.interval.seconds	10 minute(s)	refresh.topics.interval.seconds	true
SRM_HEAP_OPTS	Memory heap params.	SRM_HEAP_OPTS	-Xmx8G -Xms1G	SRM_HEAP_OPTS	false
SRM_JMX_OPTS	Change parameters to setup jmxremote.	SRM_JMX_OPTS	-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false	SRM_JMX_OPTS	false
SRM_JVM_PERF_OPTS	SRM JVM perf and gc opts.	SRM_JVM_PERF_OPTS	-server -XX: +UseG1GC -XX:MaxGCPauseMillis=20 -XX:InitiatingHeapOccupancyPercent=35 -XX: +ExplicitGCInvokesConcurrent -Djava.awt.headless=true	SRM_JVM_PERF_OPTS	false

Display Name	Description	Related Name	Default Value	API Name	Required
streams.replication.manager.replication configs	Specifies cluster bootstrap server information and cluster replication pairs. Bootstrap information for each alias defined in the clusters property has to be added as follows: 'primary.bootstrap.servers=mycluster1.example.com:9092,primary.bootstrap.servers=mycluster2.example.com:9092'.. Each cluster has to be added in a new line, if cluster has multiple nodes, add all nodes to the same line and delimit each with a comma. Cluster replication pairs can be defined as follows: 'primary->backup.enabled=true'.. Each unique replication pair has to be added in a new line.	streams.replication.manager.config		streams.replication.manager.config	true
Streams Replication Manager Configuration Directory	Directory for additional libs, jars, and other resources.	streams.replication.manager.working.directory	/var/lib/streamsrepmgr	streams.replication.manager.working.directory	false
sync topic acls enabled	Enables the monitoring of the source cluster for ACL changes.	sync.topic.acls.enabled	false	sync.topic.acls.enabled	true
sync topic acls interval seconds	The interval at which SRM checks the source cluster for ACL changes.	sync.topic.acls.interval.seconds	30 second(s)	sync.topic.acls.interval.seconds	true
sync topic configs interval seconds	The interval at which SRM checks the source cluster for configuration changes.	sync.topic.configs.interval.seconds	10 minute(s)	sync.topic.configs.interval.seconds	true
tasks.max	Maximum number of tasks for replication between clusters.	tasks.max	3	tasks.max	true

Table 4: Suppressions

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Parameter Validation: Streams Replication Manager Cluster alias.	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Cluster alias. parameter.		false	service_config_suppression_clusters	true
Suppress Parameter Validation: System Group	Whether to suppress configuration warnings produced by the built-in parameter validation for the System Group parameter.		false	service_config_suppression_process_groupname	true
Suppress Parameter Validation: System User	Whether to suppress configuration warnings produced by the built-in parameter validation for the System User parameter.		false	service_config_suppression_process_username	true
Suppress Parameter Validation: Service Triggers	Whether to suppress configuration warnings produced by the built-in parameter validation for the Service Triggers parameter.		false	service_config_suppression_service_triggers	true
Suppress Parameter Validation: Service Monitor Derived Configs Advanced Configuration Snippet (Safety Valve)	Whether to suppress configuration warnings produced by the built-in parameter validation for the Service Monitor Derived Configs Advanced Configuration Snippet (Safety Valve) parameter.		false	service_config_suppression_smon_derived_configs_safety_valve	true

Display Name	Description	Related Name	Default Value	API Name	Require
Suppress Parameter Validation: SRM_HEAP_OPTS	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM_HEAP_OPTS parameter.		false	service_config_suppressi on_srm_heap_opts	true
Suppress Parameter Validation: SRM_JMX_OPTS	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM_JMX_OPTS parameter.		false	service_config_suppressi on_srm_jmx_opts	true
Suppress Parameter Validation: SRM_JVM_PERF_OPTS	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM_JVM_PERF_OPTS parameter.		false	service_config_suppressi on_srm_jvm_perf_opts	true
Suppress Parameter Validation: Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/globals.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/globals.xml parameter.		false	service_config_suppr ession_staging/globals.x ml_service_safety_valve	true
Suppress Parameter Validation: Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml parameter.		false	service_config_suppressi on_staging/srm-commo n.xml_service_safety _valve	true
Suppress Parameter Validation: streams.replication.manager.replication configs	Whether to suppress configuration warnings produced by the built-in parameter validation for the streams.replication.manager's replication configs parameter.		false	service_config_suppressi on_streams.replication.m anager.config	true
Suppress Parameter Validation: Streams Replication Manager Configuration Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Configuration Directory parameter.		false	service_config_suppressi on_streams.replication.m anager.working.directory	true
Suppress Configuration Validator: Streams Replication Manager Driver Count Validator	Whether to suppress configuration warnings produced by the Streams Replication Manager Driver Count Validator configuration validator.		false	service_config_suppr ession_streams_replicati on_manager_driver_co unt_validator	true
Suppress Configuration Validator: Streams Replication Manager Service Count Validator	Whether to suppress configuration warnings produced by the Streams Replication Manager Service Count Validator configuration validator.		false	service_config_suppr ession_streams_replicati on_manager_service_c ount_validator	true
Suppress Parameter Validation: Streams Replication Manager Service Environment Advanced Configuration Snippet (Safety Valve)	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Environment Advanced Configuration Snippet (Safety Valve) parameter.		false	service_config_suppr ession_streams_replicati on_manager_service_e nv_safety_valve	true

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Health Test: Streams Replication Manager Driver Health	Whether to suppress the results of the Streams Replication Manager Driver Health health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	service_health_suppression_streams_replication_manager_streams_replication_manager_driver_health	true
Suppress Health Test: Streams Replication Manager Service Health	Whether to suppress the results of the Streams Replication Manager Service Health health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	service_health_suppression_streams_replication_manager_streams_replication_manager_service_health	true

Streams Replication Manager Driver

Table 5: Advanced

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Driver XML Override	For advanced use only, replace entire XML in the logback configuration file for Streams Replication Manager Driver, ignoring all logging configuration.	logback_safety_valve		logback_safety_valve	false
Heap Dump Directory	Path to directory where heap dumps are generated when <code>java.lang.OutOfMemoryError</code> error is thrown. This directory is automatically created if it does not exist. If this directory already exists, role user must have write access to this directory. If this directory is shared among multiple roles, it should have 1777 permissions. The heap dump files are created with 600 permissions and are owned by the role user. The amount of free space in this directory should be greater than the maximum Java Process heap size configured for this role.	oom_heap_dump_dir	/tmp	oom_heap_dump_dir	false
Dump Heap When Out of Memory	When set, generates heap dump file when <code>java.lang.OutOfMemoryError</code> is thrown.		true	oom_heap_dump_enabled	true
Kill When Out of Memory	When set, a SIGKILL signal is sent to the role process when <code>java.lang.OutOfMemoryError</code> is thrown.		true	oom_sigkill_enabled	true
Automatically Restart Process	When set, this role's process is automatically (and transparently) restarted in the event of an unexpected failure.		false	process_auto_restart	true

Display Name	Description	Related Name	Default Value	API Name	Required
Enable Metric Collection	Cloudera Manager agent monitors each service and each of its role by publishing metrics to the Cloudera Manager Service Monitor. Setting it to false will stop Cloudera Manager agent from publishing any metric for corresponding service/roles. This is usually helpful for services that generate large amount of metrics which Service Monitor is not able to process.		true	process_should_monitor	true
Streams Replication Manager Driver Advanced Configuration Snippet (Safety Valve) for staging/globals.xml	For advanced use only. A string to be inserted into staging/globals.xml for this role only.			staging/globals.xml_role_safety_valve	false
Streams Replication Manager Driver Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml	For advanced use only. A string to be inserted into staging/srm-common.xml for this role only.			staging/srm-common.xml_role_safety_valve	false
Streams Replication Manager Driver Environment Advanced Configuration Snippet (Safety Valve)	For advanced use only, key-value pairs (one on each line) to be inserted into a role's environment. Applies to configurations of this role except client configuration.			STREAMS_REPLICATION_MANAGER_DRIVER_role_env_safety_valve	false

Table 6: Logs

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Driver Log Directory	The log directory for log files of the role Streams Replication Manager Driver.	log_dir	/var/log/streams-replication-manager	log_dir	false
Streams Replication Manager Driver Logging Threshold	The minimum log level for Streams Replication Manager Driver logs		INFO	log_threshold	false
Streams Replication Manager Driver Maximum Log File Backups	The maximum number of rolled log files to keep for Streams Replication Manager Driver logs. Typically used by log4j or logback.		10	max_log_backup_index	false
Streams Replication Manager Driver Max Log Size	The maximum size, in megabytes, per log file for Streams Replication Manager Driver logs. Typically used by log4j or logback.		200 MiB	max_log_size	false

Table 7: Monitoring

Display Name	Description	Related Name	Default Value	API Name	Required
Enable Health Alerts for this Role	When set, Cloudera Manager will send alerts when the health of this role reaches the threshold specified by the EventServer setting eventserver_health_events_alert_threshold		true	enable_alerts	false

Display Name	Description	Related Name	Default Value	API Name	Required
Enable Configuration Change Alerts	When set, Cloudera Manager will send alerts when this entity's configuration changes.		false	enable_config_alerts	false
Log Directory Free Space Monitoring Absolute Thresholds	The health test thresholds for monitoring of free space on the filesystem that contains this role's log directory.		Warning: 10 GiB, Critical: 5 GiB	log_directory_free_space_absolute_thresholds	false
Log Directory Free Space Monitoring Percentage Thresholds	The health test thresholds for monitoring of free space on the filesystem that contains this role's log directory. Specified as a percentage of the capacity on that filesystem. This setting is not used if a Log Directory Free Space Monitoring Absolute Thresholds setting is configured.		Warning: Never, Critical: Never	log_directory_free_space_percentage_thresholds	false
Process Swap Memory Thresholds	The health test thresholds on the swap memory usage of the process. This takes precedence over the host level threshold.		Warning: 200 B, Critical: Never	process_swap_memory_thresholds	false

Display Name	Description	Related Name	Default Value	API Name	Required
Role Triggers	<p>The configured triggers for this role. This is a JSON-formatted list of triggers. These triggers are evaluated as part of the health system. Every trigger expression is parsed, and if the trigger condition is met, the list of actions provided in the trigger expression is executed. Each trigger has the following fields:</p> <ul style="list-style-type: none"> triggerName (mandatory) - The name of the trigger. This value must be unique for the specific role. triggerExpression (mandatory) - A tsquery expression representing the trigger. streamThreshold (optional) - The maximum number of streams that can satisfy a condition of a trigger before the condition fires. By default set to 0, and any stream returned causes the condition to fire. enabled (optional) - By default set to 'true'. If set to 'false', the trigger is not evaluated. expressionEditorConfig (optional) - Metadata for the trigger editor. If present, the trigger should only be edited from the Edit Trigger page; editing the trigger here can lead to inconsistencies. <p>For example, the following JSON formatted trigger configured for a DataNode fires if the DataNode has more than 1500 file descriptors opened: [{"triggerName": "sample-trigger", "triggerExpression": "IF (SELECT fd_open WHERE roleName = \$ROLENAME and last(fd_open) > 1500) DO health:bad", "streamThreshold": 0, "enabled": "true"}] See the trigger rules documentation for more details on how to write triggers using tsquery. The JSON format is evolving and may change and, as a result, backward compatibility is not guaranteed between releases.</p>		[]	role_triggers	true
File Descriptor Monitoring Thresholds	The health test thresholds of the number of file descriptors used. Specified as a percentage of file descriptor limit.		Warning: 50.0 %, Critical: 70.0 %	streams_replication_manager_driver_fd_thresholds	false
Streams Replication Manager Driver Host Health Test	When computing the overall Streams Replication Manager Driver health, consider the host's health.		true	streams_replication_manager_driver_host_health_enabled	false

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Driver Process Health Test	Enables the health test that the Streams Replication Manager Driver's process state is consistent with the role configuration		true	streams_replication_manager_driver_scm_health_enabled	false
Unexpected Exits Thresholds	The health test thresholds for unexpected exits encountered within a recent period specified by the unexpected_exits_window configuration for the role.		Warning: Never, Critical: Any	unexpected_exits_thresholds	false
Unexpected Exits Monitoring Period	The period to review when computing unexpected exits.		5 minute(s)	unexpected_exits_window	false

Table 8: Other

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Configuration Directory	Directory with all configurations for SRM's daemons.	streams.replication.manager.conf.directory	\$CONF_DIR	streams.replication.manager.conf.directory	true
streams replication manager driver target cluster	Target cluster aliases for the srm-driver. When set, the srm-driver will collect data from all clusters, but will only write to the clusters specified here. Cluster aliases are comma separated.	streams.replication.manager.driver.target.cluster		streams.replication.manager.driver.target.cluster	false
Java Home Path Override	Java Home Path Override for Streams Replication Manager. If left empty, the java shipped with CM is used.	streams.replication.manager.jdk.home		streams.replication.manager.jdk.home	false

Table 9: Performance

Display Name	Description	Related Name	Default Value	API Name	Required
Maximum Process File Descriptors	If configured, overrides the process soft and hard rlimits (also called ulimits) for file descriptors to the configured value.			rlimit_fds	false

Table 10: Resource Management

Display Name	Description	Related Name	Default Value	API Name	Required
Cgroup CPU Shares	Number of CPU shares to assign to this role. The greater the number of shares, the larger the share of the host's CPUs that will be given to this role when the host experiences CPU contention. Must be between 2 and 262144. Defaults to 1024 for processes not managed by Cloudera Manager.	cpu.shares	1024	rm_cpu_shares	true
Cgroup I/O Weight	Weight for the read I/O requests issued by this role. The greater the weight, the higher the priority of the requests when the host experiences I/O contention. Must be between 100 and 1000. Defaults to 1000 for processes not managed by Cloudera Manager.	blkio.weight	500	rm_io_weight	true

Display Name	Description	Related Name	Default Value	API Name	Required
Cgroup Memory Hard Limit	Hard memory limit to assign to this role, enforced by the Linux kernel. When the limit is reached, the kernel will reclaim pages charged to the process. If reclaiming fails, the kernel may kill the process. Both anonymous as well as page cache pages contribute to the limit. Use a value of -1 to specify no limit. By default processes not managed by Cloudera Manager will have no limit. If the value is -1, Cloudera Manager will not monitor Cgroup memory usage therefore some of the charts will show 'No Data'	memory.limit_in_bytes	-1 MiB	rm_memory_hard_limit	true
Cgroup Memory Soft Limit	Soft memory limit to assign to this role, enforced by the Linux kernel. When the limit is reached, the kernel will reclaim pages charged to the process if and only if the host is facing memory pressure. If reclaiming fails, the kernel may kill the process. Both anonymous as well as page cache pages contribute to the limit. Use a value of -1 to specify no limit. By default processes not managed by Cloudera Manager will have no limit. If the value is -1, Cloudera Manager will not monitor Cgroup memory usage therefore some of the charts will show 'No Data'	memory.soft_limit_in_bytes	-1 MiB	rm_memory_soft_limit	true

Table 11: Security

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Driver TLS/SSL Client Trust Store File	The location on disk of the trust store, in .jks format, used to confirm the authenticity of TLS/SSL servers that Streams Replication Manager Driver might connect to. This is used when Streams Replication Manager Driver is the client in a TLS/SSL connection. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.	ssl.truststore.location		ssl_client_truststore_location	false
Streams Replication Manager Driver TLS/SSL Client Trust Store Password	The password for the Streams Replication Manager Driver TLS/SSL Certificate Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file. The contents of trust stores are certificates, and certificates are public information.	ssl.truststore.password		ssl_client_truststore_password	false

Display Name	Description	Related Name	Default Value	API Name	Required
Enable TLS/SSL for Streams Replication Manager Driver	Encrypt communication between clients and Streams Replication Manager Driver using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)).	ssl_enabled	false	ssl_enabled	false
Streams Replication Manager Driver TLS/SSL Server JKS Keystore Key Password	The password that protects the private key contained in the JKS keystore used when Streams Replication Manager Driver is acting as a TLS/SSL server.	ssl.key.password		ssl_server_keystore_keypassword	false
Streams Replication Manager Driver TLS/SSL Server JKS Keystore File Location	The path to the TLS/SSL keystore file containing the server certificate and private key used for TLS/SSL. Used when Streams Replication Manager Driver is acting as a TLS/SSL server. The keystore must be in JKS format.	ssl.keystore.location		ssl_server_keystore_location	false
Streams Replication Manager Driver TLS/SSL Server JKS Keystore File Password	The password for the Streams Replication Manager Driver JKS keystore file.	ssl.keystore.password		ssl_server_keystore_password	false

Table 12: Stacks Collections

Display Name	Description	Related Name	Default Value	API Name	Required
Stacks Collection Data Retention	The amount of stacks data that is retained. After the retention limit is reached, the oldest data is deleted.	stacks_collection_data_retention	100 MiB	stacks_collection_data_retention	false
Stacks Collection Directory	The directory in which stacks logs are placed. If not set, stacks are logged into a stacks subdirectory of the role's log directory.	stacks_collection_directory		stacks_collection_directory	false
Stacks Collection Enabled	Whether or not periodic stacks collection is enabled.	stacks_collection_enabled	false	stacks_collection_enabled	true
Stacks Collection Frequency	The frequency with which stacks are collected.	stacks_collection_frequency	5.0 second(s)	stacks_collection_frequency	false
Stacks Collection Method	The method used to collect stacks. The jstack option involves periodically running the jstack command against the role's daemon process. The servlet method is available for those roles that have an HTTP server endpoint exposing the current stacks traces of all threads. When the servlet method is selected, that HTTP endpoint is periodically scraped.	stacks_collection_method	jstack	stacks_collection_method	false

Table 13: Suppressions

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Configuration Validator: CDH Version Validator	Whether to suppress configuration warnings produced by the CDH Version Validator configuration validator.		false	role_config_suppression_cdh_version_validator	true

Display Name	Description	Related Name	Default Value	API Name	Require
Suppress Parameter Validation: Streams Replication Manager Driver Log Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver Log Directory parameter.		false	role_config_suppress ion_log_dir	true
Suppress Parameter Validation: Streams Replication Manager Driver XML Override	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver XML Override parameter.		false	role_config_suppress ion_logback_safety_v alve	true
Suppress Parameter Validation: Heap Dump Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Heap Dump Directory parameter.		false	role_config_suppress ion_oom_heap_dum p_dir	true
Suppress Parameter Validation: Role Triggers	Whether to suppress configuration warnings produced by the built-in parameter validation for the Role Triggers parameter.		false	role_config_suppress ion_role_triggers	true
Suppress Parameter Validation: Streams Replication Manager Driver TLS/SSL Client Trust Store File	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver TLS/SSL Client Trust Store File parameter.		false	role_config_suppress ion_ssl_client_truststor e_location	true
Suppress Parameter Validation: Streams Replication Manager Driver TLS/SSL Client Trust Store Password	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver TLS/SSL Client Trust Store Password parameter.		false	role_config_suppress ion_ssl_client_truststor e_password	true
Suppress Parameter Validation: Streams Replication Manager Driver TLS/SSL Server JKS Keystore Key Password	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver TLS/SSL Server JKS Keystore Key Password parameter.		false	role_config_suppress ion_ssl_server_keystore_ keypassword	true
Suppress Parameter Validation: Streams Replication Manager Driver TLS/SSL Server JKS Keystore File Location	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver TLS/SSL Server JKS Keystore File Location parameter.		false	role_config_suppress ion_ssl_server_keystore_ location	true
Suppress Parameter Validation: Streams Replication Manager Driver TLS/SSL Server JKS Keystore File Password	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver TLS/SSL Server JKS Keystore File Password parameter.		false	role_config_suppress ion_ssl_server_keystore_ password	true
Suppress Parameter Validation: Stacks Collection Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Stacks Collection Directory parameter.		false	role_config_suppress ion_stacks_collection_di rectory	true

Display Name	Description	Related Name	Default Value	API Name	Require
Suppress Parameter Validation: Streams Replication Manager Driver Advanced Configuration Snippet (Safety Valve) for staging/globals.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver Advanced Configuration Snippet (Safety Valve) for staging/globals.xml parameter.		false	role_config_suppression_staging/globals.xml_role_safety_valve	true
Suppress Parameter Validation: Streams Replication Manager Driver Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml parameter.		false	role_config_suppression_staging/srm-common.xml_role_safety_valve	true
Suppress Parameter Validation: Streams Replication Manager Configuration Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Configuration Directory parameter.		false	role_config_suppression_streams.replication.manager.conf.directory	true
Suppress Parameter Validation: streams replication manager driver target cluster	Whether to suppress configuration warnings produced by the built-in parameter validation for the streams replication manager driver target cluster parameter.		false	role_config_suppression_streams.replication.manager.driver.target_cluster	true
Suppress Parameter Validation: Java Home Path Override	Whether to suppress configuration warnings produced by the built-in parameter validation for the Java Home Path Override parameter.		false	role_config_suppression_streams.replication.manager.jdk.home	true
Suppress Parameter Validation: Streams Replication Manager Driver Environment Advanced Configuration Snippet (Safety Valve)	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Driver Environment Advanced Configuration Snippet (Safety Valve) parameter.		false	role_config_suppression_streams_replication_manager_driver_role_env_safety_valve	true
Suppress Health Test: Audit Pipeline Test	Whether to suppress the results of the Audit Pipeline Test health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_audit_health	true
Suppress Health Test: File Descriptors	Whether to suppress the results of the File Descriptors health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_file_descriptor	true
Suppress Health Test: Host Health	Whether to suppress the results of the Host Health health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_host_health	true

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Health Test: Log Directory Free Space	Whether to suppress the results of the Log Directory Free Space health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_log_directory_free_space	true
Suppress Health Test: Process Status	Whether to suppress the results of the Process Status health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_scm_health	true
Suppress Health Test: Swap Memory Usage	Whether to suppress the results of the Swap Memory Usage health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_swap_memory_usage	true
Suppress Health Test: Unexpected Exits	Whether to suppress the results of the Unexpected Exits health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_driver_unexpected_exits	true

Streams Replication Manager Service

Table 14: Advanced

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Service XML Override	For advanced use only, replace entire XML in the logback configuration file for Streams Replication Manager Service, ignoring all logging configuration.	logback_safety_valve		logback_safety_valve	false
Heap Dump Directory	Path to directory where heap dumps are generated when <code>java.lang.OutOfMemoryError</code> error is thrown. This directory is automatically created if it does not exist. If this directory already exists, role user must have write access to this directory. If this directory is shared among multiple roles, it should have 1777 permissions. The heap dump files are created with 600 permissions and are owned by the role user. The amount of free space in this directory should be greater than the maximum Java Process heap size configured for this role.	oom_heap_dump_dir	/tmp	oom_heap_dump_dir	false

Display Name	Description	Related Name	Default Value	API Name	Required
Dump Heap When Out of Memory	When set, generates heap dump file when <code>java.lang.OutOfMemoryError</code> is thrown.		true	<code>oom_heap_dump_enabled</code>	true
Kill When Out of Memory	When set, a SIGKILL signal is sent to the role process when <code>java.lang.OutOfMemoryError</code> is thrown.		true	<code>oom_sigkill_enabled</code>	true
Automatically Restart Process	When set, this role's process is automatically (and transparently) restarted in the event of an unexpected failure.		false	<code>process_auto_restart</code>	true
Enable Metric Collection	Cloudera Manager agent monitors each service and each of its role by publishing metrics to the Cloudera Manager Service Monitor. Setting it to false will stop Cloudera Manager agent from publishing any metric for corresponding service/roles. This is usually helpful for services that generate large amount of metrics which Service Monitor is not able to process.		true	<code>process_should_monitor</code>	true
Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for <code>staging/globals.xml</code>	For advanced use only. A string to be inserted into <code>staging/globals.xml</code> for this role only.			<code>staging/globals.xml_role_safety_valve</code>	false
Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for <code>staging/srm-common.xml</code>	For advanced use only. A string to be inserted into <code>staging/srm-common.xml</code> for this role only.			<code>staging/srm-common.xml_role_safety_valve</code>	false
Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for <code>staging/srm-service-common.xml</code>	For advanced use only. A string to be inserted into <code>staging/srm-service-common.xml</code> for this role only.			<code>staging/srm-service-common.xml_role_safety_valve</code>	false
Streams Replication Manager Service Environment Advanced Configuration Snippet (Safety Valve)	For advanced use only, key-value pairs (one on each line) to be inserted into a role's environment. Applies to configurations of this role except client configuration.			<code>STREAMS_REPLICATION_MANAGER_SERVICE_role_env_safety_valve</code>	false

Table 15: Logs

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Service Log Directory	The log directory for log files of the role Streams Replication Manager Service.	<code>log_dir</code>	<code>/var/log/streams-replication-manager</code>	<code>log_dir</code>	false
Streams Replication Manager Service Logging Threshold	The minimum log level for Streams Replication Manager Service logs		INFO	<code>log_threshold</code>	false

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Service Maximum Log File Backups	The maximum number of rolled log files to keep for Streams Replication Manager Service logs. Typically used by log4j or logback.		10	max_log_backup_index	false
Streams Replication Manager Service Max Log Size	The maximum size, in megabytes, per log file for Streams Replication Manager Service logs. Typically used by log4j or logback.		200 MiB	max_log_size	false

Table 16: Monitoring

Display Name	Description	Related Name	Default Value	API Name	Required
Enable Health Alerts for this Role	When set, Cloudera Manager will send alerts when the health of this role reaches the threshold specified by the EventServer setting eventserver_health_events_alert_threshold		true	enable_alerts	false
Enable Configuration Change Alerts	When set, Cloudera Manager will send alerts when this entity's configuration changes.		false	enable_config_alerts	false
Log Directory Free Space Monitoring Absolute Thresholds	The health test thresholds for monitoring of free space on the filesystem that contains this role's log directory.		Warning: 10 GiB, Critical: 5 GiB	log_directory_free_space_absolute_thresholds	false
Log Directory Free Space Monitoring Percentage Thresholds	The health test thresholds for monitoring of free space on the filesystem that contains this role's log directory. Specified as a percentage of the capacity on that filesystem. This setting is not used if a Log Directory Free Space Monitoring Absolute Thresholds setting is configured.		Warning: Never, Critical: Never	log_directory_free_space_percentage_thresholds	false
Process Swap Memory Thresholds	The health test thresholds on the swap memory usage of the process. This takes precedence over the host level threshold.		Warning: 200 B, Critical: Never	process_swap_memory_thresholds	false

Display Name	Description	Related Name	Default Value	API Name	Required
Role Triggers	<p>The configured triggers for this role. This is a JSON-formatted list of triggers. These triggers are evaluated as part of the health system. Every trigger expression is parsed, and if the trigger condition is met, the list of actions provided in the trigger expression is executed. Each trigger has the following fields:</p> <ul style="list-style-type: none"> triggerName (mandatory) - The name of the trigger. This value must be unique for the specific role. triggerExpression (mandatory) - A tsquery expression representing the trigger. streamThreshold (optional) - The maximum number of streams that can satisfy a condition of a trigger before the condition fires. By default set to 0, and any stream returned causes the condition to fire. enabled (optional) - By default set to 'true'. If set to 'false', the trigger is not evaluated. expressionEditorConfig (optional) - Metadata for the trigger editor. If present, the trigger should only be edited from the Edit Trigger page; editing the trigger here can lead to inconsistencies. <p>For example, the following JSON formatted trigger configured for a DataNode fires if the DataNode has more than 1500 file descriptors opened: [{"triggerName": "sample-trigger", "triggerExpression": "IF (SELECT fd_open WHERE roleName=\$ROLENAME and last(fd_open) > 1500) DO health:bad", "streamThreshold": 0, "enabled": "true"}] See the trigger rules documentation for more details on how to write triggers using tsquery. The JSON format is evolving and may change and, as a result, backward compatibility is not guaranteed between releases.</p>		[]	role_triggers	true
File Descriptor Monitoring Thresholds	The health test thresholds of the number of file descriptors used. Specified as a percentage of file descriptor limit.		Warning: 50.0 %, Critical: 70.0 %	streams_replication_manager_service_fd_thresholds	false
Streams Replication Manager Service Host Health Test	When computing the overall Streams Replication Manager Service health, consider the host's health.		true	streams_replication_manager_service_host_health_enabled	false

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Service Process Health Test	Enables the health test that the Streams Replication Manager Service's process state is consistent with the role configuration		true	streams_replication_manager_service_scm_health_enabled	false
Unexpected Exits Thresholds	The health test thresholds for unexpected exits encountered within a recent period specified by the unexpected_exits_window configuration for the role.		Warning: Never, Critical: Any	unexpected_exits_thresholds	false
Unexpected Exits Monitoring Period	The period to review when computing unexpected exits.		5 minute(s)	unexpected_exits_window	false

Table 17: Other

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Configuration Directory	Directory with all configurations for SRM's daemons.	streams.replication.manager.conf.directory	\$CONF_DIR	streams.replication.manager.conf.directory	false
SRM Service port	SRM Service port.	streams.replication.manager.service.port	6670	streams.replication.manager.service.port	true
SRM Service rpc port	SRM Service rpc port.	streams.replication.manager.service.rpcPort	\$streams.replication.manager.service.rpcPort	streams.replication.manager.service.rpcPort	true
SRM Service https port	SRM Service https port.	streams.replication.manager.service.ssl.port	6671	streams.replication.manager.service.ssl.port	true
SRM Service https rpc port	SRM Service https rpc port.	streams.replication.manager.service.ssl.rpcPort	\$streams.replication.manager.service.ssl.rpcPort	streams.replication.manager.service.ssl.rpcPort	true
streams replication manager service target cluster	Target cluster from which the SRM Service collects data.	streams.replication.manager.service.target.cluster		streams.replication.manager.service.target.cluster	true
Keystore Type	The keystore type. Required if SSL is enabled for the SRM Service. For example, PKCS12 or JKS.	streams.replication.manager.ssl.keyStoreType	JKS	streams.replication.manager.ssl.keyStoreType	false
TrustStore Type	The truststore type. Required if SSL is enabled for the SRM Service. For example, PKCS12 or JKS.	streams.replication.manager.ssl.trustStoreType	JKS	streams.replication.manager.ssl.trustStoreType	false
validateCerts	Whether or not to validate TLS certificates before starting. If enabled, SRM will not start with expired or otherwise invalid certificates.	streams.replication.manager.ssl.validateCerts	false	streams.replication.manager.ssl.validateCerts	false
validatePeers	Whether or not to validate TLS peer certificates.	streams.replication.manager.ssl.validatePeers	false	streams.replication.manager.ssl.validatePeers	false

Table 18: Performance

Display Name	Description	Related Name	Default Value	API Name	Required
Maximum Process File Descriptors	If configured, overrides the process soft and hard rlimits (also called ulimits) for file descriptors to the configured value.			rlimit_fds	false

Table 19: Resource Management

Display Name	Description	Related Name	Default Value	API Name	Required
Cgroup CPU Shares	Number of CPU shares to assign to this role. The greater the number of shares, the larger the share of the host's CPUs that will be given to this role when the host experiences CPU contention. Must be between 2 and 262144. Defaults to 1024 for processes not managed by Cloudera Manager.	cpu.shares	1024	rm_cpu_shares	true
Cgroup I/O Weight	Weight for the read I/O requests issued by this role. The greater the weight, the higher the priority of the requests when the host experiences I/O contention. Must be between 100 and 1000. Defaults to 1000 for processes not managed by Cloudera Manager.	blkio.weight	500	rm_io_weight	true
Cgroup Memory Hard Limit	Hard memory limit to assign to this role, enforced by the Linux kernel. When the limit is reached, the kernel will reclaim pages charged to the process. If reclaiming fails, the kernel may kill the process. Both anonymous as well as page cache pages contribute to the limit. Use a value of -1 to specify no limit. By default processes not managed by Cloudera Manager will have no limit. If the value is -1, Cloudera Manager will not monitor Cgroup memory usage therefore some of the charts will show 'No Data'	memory.limit_in_bytes	-1 MiB	rm_memory_hard_limit	true
Cgroup Memory Soft Limit	Soft memory limit to assign to this role, enforced by the Linux kernel. When the limit is reached, the kernel will reclaim pages charged to the process if and only if the host is facing memory pressure. If reclaiming fails, the kernel may kill the process. Both anonymous as well as page cache pages contribute to the limit. Use a value of -1 to specify no limit. By default processes not managed by Cloudera Manager will have no limit. If the value is -1, Cloudera Manager will not monitor Cgroup memory usage therefore some of the charts will show 'No Data'	memory.soft_limit_in_bytes	-1 MiB	rm_memory_soft_limit	true

Table 20: Security

Display Name	Description	Related Name	Default Value	API Name	Required
Streams Replication Manager Service TLS/SSL Client Trust Store File	The location on disk of the trust store, in .jks format, used to confirm the authenticity of TLS/SSL servers that Streams Replication Manager Service might connect to. This is used when Streams Replication Manager Service is the client in a TLS/SSL connection. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.	ssl.truststore.location		ssl_client_truststore_location	false
Streams Replication Manager Service TLS/SSL Client Trust Store Password	The password for the Streams Replication Manager Service TLS/SSL Certificate Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file. The contents of trust stores are certificates, and certificates are public information.	ssl.truststore.password		ssl_client_truststore_password	false
Enable TLS/SSL for Streams Replication Manager Service	Encrypt communication between clients and Streams Replication Manager Service using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)).	ssl.enabled	false	ssl.enabled	false
Streams Replication Manager Service TLS/SSL Server JKS Keystore Key Password	The password that protects the private key contained in the JKS keystore used when Streams Replication Manager Service is acting as a TLS/SSL server.	ssl.key.password		ssl_server_keystore_keypassword	false
Streams Replication Manager Service TLS/SSL Server JKS Keystore File Location	The path to the TLS/SSL keystore file containing the server certificate and private key used for TLS/SSL. Used when Streams Replication Manager Service is acting as a TLS/SSL server. The keystore must be in JKS format.	ssl.keystore.location		ssl_server_keystore_location	false
Streams Replication Manager Service TLS/SSL Server JKS Keystore File Password	The password for the Streams Replication Manager Service JKS keystore file.	ssl.keystore.password		ssl_server_keystore_password	false

Table 21: Stack Collection

Display Name	Description	Related Name	Default Value	API Name	Required
Stacks Collection Data Retention	The amount of stacks data that is retained. After the retention limit is reached, the oldest data is deleted.	stacks_collection_data_retention	100 MiB	stacks_collection_data_retention	false

Display Name	Description	Related Name	Default Value	API Name	Required
Stacks Collection Directory	The directory in which stacks logs are placed. If not set, stacks are logged into a stacks subdirectory of the role's log directory.	stacks_collection_directory		stacks_collection_directory	false
Stacks Collection Enabled	Whether or not periodic stacks collection is enabled.	stacks_collection_enabled	false	stacks_collection_enabled	true
Stacks Collection Frequency	The frequency with which stacks are collected.	stacks_collection_frequency	5.0 second(s)	stacks_collection_frequency	false
Stacks Collection Method	The method used to collect stacks. The jstack option involves periodically running the jstack command against the role's daemon process. The servlet method is available for those roles that have an HTTP server endpoint exposing the current stacks traces of all threads. When the servlet method is selected, that HTTP endpoint is periodically scraped.	stacks_collection_method	jstack	stacks_collection_method	false

Table 22: Suppressions

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Configuration Validator: CDH Version Validator	Whether to suppress configuration warnings produced by the CDH Version Validator configuration validator.		false	role_config_suppression_cdh_version_validator	true
Suppress Parameter Validation: Streams Replication Manager Service Log Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Log Directory parameter.		false	role_config_suppression_log_dir	true
Suppress Parameter Validation: Streams Replication Manager Service XML Override	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service XML Override parameter.		false	role_config_suppression_logback_safety_valve	true
Suppress Parameter Validation: Heap Dump Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Heap Dump Directory parameter.		false	role_config_suppression_oom_heap_dump_dir	true
Suppress Parameter Validation: Role Triggers	Whether to suppress configuration warnings produced by the built-in parameter validation for the Role Triggers parameter.		false	role_config_suppression_role_triggers	true
Suppress Parameter Validation: Streams Replication Manager Service TLS/SSL Client Trust Store File	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service TLS/SSL Client Trust Store File parameter.		false	role_config_suppression_ssl_client_truststore_location	true
Suppress Parameter Validation: Streams Replication Manager Service TLS/SSL Client Trust Store Password	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service TLS/SSL Client Trust Store Password parameter.		false	role_config_suppression_ssl_client_truststore_password	true

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Parameter Validation: Streams Replication Manager Service TLS/SSL Server JKS Keystore Key Password	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service TLS/SSL Server JKS Keystore Key Password parameter.		false	role_config_suppression_ssl_server_keystore_keypassword	true
Suppress Parameter Validation: Streams Replication Manager Service TLS/SSL Server JKS Keystore File Location	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service TLS/SSL Server JKS Keystore File Location parameter.		false	role_config_suppression_ssl_server_keystore_location	true
Suppress Parameter Validation: Streams Replication Manager Service TLS/SSL Server JKS Keystore File Password	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service TLS/SSL Server JKS Keystore File Password parameter.		false	role_config_suppression_ssl_server_keystore_password	true
Suppress Parameter Validation: Stacks Collection Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Stacks Collection Directory parameter.		false	role_config_suppression_stacks_collection_directory	true
Suppress Parameter Validation: Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/globals.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/globals.xml parameter.		false	role_config_suppression_staging_globals_xml_role_safety_valve	true
Suppress Parameter Validation: Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-common.xml parameter.		false	role_config_suppression_staging_srm_common_xml_role_safety_valve	true
Suppress Parameter Validation: Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-service-common.xml	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Advanced Configuration Snippet (Safety Valve) for staging/srm-service-common.xml parameter.		false	role_config_suppression_staging_srm_service_common_xml_role_safety_valve	true
Suppress Parameter Validation: Streams Replication Manager Configuration Directory	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Configuration Directory parameter.		false	role_config_suppression_streams_replication_manager_conf_directory	true
Suppress Parameter Validation: SRM Service port	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM Service port parameter.		false	role_config_suppression_streams_replication_manager_service_port	true

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Parameter Validation: SRM Service rpc port	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM Service rpc port parameter.		false	role_config_suppression_streams.replication.manager.service.rpcport	true
Suppress Parameter Validation: SRM Service https port	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM Service https port parameter.		false	role_config_suppression_streams.replication.manager.service.ssl.port	true
Suppress Parameter Validation: SRM Service https rpc port	Whether to suppress configuration warnings produced by the built-in parameter validation for the SRM Service https rpc port parameter.		false	role_config_suppression_streams.replication.manager.service.ssl.rpcport	true
Suppress Parameter Validation: streams replication manager service target cluster	Whether to suppress configuration warnings produced by the built-in parameter validation for the streams replication manager service target cluster parameter.		false	role_config_suppression_streams.replication.manager.service.target.cluster	true
Suppress Parameter Validation: Keystore Type	Whether to suppress configuration warnings produced by the built-in parameter validation for the Keystore Type parameter.		false	role_config_suppression_streams.replication.manager.ssl.keystoretype	true
Suppress Parameter Validation: TrustStore Type	Whether to suppress configuration warnings produced by the built-in parameter validation for the TrustStore Type parameter.		false	role_config_suppression_streams.replication.manager.ssl.truststoretype	true
Suppress Parameter Validation: validateCerts	Whether to suppress configuration warnings produced by the built-in parameter validation for the validateCerts parameter.		false	role_config_suppression_streams.replication.manager.ssl.validatecerts	true
Suppress Parameter Validation: validatePeers	Whether to suppress configuration warnings produced by the built-in parameter validation for the validatePeers parameter.		false	role_config_suppression_streams.replication.manager.ssl.validatepeers	true
Suppress Parameter Validation: Streams Replication Manager Service Environment Advanced Configuration Snippet (Safety Valve)	Whether to suppress configuration warnings produced by the built-in parameter validation for the Streams Replication Manager Service Environment Advanced Configuration Snippet (Safety Valve) parameter.		false	role_config_suppression_streams_replication_manager_service_role_env_safety_valve	true
Suppress Health Test: Audit Pipeline Test	Whether to suppress the results of the Audit Pipeline Test health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_service_audit_health	true
Suppress Health Test: File Descriptors	Whether to suppress the results of the File Descriptors health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppression_streams_replication_manager_streams_replication_manager_service_file_descriptor	true

Display Name	Description	Related Name	Default Value	API Name	Required
Suppress Health Test: Host Health	Whether to suppress the results of the Host Health health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppress ion_streams_replication_ manager_streams_repl ication_manager_serv ice_host_health	true
Suppress Health Test: Log Directory Free Space	Whether to suppress the results of the Log Directory Free Space health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppress ion_streams_replication_ manager_streams_repl ication_manager_serv ice_log_directory_free_s pace	true
Suppress Health Test: Process Status	Whether to suppress the results of the Process Status health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppress ion_streams_replication_ manager_streams_repl ication_manager_serv ice_scm_health	true
Suppress Health Test: Swap Memory Usage	Whether to suppress the results of the Swap Memory Usage health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppress ion_streams_replicat ion_manager_streams_ replication_manager_ service_swap_memory_ usage	true
Suppress Health Test: Unexpected Exits	Whether to suppress the results of the Unexpected Exits health test. The results of suppressed health tests are ignored when computing the overall health of the associated host, role or service, so suppressed health tests will not generate alerts.		false	role_health_suppress ion_streams_replication_ manager_streams_repl ication_manager_serv ice_unexpected_exits	true