

Data Collector Reference - Services

Data Infrastructure Insights

NetApp August 28, 2025

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructureinsights/task_config_telegraf_node.html on August 28, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Data Collector Reference - Services
Node Data Collection
Installation
Objects and Counters
Setup
ActiveMQ Data Collector
Installation
Setup
Objects and Counters
Troubleshooting
Apache Data Collector
Installation
Setup
Objects and Counters
Troubleshooting
Consul Data Collector
Installation
Setup
Objects and Counters for consul
Troubleshooting
Couchbase Data Collector 9
Installation
Setup
Objects and Counters
Troubleshooting
CouchDB Data Collector
Installation
Setup
Objects and Counters
Troubleshooting
Docker Data Collector
Installation
Setup
Objects and Counters
Troubleshooting
Elasticsearch Data Collector
Setup
Objects and Counters
Troubleshooting
Flink Data Collector
Installation
Setup
Objects and Counters

Troubleshooting	. 30
Hadoop Data Collector	. 30
Installation	. 30
Setup	. 31
Objects and Counters	35
Troubleshooting	. 36
HAProxy Data Collector	. 36
Installation	36
Setup	. 37
Objects and Counters	. 39
Troubleshooting	41
JVM Data Collector	42
Installation	42
Setup	43
Objects and Counters	43
Troubleshooting	46
Kafka Data Collector	46
Installation	46
Setup	47
Objects and Counters	48
Troubleshooting	48
Kibana Data Collector	49
Installation	49
Setup	50
Objects and Counters	50
Troubleshooting	51
Kubernetes Monitoring Operator Installation and Configuration	51
Before installing the Kubernetes Monitoring Operator	51
Installing the Kubernetes Monitoring Operator	51
Kubernetes Monitoring Components	54
Upgrading to the latest Kubernetes Monitoring Operator	54
Stopping and Starting the Kubernetes Monitoring Operator	56
Uninstalling.	56
About Kube-state-metrics	57
Configuring/Customizing the Operator	57
A Note About Secrets	61
Verifying Kubernetes Monitoring Operator Image Signatures	61
Troubleshooting	62
Memcached Data Collector	. 71
Installation	. 71
Setup	72
Objects and Counters	72
Troubleshooting	74
MongoDB Data Collector	. 74
Installation	. 74

Setup	ō
Objects and Counters	ō
Troubleshooting	6
MySQL Data Collector	6
Installation	6
Setup	7
Objects and Counters	3
Troubleshooting	1
Netstat Data Collector	1
Installation	1
Setup	2
Objects and Counters	2
Troubleshooting	2
Nginx Data Collector	2
Installation	3
Setup	4
Objects and Counters	4
Troubleshooting	ō
PostgreSQL Data Collector	ō
Installation	5
Setup	6
Objects and Counters	6
Troubleshooting	7
Puppet Agent Data Collector	7
Installation	7
Setup	3
Objects and Counters	3
Troubleshooting	9
Redis Data Collector	9
Installation	9
Setup)
Objects and Counters	1
Troubleshooting	1

Data Collector Reference - Services

Node Data Collection

Data Infrastructure Insights gathers metrics from the node on which you install an agent.

Installation

- 1. From **Observability > Collectors**, choose an operating system/platform. Note that installing any integration data collector (Kubernetes, Docker, Apache, etc.) will also configure node data collection.
- 2. Follow the instructions to configure the agent. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

Objects and Counters

The following objects and their counters are collected as Node metrics:

Object:	Identifiers:	Attributes:	Datapoints:
Node Filesystem	Node UUID Device Path Type	Node IP Node Name Node OS Mode	Free Inodes Free Inodes Total Inodes Used Total Used Total Used
Node Disk	Node UUID Disk	Node IP Node Name Node OS	IO Time Total IOPS In Progress Read Bytes (per sec) Read Time Total Reads (per sec) Weighted IO Time Total Write Bytes (per sec) Write Time Total Writes (per sec) Current Disk Queue Length Write Time Read Time IO Time
Node CPU	Node UUID CPU	Node IP Node Name Node OS	System CPU Usage User CPU Usage Idle CPU Usage Processor CPU Usage Interrupt CPU Usage DPC CPU Usage

Object:	Identifiers:	Attributes:	Datapoints:
Node	Node UUID	Node IP Node OS	Kernel Boot Time Kernel Context Switches (per sec) Kernel Entropy Available Kernel Interrupts (per sec) Kernel Processes Forked (per sec) Memory Active Memory Available Total Memory Available Total Memory Cached Memory Cached Memory Commit Limit Memory Committed As Memory Dirty Memory Free Memory High Free Memory High Total Memory Huge Pages Size Memory Huge Pages Total Memory Low Free Memory Low Free Memory Low Total Memory Shared Memory Shared Memory Shared Memory Swap Cached Memory Swap Free Memory Swap Free Memory Used Total Memory Used Total Memory Used Total Memory Vmalloc Chunk Memory Vmalloc Chunk Memory Wired Memory Wirteback Total Memory Wirteback Total Memory Wirteback Total Memory Wirteback Total Memory Wirteback Total Memory Wirteback Total Memory Page Faults Memory Page Faults Memory Pages Memory Paged Memory Paged Memory Paged Memory Paged Memory Paged Memory Paged Memory Standby Cache Reserve Memory Standby Cache Reserve Memory Transition Faults Processes Blocked Processes Dead

Object:	Identifiers:	Attributes:	Datapoints:
Node Network	Network Interface Node UUID	Node Name Node IP Node OS	Bytes Received Bytes Sent Packets Outboud Discarded Packets Outboud Errors Packets Received Discarded Packets Received Errors Packets Received Packets Sent

Setup and Troubleshooting information can be found on the Configuring an Agent page.

ActiveMQ Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from ActiveMQ.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose ActiveMQ.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

	ActiveMQ Configuration Gathers ActiveMQ metrics.	
What Operating Sy	stem or Platform Are You Using?	Need Help?
e Windows	•	
Select existing Age	ent Access Key or create a new one	
Default (405fb5ec-d4cb-4404	4-977b-71fa931e1ad3) Agent Access Key	
*Please ensure that you ha	ve a Telegraf Agent in you environment before configuring Show Instruction	ans
ollow Configuration	on Steps	Need Help
[[inputs.act ## Required	ivemq]] i ActiveMO Endpoint, port	1
<pre>[[inputs.acti ## Required ## USER-ACT server = "4 port = <ins <insert_a="" a<="" address,="" machine="" pre="" replace=""></ins></pre>	ivemq]] J ActiveMQ Endpoint, port IION: Provide address of ActiveMQ, HTTP port for ActiveMQ INSERT_ACTIVEMQ_ADDRESS>" SERT_ACTIVEMQ_PORT> CTIVEMQ_ADDRESS> with the applicable ActiveMQ server ad nd refrain from using a loopback address.	ddress. Please specify a real
<pre>[[inputs.act] ## Required ## USER-ACT server = "d port = <ins <insert_a="" <insert_a<="" address,="" al="" machine="" pre="" replace=""></ins></pre>	ivemq]] d ActiveMQ Endpoint, port IIN: Provide address of ActiveMQ, HTTP port for ActiveMQ (INSERT_ACTIVEMQ_ADDRESS>" SERT_ACTIVEMQ_PORT> CTIVEMQ_ADDRESS> with the applicable ActiveMQ server ad nd refrain from using a loopback address. CTIVEMQ_PORT> with the applicable ActiveMQ server HTTP	ddress. Please specify a real
 Imputs.active Required ## Required ## USER-ACTIVE Server = "Content of the server and the server a	ivemq]] d ActiveMQ Endpoint, port ITON: Provide address of ActiveMQ, HTTP port for ActiveMQ (INSERT_ACTIVEMQ_ADDRESS>" SERT_ACTIVEMQ_PORT> CTIVEMQ_ADDRESS> with the applicable ActiveMQ server ad nd refrain from using a loopback address. CTIVEMQ_PORT> with the applicable ActiveMQ server HTTP CTIVEMQ_USERNAME> and <insert_activemq_passwor ils.</insert_activemq_passwor 	ddress. Please specify a real port. D> with the applicable
 Imputs.active Required ## Required ## USER-Active Server = "does not be a constructed on the server of the server o	ivemq]] d ActiveMQ Endpoint, port TION: Provide address of ActiveMQ, HTTP port for ActiveMQ CINSERT_ACTIVEMQ_ADDRESS>" SERT_ACTIVEMQ_PORT> CTIVEMQ_ADDRESS> with the applicable ActiveMQ server ad nd refrain from using a loopback address. CTIVEMQ_PORT> with the applicable ActiveMQ server HTTP CTIVEMQ_USERNAME> and <insert_activemq_passwor ils. if needed (if ActiveMQ server changes web admin root path).</insert_activemq_passwor 	ddress. Please specify a real port. D> with the applicable
 Imputs.action Imputs.action Imputs.action Imputs.action Imputs.action Imputs.action Replace <insert_a< li=""> </insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<></insert_a<>	<pre>ivemq]] d ActiveMQ Endpoint, port TTON: Provide address of ActiveMQ, HTTP port for ActiveMQ CINSERT_ACTIVEMQ_ADDRESS>" SERT_ACTIVEMQ_ADDRESS> with the applicable ActiveMQ server ad nd refrain from using a loopback address. CTIVEMQ_PORT> with the applicable ActiveMQ server HTTP CTIVEMQ_USERNAME> and <insert_activemq_passwor '="" (if="" (to="" activemq="" admin="" avoid="" changes="" clashes)<="" disambiguation="" for="" if="" ils.="" name="" needed="" path).="" pre="" root="" server="" web=""></insert_activemq_passwor></pre>	ddress. Please specify a real port. D> with the applicable
 [[inputs.act] ## Required ## USER-ACI server = "d port = <ins< li=""> Replace <insert_a machine address, and Replace <insert_a< li=""> Replace <insert_a ActiveMQ credentia</insert_a Modify 'webadmin' Modify 'Namespace Restart the Telegraf </insert_a<></insert_a </ins<>	<pre>ivemq]] d ActiveNQ Endpoint, port IION: Provide address of ActiveNQ, HTTP port for ActiveMQ CINSERT_ACTIVENQ_ADDRESS>" SERT_ACTIVENQ_PORT> CTIVEMQ_ADDRESS> with the applicable ActiveMQ server ad nd refrain from using a loopback address. CTIVEMQ_PORT> with the applicable ActiveMQ server HTTP CTIVEMQ_USERNAME> and <insert_activemq_passwor 'if="" (if="" (to="" activemq="" admin="" avoid="" changes="" clashes)="" disambiguation="" for="" ils.="" name="" needed="" path).="" pre="" root="" server="" service.<="" web=""></insert_activemq_passwor></pre>	ddress. Please specify a real port. D> with the applicable

Information may be found in the ActiveMQ documentation

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:	Datapoints:
ActiveMQ Queue	Namespace Queue Port Server	Node Name Node IP Node UUID	Consumer Count Dequeue Count Enqueue Count Queue Size
ActiveMQ Subscriber	Client ID Connection ID Port Server Namespace	Is Active Destination Node Name Node IP Node UUID Node OS Selector Subscription	Dequeue Count Dispatched Count Dispatched Queue Size Enqueue Count Pending Queue Size
ActiveMQ Topic	Topic Port Server Namespace	Node Name Node IP Node UUID Node OS	Consumer Count Dequeue Count Enqueue Count Size

Troubleshooting

Additional information may be found from the Support page.

Apache Data Collector

This data collector allows collection of data from Apache servers on your tenant.

Pre-requisites

- · You must have your Apache HTTP Server set up and properly running
- · You must have sudo or administrator permissions on your agent host/VM
- Typically, the Apache *mod_status* module is configured to expose a page at the '/server-status?auto' location of the Apache server. The *ExtendedStatus* option must be enabled in order to collect all available fields. For information about how to configure your server, see the Apache module documentation: https://httpd.apache.org/docs/2.4/mod/mod_status.html#enable

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Apache.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the + Agent Access Key button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.

4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

	Apache Configuration Gathers Apache metrics.	
What Operating S	stem or Platform Are You Using?	Need Help?
🔞 Ubuntu & Debian		
Select existing Age	nt Access Key or create a new one	
Default (405fb5ec-d4cb-440	-977b-71fa931e1ad3)	
*Please ensure that you ha	ve a Telegraf Agent in you environment before configuring. Show In	Istructions
Follow Configurati	on Steps	Need Help?
• • • • • • • • • •		
 Ensure that the Apa enabled and exposit Copy the contents I contents to the /etc 	che HTTP Server system you're going to gather metrics ed. For details refer to the following document. elow into a new .conf file under the /etc/telegraf/telegr /telegraf/telegraf.d/cloudinsights-apache.conf file.	on has the 'mod_status' module af.d/ directory. For example, copy the
 Ensure that the Apa enabled and exposite contents to the /etc Copy the contents R contents to the /etc [[inputs.apach ## An array ## readable ## USER-ACTI server-status. -* Pi 	che HTTP Server system you're going to gather metrics ed. For details refer to the following document. elow into a new conf file under the /etc/telegraf/telegr /telegraf/telegraf.d/cloudinsights-apache.conf file.]] If URLs to gather from, must be directed at the machi rersion of the mod_status page including the auto que W: Provide address of apache server, port for apache	on has the 'mod_status' module af.d/directory. For example, copy the ry string. server, confirm path for
 Ensure that the Apa enabled and expose Copy the contents I contents to the /etc Copy the contents to the /etc Copy th	che HTTP Server system you're going to gather metrics ed. For details refer to the following document. eelow into a new conf file under the /etc/telegraf/telegr /telegraf/telegraf.d/cloudinsights-apache.conf file. [] If URLs to gather from, must be directed at the machi rersion of the mod_status page including the auto que W: Provide address of apache server, port for apache EECHE_ADDRESS> with the applicable Apache server of from using a loopback address.	on has the 'mod_status' module af.d/ directory. For example, copy the ry string. server, confirm path for 1toric office address. Please specify a real machine
 Ensure that the Apa enabled and expose contents to the /etc Copy the contents I contents to the /etc [[inputs.apach ## An array ## readable ## USER-ACTI server-status. "***********************************	che HTTP Server system you're going to gather metrics ed. For details refer to the following document. eelow into a new conf file under the /etc/telegraf/telegr /telegraf/telegraf.d/cloudinsights-apache.conf file. e]] if URLs to gather from, must be directed at the machi rersion of the mod_status page including the auto que W: Provide address of apache server, port for apache efforted at the machine from the applicable Apache server of from using a loopback address. PACHE_PORT> with the applicable Apache server port.	on has the 'mod_status' module af.d/directory. For example, copy the ry string. server; confirm path for ?t! -tf (f - address. Please specify a real machine
 Ensure that the Apa enabled and expose contents to the /etc (inputs.apach ## An array ## readable ## USER-ACTI server-status. nt	che HTTP Server system you're going to gather metrics ed. For details refer to the following document. eelow into a new conf file under the /etc/telegraf/telegr /telegraf/telegraf.d/cloudinsights-apache.conf file. [] f URLs to gather from, must be directed at the machi ersion of the mod_status page including the auto que N: Provide address of apache server, port for apache effected at the machine from using a loopback address. PACHE_PORT> with the applicable Apache server port. tatus' path in accordance to the Apache server configu	on has the 'mod_status' module af.d/ directory. For example, copy the re ry string. server, confirm path for
 Ensure that the Apa enabled and expose contents to the /etc (inputs.apach ## An array ## readable ## USER-ACTI server=status. 	che HTTP Server system you're going to gather metrics ed. For details refer to the following document. eelow into a new conf file under the /etc/telegraf/telegr /telegraf/telegraf.d/cloudinsights-apache.conf file.	on has the 'mod_status' module af.d/directory. For example, copy the re ry string. server, confirm path for ?to- re address. Please specify a real machine aration.

Setup

Telegraf's plugin for Apache's HTTP Server relies on the 'mod_status' module to be enabled. When this is enabled, Apache's HTTP Server will expose an HTML endpoint that can be viewed on your browser or scraped for extraction of status of all Apache's HTTP Server configuration.

Compatibility:

Configuration was developed against Apache's HTTP Server version 2.4.38.

Enabling mod_status:

Enabling and exposing the 'mod_status' modules involves two steps:

- · Enabling module
- Exposing stats from module

Enabling module:

The loading of modules is controlled by the config file under '/usr/local/apache/conf/httpd.conf'. Edit the config file and uncomment the following lines:

```
LoadModule status_module modules/mod_status.so
```

```
Include conf/extra/httpd-info.conf
```

Exposing stats from module:

The exposing of 'mod_status' is controlled by the config file under '/usr/local/apache2/conf/extra/httpdinfo.conf'. Make sure you have the following in that configuration file (at least, other directives will be there):

```
# Allow server status reports generated by mod_status,
# with the URL of http://servername/server-status
<Location /server-status>
    SetHandler server-status
</Location>
#
# # ExtendedStatus controls whether Apache will generate "full" status
# information (ExtendedStatus On) or just basic information
(ExtendedStatus
# Off) when the "server-status" handler is called. The default is Off.
#
ExtendedStatus On
```

For detailed instructions on the 'mod_status' module, see the Apache documentation

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
Apache	Namespace Server	Node IP Node Name Port Parent Server Config Generation Parent Server MPM Generation Server Uptime Is Stopping	Busy Workers Bytes per Request Bytes per Second CPU Children System CPU Children User CPU Load CPU System CPU User Asynchronous Connections Closing Asynchronous Connections Keep Alive Asynchronous Connections Writing Connections Total Duration per Request Idle Workers Load Average (last 1m) Load Average (last 15m) Load Average (last 5m) Processes Requests per Second Total Accesses Total Duration Total KBytes Scoreboard Closing Scoreboard Idle Cleanup Scoreboard Idle Cleanup Scoreboard Idle Cleanup Scoreboard Logging Scoreboard Sending Scoreboard Sending Scoreboard Sending Scoreboard Sending Scoreboard Starting Scoreboard Waiting

Additional information may be found from the Support page.

Consul Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Consul.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Consul.

If you haven't configured an Agent for collection, you are prompted to install an agent on your tenant.

If you have an agent already configured, select the appropriate Operating System or Platform and click **Continue**.

2. Follow the instructions in the Consul Configuration screen to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

Setup

Information may be found in the Consul documentation.

Objects and Counters for consul

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:	Datapoints:
Consul	Namespace Check ID Service Node	Node IP Node OS Node UUID Node Name Service Name Check Name Service ID Status	Critical Passing Warning

Troubleshooting

Additional information may be found from the Support page.

Couchbase Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Couchbase.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Couchbase.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.



Information may be found in the Couchbase documentation.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
Couchbase Node	Namespace Cluster Couchbase Node Hostname	Node Name Node IP	Memory Free Memory Total
Couchbase Bucket	Namespace Bucket Cluster	Node Name Node IP	Data Used Data Fetches Disk Used Item Count Memory Used Operations Per Second Quota Used

Additional information may be found from the Support page.

CouchDB Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from CouchDB.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose CouchDB.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

🔚 CouchDB	CouchDB Configuration Gathers CouchDB metrics.	
What Operating S	ystem or Platform Are You Using?	Need Help?
🛟 RHEL & CentOS	. . .	
Select existing Age	ent Access Key or create a new one	
Default (405fb5ec-d4cb-440	14-977b-71fa931e1ad3) + Agent Access Key	
*Please ensure that you h	ave a Telegraf Agent in you environment before configuring Show Instruction	ons
Follow Configurat	ion Steps	Need Help?
Copy the contents contents to the /etc	ion Steps below into a new .conf file under the /etc/telegraf/telegraf.d/ c c/telegraf/telegraf.d/cloudinsights-couchdb.conf file.	Need Help? lirectory. For example, copy the
Follow Configurat Copy the contents contents to the /etc ## Read CouchD [[inputs.couch ## Works wit ## Wultiple ## USER-ACTI co (ISCS 1077	ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/c c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. B Stats from one or more servers db]] ch CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats: CON: Provide comma-separated list of couchdb IP(s) and port	Need Help? lirectory. For example, copy the
 Follow Configurat Copy the contents contents to the /etcl contents to the /etcl finguts.couch fill Works with fill w	ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/c c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. B Stats from one or more servers dbj] ch CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats: CON: Provide comma-separated list of couchdb IP(s) and port to the box of the box of the box Hosts from which to read CouchDB stats: CON: Provide comma-separated list of couchdb IP(s) and port to the box of the box of the box of the box Hosts from which to read CouchDB stats: COUCHDB_ADDRESS> with the applicable CouchDB address.	Need Help? lirectory. For example, copy the (a) + Please specify a real machine
 Follow Configurat Copy the contents contents to the /etcl finance of the formation of the formation	<pre>ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/co c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. B Stats from one or more servers heb1] th CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats: CON: Provide comma-separated list of couchdb IP(s) and port COUCHDB_ADDRESS> with the applicable CouchDB address. COUCHDB_PORT> with the applicable CouchDB port.</pre>	Need Help? lirectory. For example, copy the (*) * Please specify a real machine
Follow Configurat Copy the contents contents to the /etc ## Read CouchD [[inputs.couch # Works wit # Multiple ## USER-ACTI COURCE CONSERT_C address, and refrain Replace <insert_c c<="" if="" modify="" td="" the="" url=""><td><pre>ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/co c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. B Stats from one or more servers bbj] ch CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats; CON: Provide comma-separated list of couchdb IP(s) and port COUCHDB_ADDRESS> with the applicable CouchDB address. COUCHDB_PORT> with the applicable CouchDB port. CouchDB monitoring is exposed at different path</pre></td><td>Need Help? lirectory. For example, copy the</td></insert_c>	<pre>ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/co c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. B Stats from one or more servers bbj] ch CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats; CON: Provide comma-separated list of couchdb IP(s) and port COUCHDB_ADDRESS> with the applicable CouchDB address. COUCHDB_PORT> with the applicable CouchDB port. CouchDB monitoring is exposed at different path</pre>	Need Help? lirectory. For example, copy the
Follow Configurat Copy the contents contents to the /etc ## Read CouchD [[inputs.couch ## Works wit ## Multiple ## USER-ACT] Contents Replace <insert_c <insert_c="" address,="" and="" c="" if="" modify="" refrain="" replace="" restart="" td="" telegra<="" the="" url=""><td>ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/c c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. 8 Stats from one or more servers dbj1 th CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats: ON: Provide comma-separated list of couchdb IP(s) and port couchDB_ADDRESS> with the applicable CouchDB address. the form using a loopback address. COUCHDB_PORT> with the applicable CouchDB port. COUCHDB monitoring is exposed at different path f service.</td><td>Need Help? lirectory. For example, copy the</td></insert_c>	ion Steps below into a new conf file under the /etc/telegraf/telegraf.d/c c/telegraf/telegraf.d/cloudinsights-couchdb.conf file. 8 Stats from one or more servers dbj1 th CouchDB stats endpoints out of the box Hosts from which to read CouchDB stats: ON: Provide comma-separated list of couchdb IP(s) and port couchDB_ADDRESS> with the applicable CouchDB address. the form using a loopback address. COUCHDB_PORT> with the applicable CouchDB port. COUCHDB monitoring is exposed at different path f service.	Need Help? lirectory. For example, copy the

Information may be found in the CouchDB documentation.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
CouchDB	Namespace Server	Node IP	Authentication Cache Hits Authentication Cache Miss Database Reads Database Writes Databases Open Open OS Files Max Request Time Min Request Time Httpd Request Methods Copy Httpd Request Methods Delete Httpd Request Methods Get Httpd Request Methods Head Httpd Request Methods Post Httpd Request Methods Post Httpd Request Methods Put Status Codes 200 Status Codes 201 Status Codes 202 Status Codes 301 Status Codes 304 Status Codes 400 Status Codes 403 Status Codes 403 Status Codes 404 Status Codes 409 Status Codes 412 Status Codes 412 Status Codes 500

Additional information may be found from the Support page.

Docker Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Docker.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Docker.

If you haven't configured an Agent for collection, you are prompted to install an agent on your tenant.

If you have an agent already configured, select the appropriate Operating System or Platform and click **Continue**.

2. Follow the instructions in the Docker Configuration screen to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

- docker	Docker Configuration Gathers Docker metrics.	
What Operating Sy	stem or Platform Are You Using?	Need Help?
RHEL & CentOS	~	
Select existing Age	4-977b-71fa931e1ad3) Agent Access Key	
*Please ensure that you ha	we a Telegraf Agent in you environment before configuring. Show Instructions	
Follow Configuration	on Steps	Need Help
Copy the contents to contents to contents to the /etc	on Steps below into a new .conf file under the /etc/telegraf/telegraf.d/ dim /telegraf/telegraf.d/cloudinsights-docker.conf file.	Need Help ectory. For example, copy the
Copy the contents to contents to the /etc [[inputs.dock ## Docker E ## Docker E ## To use T unencrypted and S ## To use T	on Steps below into a new conf file under the /etc/telegraf/telegraf.d/ dim /telegraf/telegraf.d/cloudinsights-docker.conf file. ker]] Endpoint TCP, set endpoint = "tcp://[ip]:[port]". By default, Docker cars for encrypted environment variables (ie, docker-machine), set endpoint = "i	Need Help ectory. For example, copy the uses port 2375 for ENV'''
Copy the contents b contents to the /etc [[inputs.dock ## Docker] ## To use] unencrypted and 2 ## To use a #	<pre>on Steps pelow into a new .conf file under the /etc/telegraf/telegraf.d/ dir /telegraf/telegraf.d/cloudinsights-docker.conf file. ker]] Endpoint TCP, set endpoint = "tcp://[ip]:[port]". By default, Docker 2376 for encrypted environment variables (ie, docker-machine), set endpoint = "i</pre>	Need Help ectory. For example, copy the uses port 2375 for
 Copy the contents be contents to the /etc Copy the contents be contents to the /etc ([inputs.docker B ## Docker B ## To use I ## I ## To use I ## I	on Steps below into a new conf file under the /etc/telegraf/telegraf.d/ dim /telegraf/telegraf.d/cloudinsights-docker.conf file. ker]] Endpoint TCP, set endpoint = "tcp://[ip]:[port]". By default, Docker 276 for encrypted environment variables (ie, docker-machine), set endpoint = "[DOCKER_ENDPOINT> with the applicable Docker endpoint. ' if needed for server disambiguation (to avoid name clashes).	Need Help ectory. For example, copy the uses port 2375 for
 Copy the contents b contents to the /etc Copy the contents b contents to the /etc [[inputs.dock ## To use 1 ## To use 1 unencrypted and 2 ## To use 2 Replace <insert_d< li=""> Modify 'Namespace Restart the Telegraf </insert_d<>	<pre>on Steps pelow into a new .conf file under the /etc/telegraf/telegraf.d/ dim /telegraf/telegraf.d/cloudinsights-docker.conf file. ker]] Endpoint TCP, set endpoint = "tcp://[ip]:[port]". By default, Docker anvironment variables (ie, docker-machine), set endpoint = "(</pre>	Need Help ectory. For example, copy the uses port 2375 for ENV"

Setup

The Telegraf input plugin for Docker collects metrics through a specified UNIX socket or a TCP endpoint.

Compatibility

Configuration was developed against Docker version 1.12.6.

Setting Up

Accessing Docker through a UNIX socket

If the Telegraf agent is running on baremetal, add the telegraf Unix user to the docker Unix group by running the following:

```
sudo usermod -aG docker telegraf
```

If the Telegraf agent is running within a Kubernetes pod, expose the Docker Unix socket by mapping the socket into the pod as a volume and then mounting that volume to /var/run/docker.sock. For example, add the following to the PodSpec:

```
volumes:
...
- name: docker-sock
hostPath:
path: /var/run/docker.sock
type: File
```

Then, add the following to the Container:

```
volumeMounts:
...
- name: docker-sock
mountPath: /var/run/docker.sock
```

Note that the Data Infrastructure Insights installer provided for the Kubernetes platform takes care of this mapping automatically.

Access Docker through a TCP endpoint

By default, Docker uses port 2375 for unencrypted access and port 2376 for encrypted access.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
Docker Engine	Namespace Docker Engine	Node Name Node IP Node UUID Node OS Kubernetes Cluster Docker Version Unit	Memory Containers Containers Paused Containers Running Containers Stopped CPUs Go Routines Images Listener Events Used File Descriptors Data Available Data Total Data Used Metadata Available Metadata Total Metadata Used Pool Blocksize

Object:	Identifiers:	Attributes:	Datapoints:
Docker Container	Namespace Container Name Docker Engine	Autorities.Kubernetes ContainerHashKubernetes ContainerPortsKubernetes ContainerRestart CountKubernetes ContainerTermination MessagePathKubernetes ContainerTermination MessagePolicyKubernetes PodTermination Grace PeriodContainer ImageContainer StatusContainer VersionNode NameKubernetes Container LogPathKubernetes Container LogPathKubernetes Container LogPathKubernetes Container LogPathKubernetes Pod NameKubernetes Pod NameKubernetes Pod NameKubernetes Pod UIDKubernetes Sandbox IDNode IPNode UUIDDocker VersionKubernetes IO ConfigSeenKubernetes DescriptionKubernetes Display NameOpenShift IO SCCKubernetes Display NameOpenShift TagsKompose ServicePod Template HashController Revision HashPod Template GenerationLicenseSchema Build DateSchema URLSchema VersionSchema Version<	Memory Active Anonymous Memory Active File Memory Cache Memory Hierarchical Limit Memory Inactive Anonymous Memory Inactive File Memory Limit Memory Mapped File Memory Mapped File Memory Page Fault Memory Page Major Fault Memory Paged Out Memory Paged Out Memory Resident Set Size Memory Resident Set Size Huge Memory Total Active Anonymous Memory Total Active Anonymous Memory Total Inactive Anonymous Memory Total Inactive Anonymous Memory Total Inactive Anonymous Memory Total Page Fault Memory Total Page Fault Memory Total Page Major Fault Memory Total Paged In Memory Total Paged Out Memory Total Resident Set Size Memory Total Resident Set Size Memory Total Unevictable Memory Usage Memory Usage Memory Usage Memory Usage Percent Exit Code OOM Killed PID Started At Failing Streak
		Customer Pod	

Object:	Identifiers:	Attributes:	Datapoints:
Object: Docker Container Block IO	Identifiers: Namespace Container Name Device Docker Engine	Attributes: Kubernetes Container Hash Kubernetes Container Ports Kubernetes Container Restart Count Kubernetes Container Termination Message Path Kubernetes Container Termination Message Policy Kubernetes Pod Termination Grace Period Container Image Container Status Container Version Node Name Kubernetes Container Log Path Kubernetes Container Log Path Kubernetes Pod Name Kubernetes Pod Name Kubernetes Pod Name Kubernetes Pod Name Kubernetes Pod VID Kubernetes Pod UID Kubernetes Pod UID Node IP Node UUID Docker Version Kubernetes Config Seen Kubernetes Config Seen Kubernetes Description Kubernetes Display Name OpenShift SCC Kubernetes Display Name OpenShift Tags Schema Schema Version Pod Template Generation Kompose Service Schema Build Date Schema License Schema Name Schema Vendor Customer Pod Kubernetes StatefulSet Pod Name Tenant Webconsole	Datapoints: IO Service Bytes Recursive Read IO Service Bytes Recursive Sync IO Service Bytes Recursive Total IO Serviced Recursive Async IO Serviced Recursive Read IO Serviced Recursive Sync IO Serviced Recursive Write
		License Vendor	

Object:	Identifiers:	Attributes:	Datapoints:
Docker Container Network	Namespace Container Name Network Docker Engine	Container Image Container Status Container Version Node Name Node IP Node UUID Node OS K8s Cluster Docker Version Container ID	RX Dropped RX Bytes RX Errors RX Packets TX Dropped TX Bytes TX Errors TX Packets

Object:	Identifiers:	Attributes:	Datapoints:
Docker Container CPU	Namespace Container Name CPU Docker Engine	Kubernetes Container Hash Kubernetes Container Ports Kubernetes Container Restart Count Kubernetes Container Termination Message Path Kubernetes Container Termination Message Policy Kubernetes Pod Termination Grace Period Kubernetes Config Seen Kubernetes Config Seen Kubernetes Config Seen Kubernetes Config Source OpenShift SCC Container Image Container Version Node Name Kubernetes Container Log Path Kubernetes Container Log Path Kubernetes Container Log Path Kubernetes Pod Name Kubernetes Pod Name Kubernetes Pod Name Kubernetes Pod VID Kubernetes Pod UID Node IP Node UUID Node OS Kubernetes Cluster Docker Version Kubernetes Description Kubernetes Display Name OpenShift Tags Schema Version Pod Template Hash Controller Revision Hash Pod Template Generation Kompose Service Schema Build Date Schema Vendor Customer Pod Kubernetes StatefulSet Pod Name Tanant	Throttling Periods Throttling Throttled Periods Throttling Throttled Time Usage In Kernel Mode Usage Percent Usage System Usage Total
		Webconsole Build Date	

Problem:	Try this:
I do not see my Docker metrics in Data Infrastructure Insights after following the instructions on the configuration page.	Check the Telegraf agent logs to see if it reports the following error:
	E! Error in plugin [inputs.docker]: Got permission denied while trying to connect to the Docker daemon socket
	If it does, take the necessary steps to provide the Telegraf agent access to the Docker Unix socket as specified above.

Additional information may be found from the Support page.

Elasticsearch Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Elasticsearch.

1. From **Observability > Collectors**, click **+Data Collector**. Choose Elasticsearch.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.



Information may be found in the Elasticsearch documentation.

Objects and Counters

Object:	Identifiers:	Attributes:
Elasticsearch Cluster	Namespace Cluster	Node IP Node Name Cluster Status

Object:	Identifiers:	Attributes:
Elasticsearch Node	Namespace Cluster ES Node ID ES Node IP ES Node	Zone ID

Additional information may be found from the Support page.

Flink Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Flink.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Flink.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

Contract Flink	Flink Configuration Gathers Flink metrics.	
What Operating S	System or Platform Are You Using?	Need Help?
H Windows	*	
Select existing Ag	gent Access Key or create a new one	
Default (405fb5ec-d4cb-44	104-977b-71fa931e1ad3) Agent Access Key	
*Please ensure that you h	have a Telegraf Agent in you environment before configuring. Show Instructions	
Follow Configurat	tion Steps	Need Help?
 Install Jolokia on y document. 	our Flink JobManager(s) and Flink Task Manager(s). For details refer to) the following
2 Copy the contents copy the contents ## ######## ## JobManag ## ######## [[inputs.jo ## USER-A	s below into a new .conf file under the C:\Program Files\telegraf\telegra to the C:\Program Files\telegraf\telegraf.d\cloudinsights-flink.conf file ************************************	af.d\ folder. For example. ia, add one URL
 Copy the contents copy the contents ## ##################################	s below into a new .conf file under the C\Program Files\telegraf\telegra to the C\Program Files\telegraf\telegraf.d\cloudinsights-flink.conf file ************************************	af.d\ folder. For example.
 Copy the contents copy the contents <i>## ##################################</i>	s below into a new .conf file under the C\Program Files\telegraf\telegra to the C\Program Files\telegraf\telegraf.d\cloudinsights-flink.conf file ************************************	af.d\ folder. For example,
 Copy the contents copy the contents <i>tr terretained</i> <i>tr terretained</i> <i>tr terretained</i> Replace <insert_specify a="" li="" mac<="" real=""> Replace <insert_specify a="" li="" mac<="" real=""> Replace <insert_specify a="" li="" mac<="" real=""> </insert_specify></insert_specify></insert_specify>	s below into a new .conf file under the C\Program Files\telegraf\telegraf to the C\Program Files\telegraf\telegraf.d\cloudinsights-flink.conf file ************************************	af.d\ folder. For example,
 Copy the contents copy the contents <i>tr ####################################</i>	s below into a new .conf file under the C\Program Files\telegraf\telegraf to the C\Program Files\telegraf\telegraf.d\cloudinsights-flink.conf file ************************************	af d\ folder. For example,
 Copy the contents copy the contents <i>tr treatments</i> <i>tr treatment</i> <i>Treatment</i> <i>Treatment</i> <i>Replace <insert_specify a="" i="" mac<="" real=""></insert_specify></i> <i>Replace <insert_specify a="" i="" mac<="" real=""></insert_specify></i> <i>Replace <insert_< i=""></insert_<></i> 	s below into a new .conf file under the C.\Program Files\telegraf\telegraf to the C.\Program Files\telegraf\telegraf.dcloudinsights-flink.conf file """"""""""""""""""""""""""""""""""""	af d\ folder. For example,
 Copy the contents copy the contents <i>## ##################################</i>	s below into a new .conf file under the C\Program Files\telegraf\telegraf to the C\Program Files\telegraf\telegraf.(cloudinsights-flink.conf file """"""""""""""""""""""""""""""""""""	af.d\ folder. For example.

A full Flink deployment involves the following components:

JobManager: The Flink primary system. Coordinates a series of TaskManagers. In a High Availability setup, system will have more than one JobManager.

TaskManager: This is where Flink operators are executed.

The Flink plugin is based on the telegraf's Jolokia plugin. As such as a requirement to gather info from all Flink components, JMX needs to be configured and exposed via Jolokia on all components.

Compatibility

Configuration was developed against Flink version 1.7.0.

Setting Up

Jolokia Agent Jar

For all individual components, a version the Jolokia agent jar file must be downloaded. The version tested against was Jolokia agent 1.6.0.

Instructions below assume that downloaded jar file (jolokia-jvm-1.6.0-agent.jar) is placed under location '/opt/flink/lib/'.

JobManager

To configure JobManager to expose the Jolokia API, you can setup the following environment variable on your nodes then restart the JobManager:

```
export FLINK_ENV_JAVA_OPTS="-javaagent:/opt/flink/lib/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=0.0.0.0"
```

You can choose a different port for Jolokia (8778). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin.

TaskManager

To configure TaskManager(s) to expose the Jolokia API, you can setup the following environment variable on your nodes then restart the TaskManager:

```
export FLINK_ENV_JAVA_OPTS="-javaagent:/opt/flink/lib/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=0.0.0.0"
```

You can choose a different port for Jolokia (8778). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
Flink Task Manager	Cluster Namespace Server	Node Name Task Manager ID Node IP	Network Available Memory Segments Network Total Memory Segments Garbage Collection PS MarkSweep Count Garbage Collection PS MarkSweep Time Garbage Collection PS Scavenge Count Garbage Collection PS Scavenge Time Heap Memory Committed Heap Memory Init Heap Memory Init Heap Memory Max Heap Memory Used Thread Count Daemon Thread Count Peak Thread Count Thread Count Total Started
Flink Job	Cluster Namespace server Job ID	Node Name Job Name Node IP Last Checkpoint External Path Restarting Time	Downtime Full Restarts Last Checkpoint Alignment Buffered Last Checkpoint Duration Last Checkpoint Size Number of Completed Checkpoints Number of Failed Checkpoints Number of in Progress Checkpoints Number of Checkpoints Uptime

Object:	Identifiers:	Attributes:	Datapoints:
Flink Job Manager	Cluster Namespace Server	Node Name Node IP	Garbage Collection PS MarkSweep Count Garbage Collection PS MarkSweep Time Garbage Collection PS Scavenge Count Garbage Collection PS Scavenge Time Heap Memory Committed Heap Memory Used Heap Memory Max Heap Memory Used Number Registered Task Managers Number Running Jobs Task Slots Available Task Slots Total Thread Count Daemon Thread Count Peak Thread Count Total Started

Object:	Identifiers:	Attributes:	Datapoints:
Flink Task	Cluster Namespace Job ID Task ID	Server Node Name Job Name Sub Task Index Task Attempt ID Task Attempt Number Task Manager ID Node IP Current Input Watermark	Buffers In Pool Usage Buffers In Queue Length Buffers Out Pool Usage Buffers Out Queue Length Number Buffers In Local Per Second Count Number Buffers In Local Per Second Rate Number Buffers In Remote Number Buffers In Remote Per Second Count Number Buffers In Remote Per Second Rate Number Buffers Out Number Buffers Out Per Second Count Number Buffers Out Per Second Count Number Buffers Out Per Second Rate Number Bytes In Local Per Second Count Number Bytes In Local Per Second Rate Number Bytes In Local Per Second Rate Number Bytes In Remote Per Second Rate Number Bytes In Remote Per Second Rate Number Bytes In Remote Per Second Count Number Bytes In Remote Per Second Count Number Bytes Out Per Second Count Number Bytes Out Per Second Rate Number Bytes Out Per Second Count Number Bytes Out Per Second Count Number Bytes Out Per Second Rate Number Bytes Out Per Second Count Number Bytes Out Per Second Rate Number Records In Per Second Rate Number Records In Per Second Rate Number Records In Per Second Rate Number Records Out Per Second Count Number Records Out Per Second Rate Number Records Out Per Second Count Number Records Out Per

Object:	Identifiers:	Attributes:	Datapoints:
Flink Task Operator	Cluster Namespace Job ID Operator ID Task ID	Server Node Name Job Name Operator Name Sub Task Index Task Attempt ID Task Attempt Number Task Manager ID Node IP	Current Input Watermark Current Output Watermark Number Records In Per Second Count Number Records In Per Second Rate Number Records Out Per Second Count Number Records Out Per Second Count Number Records Out Per Second Rate Number Late Records Dropped Assigned Partitions Bytes Consumed Rate Commit Latency Avg Commit Latency Avg Commit Latency Max Commits Failed Connection Close Rate Connection Close Rate Connection Count Connection Count Connection Creation Rate Count Fetch Latency Avg Fetch Latency Max Fetch Size Avg Fetch Size Avg Fetch Size Max Fetch Throttle Time Avg Fetch Throttle Time Avg Fetch Throttle Time Avg Fetch Throttle Time Max Heartbeat Rate Incoming Byte Rate IO Ratio IO Wait Time Avg (ns) Join Rate Join Time Avg (ns) Join Rate Join Time Avg Last Heartbeat Ago Network IO Rate Outgoing Byte Rate Records Lag Max Records per Request Avg Request Size Avg Request Size Avg Request Size Avg Request Size Max Response Rate Select Rate Sync Rate Sync Rate Sync Rate

Additional information may be found from the Support page.

Hadoop Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Hadoop.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Hadoop.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

(;-hedaap	Hadoop Configuration Gathers Hadoop metrics.	
What Operating Sy	stem or Platform Are You Using?	Need Help
🛞 Ubuntu & Debian	· · · · · · · · · · · · · · · · · · ·	
Select existing Age	ent Access Key or create a new one	

Follow Configuration Steps	Need Help?
Install Jolokia on your Hadoop NameNode, Secondary NameNode, DataNod NodeManager(s) and JobHistoryServer. For details refer to the following doci	le(s). ResourceManager, ument.
Copy the contents below into a new .conf file under the /etc/telegraf/telegraf contents to the /etc/telegraf/telegraf.d/cloudinsights-hadoop.conf file.	f.d/ directory. For example, copy the
######################################	jolokia
Replace <insert_hadoop_namenode_address> with the applicable Has specify a real machine address, and refrain from using a loopback address. Replace <insert_jolokia_port> with the NameNode's assigned Jolokia port.</insert_jolokia_port></insert_hadoop_namenode_address>	adoop NameNode address. Please Replace corresponding
Replace <insert_hadoop_secondarynamenode_address> with the NameNode address. Please specify a real machine address, and refrain from corresponding <insert_jolokia_port> with the Secondary NameNode's a</insert_jolokia_port></insert_hadoop_secondarynamenode_address>	applicable Hadoop Secondary using a loopback address. Replace assigned Jolokia port.
Replace <insert_hadoop_datanode_address> with the applicable Had specify a real machine address, and refrain from using a loopback address. R <insert_jolokia_port> with the DataNode's assigned Jolokia port.</insert_jolokia_port></insert_hadoop_datanode_address>	doop DataNode address(es). Please Replace corresponding
Replace <insert_hadoop_resourcemanager_address> with the app address. Please specify a real machine address, and refrain from using a loop corresponding <insert_jolokia_port> with the ResourceManager's assig</insert_jolokia_port></insert_hadoop_resourcemanager_address>	olicable Hadoop ResourceManager oback address. Replace ned Jolokia port.
Replace <insert_hadoop_nodemanager_address> with the applicab address(es). Please specify a real machine address, and refrain from using a I corresponding <insert_jolokia_port> with the NodeManager's assigned</insert_jolokia_port></insert_hadoop_nodemanager_address>	le Hadoop NodeManager oopback address. Replace Jolokia port
Replace <insert_hadoop_jobhistoryserver_address> with the appl address. Please specify a real machine address, and refrain from using a loop corresponding <insert_jolokia_port> with the Job History Server's assign</insert_jolokia_port></insert_hadoop_jobhistoryserver_address>	icable Hadoop Job History Server oback address. Replace ned Jolokia port.
Modify 'Namespace' if needed for server disambiguation (to avoid name clas	shes).
10 Modify 'Cluster' if needed for Hadoop cluster designation.	
Restart the Telegraf service.	
systemetl restart telegraf	C III

A full Hadoop deployment involves the following components:

• NameNode: The Hadoop Distributed File System (HDFS) primary system. Coordinates a series of DataNodes.

- Secondary NameNode: a warm failover for the main NameNode. In Hadoop the promotion to NameNode does not occur automatically. Secondary NameNode gathers information from NameNode to be ready to be promoted when needed.
- DataNode: Actual owner for data.
- ResourceManager: The compute primary system (Yarn). Coordinates a series of NodeManagers.
- NodeManager: The resource for compute. Actual location for running of applications.
- JobHistoryServer: Responsible for servicing all job history related requests.

The Hadoop plugin is based on the telegraf's Jolokia plugin. As such as a requirement to gather info from all Hadoop components, JMX needs to be configured and exposed via Jolokia on all components.

Compatibility

Configuration was developed against Hadoop version 2.9.2.

Setting Up

Jolokia Agent Jar

For all individual components, a version the Jolokia agent jar file must be downloaded. The version tested against was Jolokia agent 1.6.0.

Instructions below assume that downloaded jar file (jolokia-jvm-1.6.0-agent.jar) is placed under location '/opt/hadoop/lib/'.

NameNode

To configure NameNode to expose the Jolokia API, you can setup the following in <HADOOP_HOME>/etc/hadoop/hadoop-env.sh:

```
export HADOOP_NAMENODE_OPTS="$HADOOP_NAMENODE_OPTS
-javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0
-agent.jar=port=7800,host=0.0.0.0 -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8000
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.password.file=$HADOOP_HOME/conf/jmxremote.p
assword"
You can choose a different port for JMX (8000 above) and Jolokia (7800).
If you have an internal IP to lock Jolokia onto you can replace the "catch
all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from
the telegraf plugin. You can use the option '-
Dcom.sun.management.jmxremote.authenticate=false' if you don't want to
authenticate. Use at your own risk.
```

Secondary NameNode

To configure the Secondary NameNode to expose the Jolokia API, you can setup the following in <HADOOP_HOME>/etc/hadoop/hadoop-env.sh:
```
export HADOOP_SECONDARYNAMENODE_OPTS="$HADOOP_SECONDARYNAMENODE_OPTS
-javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0
-agent.jar=port=7802,host=0.0.0.0 -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8002
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.password.file=$HADOOP_HOME/conf/jmxremote.p
assword"
You can choose a different port for JMX (8002 above) and Jolokia (7802).
If you have an internal IP to lock Jolokia onto you can replace the "catch
all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from
the telegraf plugin. You can use the option '-
Dcom.sun.management.jmxremote.authenticate=false' if you don't want to
authenticate. Use at your own risk.
```

DataNode

To configure the DataNodes to expose the Jolokia API, you can setup the following in <HADOOP_HOME>/etc/hadoop/hadoop-env.sh:

```
export HADOOP_DATANODE_OPTS="$HADOOP_DATANODE_OPTS
-javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0
-agent.jar=port=7801,host=0.0.0.0 -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8001
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.password.file=$HADOOP_HOME/conf/jmxremote.p
assword"
You can choose a different port for JMX (8001 above) and Jolokia (7801).
If you have an internal IP to lock Jolokia onto you can replace the "catch
all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from
the telegraf plugin. You can use the option '-
Dcom.sun.management.jmxremote.authenticate=false' if you don't want to
authenticate. Use at your own risk.
```

ResourceManager

To configure the ResourceManager to expose the Jolokia API, you can setup the following in <HADOOP_HOME>/etc/hadoop/hadoop-env.sh:

```
export YARN_RESOURCEMANAGER_OPTS="$YARN_RESOURCEMANAGER_OPTS
-javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0
-agent.jar=port=7803,host=0.0.0.0 -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8003
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.password.file=$HADOOP_HOME/conf/jmxremote.p
assword"
You can choose a different port for JMX (8003 above) and Jolokia (7803).
If you have an internal IP to lock Jolokia onto you can replace the "catch
all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from
the telegraf plugin. You can use the option '-
Dcom.sun.management.jmxremote.authenticate=false' if you don't want to
authenticate. Use at your own risk.
```

NodeManager

To configure the NodeManagers to expose the Jolokia API, you can setup the following in <HADOOP_HOME>/etc/hadoop/hadoop-env.sh:

```
export YARN_NODEMANAGER_OPTS="$YARN_NODEMANAGER_OPTS
-javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0
-agent.jar=port=7804,host=0.0.0.0 -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8004
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.password.file=$HADOOP_HOME/conf/jmxremote.p
assword"
You can choose a different port for JMX (8004 above) and Jolokia (7804).
If you have an internal IP to lock Jolokia onto you can replace the "catch
all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from
the telegraf plugin. You can use the option '-
Dcom.sun.management.jmxremote.authenticate=false' if you don't want to
authenticate. Use at your own risk.
```

JobHistoryServer

To configure the JobHistoryServer to expose the Jolokia API, you can setup the following in <HADOOP_HOME>/etc/hadoop/hadoop-env.sh:

```
export HADOOP_JOB_HISTORYSERVER_OPTS="$HADOOP_JOB_HISTORYSERVER_OPTS
-javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0
-agent.jar=port=7805,host=0.0.0.0 -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8005
-Dcom.sun.management.jmxremote.password.file=$HADOOP_HOME/conf/jmxremote.p
assword"
You can choose a different port for JMX (8005 above) and Jolokia (7805).
If you have an internal IP to lock Jolokia onto you can replace the "catch
all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from
the telegraf plugin. You can use the option '-
Dcom.sun.management.jmxremote.authenticate=false' if you don't want to
authenticate. Use at your own risk.
```

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:
Hadoop Secondary NameNode	Cluster Namespace Server	Node Name Node IP Compile Info Version
Hadoop NodeManager	Cluster Namespace Server	Node Name Node IP
Hadoop ResourceManager	Cluster Namespace Server	Node Name Node IP
Hadoop DataNode	Cluster Namespace Server	Node Name Node IP Cluster ID Version
Hadoop NameNode	Cluster Namespace Server	Node Name Node IP Transaction ID Last Written Time Since Last Loaded Edits HA State File System State Block Pool ID Cluster ID Compile Info Distinct Version Count Version
Hadoop JobHistoryServer	Cluster Namespace Server	Node Name Node IP

Troubleshooting

Additional information may be found from the Support page.

HAProxy Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from HAProxy.

Installation

1. From Observability > Collectors, click +Data Collector. Choose HAProxy.

Select the Operating System or Platform on which the Telegraf agent is installed.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

HAPROXY	HAProxy Configuration Gathers HAProxy metrics.	
What Operating Sy	ystem or Platform Are You Using?	Need Help?
🕘 Ubuntu & Debian		
Select existing Age	ent Access Key or create a new one	
Default (405fb5ec-d4cb-4404	4-977b-71fa931e1ad3) * + Agent Access Key	
*Please ensure that you ha	ave a Telegraf Agent in you environment before configuring Show Instru	ctions
Follow Configurati	on Steps	Need Help
Ensure that the HAI the following docur	Proxy system you're going to gather metrics on has 'stats er ment.	nable' option. For details refer to
2 Copy the contents to contents to the /etc	below into a new .conf file under the /etc/telegraf/telegraf.c :/telegraf/telegraf.d/cloudinsights-haproxy.conf file.	I/ directory. For example, copy the
# Read metrics [[inputs.hapro: ## An array o ## with optio ## Wake sure	of HAProxy, via socket or HTTP stats page xy]] of address to gather stats about. Specify an ip on hostn onal port. ie localhost. 10.10.3.33:1936. etc.	ame
	you specify the complete path to the stats endpoint	
3 Replace <insert_h machine address, a</insert_h 	you specify the complete path to the stats endpoint	ddress. Please specify a real
3 Replace <insert_h machine address, a Replace <insert_h< th=""><td>you specify the complete path to the stats endpoint HAPROXY_ADDRESS> with the applicable HAProxy server a nd refrain from using a loopback address. HAPROXY_PORT> with the applicable HAProxy server port.</td><td>ddress. Please specify a real</td></insert_h<></insert_h 	you specify the complete path to the stats endpoint HAPROXY_ADDRESS> with the applicable HAProxy server a nd refrain from using a loopback address. HAPROXY_PORT> with the applicable HAProxy server port.	ddress. Please specify a real
 3 Replace <insert_h machine address, a</insert_h 6 Replace <insert_h< li=""> 5 Modify the 'haproxy </insert_h<>	you specify the complete path to the stats endpoint APROXY_ADDRESS> with the applicable HAProxy server a nd refrain from using a loopback address. APROXY_PORT> with the applicable HAProxy server port. /?stats' path in accordance to the HAProxy server configura	ddress. Please specify a real
 3 Replace <insert_h machine address, a</insert_h 4 Replace <insert_h< li=""> 5 Modify the 'haproxy 6 Modify 'username' a </insert_h<>	you specify the complete path to the stats endpoint HAPROXY_ADDRESS> with the applicable HAProxy server a nd refrain from using a loopback address. HAPROXY_PORT> with the applicable HAProxy server port. /?stats' path in accordance to the HAProxy server configura and 'password' in accordance to the HAProxy server configura	ddress. Please specify a real tion. uration (if credentials are required).
 3 Replace <insert_h machine address, a</insert_h 4 Replace <insert_h< li=""> 5 Modify the 'haproxy 6 Modify 'username' a 7 Modify 'Namespace </insert_h<>	APROXY_ADDRESS> with the applicable HAProxy server a nd refrain from using a loopback address. APROXY_PORT> with the applicable HAProxy server port. /?stats' path in accordance to the HAProxy server configura and 'password' in accordance to the HAProxy server configura	ddress. Please specify a real tion. uration (if credentials are required). es).
 3 Replace <insert_h machine address, at</insert_h 4 Replace <insert_h< li=""> 5 Modify the 'haproxy 6 Modify 'username' at 7 Modify 'Namespace 8 Restart the Telegraf </insert_h<>	AAPROXY_ADDRESS> with the applicable HAProxy server a nd refrain from using a loopback address. APROXY_PORT> with the applicable HAProxy server port. /?stats' path in accordance to the HAProxy server configura and 'password' in accordance to the HAProxy server configura if needed for server disambiguation (to avoid name clash f service.	ddress. Please specify a real tion. uration (if credentials are required). es).

Setup

Telegraf's plugin for HAProxy relies on HAProxy Stats enablement. This is a configuration built into HAProxy but it is not enabled out of the box. When enabled, HAProxy will expose an HTML endpoint that can be viewed

on your browser or scraped for extraction of status of all HAProxy configurations.

Compatibility:

Configuration was developed against HAProxy version 1.9.4.

Setting Up:

To enable stats, edit your haproxy configuration file and add the the following lines after the 'defaults' section, using your own user/password and/or haproxy URL:

```
stats enable
stats auth myuser:mypassword
stats uri /haproxy?stats
```

The following is a simplified example configuration file with stats enabled:

```
global
  daemon
  maxconn 256
defaults
 mode http
 stats enable
 stats uri /haproxy?stats
 stats auth myuser:mypassword
 timeout connect 5000ms
 timeout client 50000ms
  timeout server 50000ms
frontend http-in
 bind *:80
 default backend servers
frontend http-in9080
 bind *:9080
  default backend servers_2
backend servers
  server server1 10.128.0.55:8080 check ssl verify none
  server server2 10.128.0.56:8080 check ssl verify none
backend servers 2
  server server3 10.128.0.57:8080 check ssl verify none
  server server4 10.128.0.58:8080 check ssl verify none
```

For complete and up to date instructions, see the HAProxy documentation.

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:	Datapoints:
HAProxy Frontend	Namespace Address Proxy	Node IP Node Name Proxy ID Mode Process id Sessions Rate Limit Server id Sessions Limit Status	Bytes In Bytes Out Cache Hits Cache Lookups Compression Bytes Bypassed Compression Bytes In Compression Bytes Out Compression Responses Connection Rate Connection Rate Max Connections Total Requests Denied by Connection Rule Requests Denied by Security Concerns Responses Denied by Security Concerns Requests Denied by Security Concerns Requests Denied by Session Rule Requests Errors Responses 1xx Responses 1xx Responses 2xx Responses 4xx Responses 5xx Responses 5xx Responses 5xx Responses 5xx Responses 5xx Responses 5xx Responses Cother Requests Intercepted Sessions Rate Sessions Rate Requests Rate Max Requests Rate Max Requests Rate Max Requests Total Sessions Max Sessions Max Sessions Total Requests Rewrites

Object:	Identifiers:	Attributes:	Datapoints:
HAProxy Backend	Namespace Address Proxy	Node IP Node Name Proxy ID Last Change Time Last Session Time Mode Process id Server id Sessions Limit Status Weight	Active Servers Backup Servers Bytes In Bytes Out Cache Hits Cache Lookups Check Downs Client Aborts Compression Bytes Bypassed Compression Bytes In Compression Bytes Out Compression Bytes Out Compression Responses Connection Average Time Downtime Total Requests Denied by Security Concerns Responses Denied by Security Concerns Responses Denied by Security Concerns Responses Ixx Responses 1xx Responses 1xx Responses 2xx Responses 4xx Responses 4xx Responses 5xx Responses 5xx Responses 5xx Responses 5xx Responses Other Server Selected Total Queue Average Time Sessions per Second Sessions per Second Max Requests Total Connection Reuse Response Time Average Sessions Max Server Transfer Aborts Sessions Total Sessions Total Time Average Requests Redispatches Requests Retries Requests Retries Requests Retries Requests Retries Requests Retries Requests Retries Requests Retries

Troubleshooting

Additional information may be found from the Support page.

JVM Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from JVM.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose JVM.

Select the Operating System or Platform on which the Telegraf agent is installed.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.



Setup

Information may be found in JVM documentation.

Objects and Counters

The following objects and their counters are collected:

Troubleshooting

Additional information may be found from the Support page.

Kafka Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Kafka.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Kafka.

Select the Operating System or Platform on which the Telegraf agent is installed.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

20 Kalka	Kafka Configuration Gathers Kafka metrics.	
What Operating S	ystem or Platform Are You Using?	Need Help
Windows		
Select existing Age	ent Access Key or create a new one	
Default (405fb5ec-d4cb-440)4-977b-71fa931e1ad3) + Agent Access Key	
*Please ensure that you h	ave a Telegraf Agent in you environment before configuring Show Instructio	ns
Follow Configurat	ion Steps	Need Help
 Install Jolokia on yo Copy the contents copy the contents 	bur Kafka brokers. For details refer to the following document. below into a new .conf file under the C.\Program Files\telegraf to the C.\Program Files\telegraf\telegraf.d\cloudinsights-kafka.c	telegraf.d\ folder. For example, :onf file.
2		
<pre># Read JMX met [[inputs.jolok ## USER-ACTI each broker in y ## Please sp }</pre>	rrics through Jolokia tia2_agent]] CON: Provide address(es) of kafka broker(s), port for joloki Your cluster Pecify actual machine IP address, and refrain from using a l	a, add one URL for oopback address (i.e.
# Read JMX met [[inputs.jolok ## USER-ACT] each broker in y ## Please sp ?	crics through Jolokia cia2_agent]] CON: Provide address(es) of kafka broker(s), port for joloki rour cluster lecify actual machine IP address, and refrain from using a l rock KAFKA_BROKER_ADDRESS> with the applicable Kafka broker and refrain from using a loopback address.	a, add one URL for oopback address (i.e. address Please specify a real
# Read JMX met [[inputs.jolok ## USER-ACT] each broker in y ## Please sp ??	trics through Jolokia (ia2_agent]] (ON: Provide address(es) of kafka broker(s), port for joloki our cluster recify actual machine IP address, and refrain from using a l (AFKA_BROKER_ADDRESS> with the applicable Kafka broker and refrain from using a loopback address.	a, add one URL for oopback address (i.e. address Please specify a real
 # Read JMX met [[inputs.jolok ## USER-ACT] each broker in y ## Please sp ?	trics through Jolokia tia2_agent]] ION: Provide address(es) of kafka broker(s), port for joloki your cluster recify actual machine IP address, and refrain from using a l CARAN KAFKA_BROKER_ADDRESS> with the applicable Kafka broker and refrain from using a loopback address. IOLOKIA_PORT> with the applicable Kafka broker jolokia port. e' if needed for server disambiguation (to avoid name clashes).	a, add one URL for oopback address (i.e. address. Please specify a real
 # Read JNX met [[inputs.jolok ## USER-ACT] each broker in y ## Please sp ?? 3 Replace <insert_ <br="">machine address, a</insert_> 4 Replace <insert_0< li=""> 5 Modify 'Namespace 6 Modify 'Cluster' if n </insert_0<>	rics through Jolokia (12_sgent]] (ON: Provide address(es) of kafka broker(s), port for joloki your cluster lecify actual machine IP address, and refrain from using a l (AFKA_BROKER_ADDRESS> with the applicable Kafka broker and refrain from using a loopback address. IOLOKIA_PORT> with the applicable Kafka broker jolokia port. e' if needed for server disambiguation (to avoid name clashes). eeded for Kafka cluster designation.	a, add one URL for oopback address (i.e. address Please specify a real
 # Read JNX met [[inputs.jolok ## USER-ACT] each broker in y ## Please sp ??	trics through Jolokia tia2_agent]] ON: Provide address(es) of kafka broker(s), port for joloki four cluster becify actual machine IP address, and refrain from using a l CARKA_BROKER_ADDRESS> with the applicable Kafka broker and refrain from using a loopback address. DOLOKIA_PORT> with the applicable Kafka broker jolokia port. e' if needed for server disambiguation (to avoid name clashes). eeded for Kafka cluster designation. f service.	a, add one URL for oopback address (i.e. address. Please specify a real

Setup

The Kafka plugin is based on the telegraf's Jolokia plugin. As such as a requirement to gather info from all Kafka brokers, JMX needs to be configured and exposed via Jolokia on all components.

Compatibility

Configuration was developed against Kafka version 0.11.0.2.

Setting up

All the instructions below assume your install location for kafka is '/opt/kafka'. You can adapt instructions below to reflect your install location.

Jolokia Agent Jar

A version the Jolokia agent jar file must be downloaded. The version tested against was Jolokia agent 1.6.0.

Instructions below assume that the downloaded jar file (jolokia-jvm-1.6.0-agent.jar) is placed under the location //opt/kafka/libs/'.

Kafka Brokers

To configure Kafka Brokers to expose the Jolokia API, you can add the following in <KAFKA_HOME>/bin/kafka-server-start.sh, just before the 'kafka-run-class.sh' call:

```
export JMX_PORT=9999
export RMI_HOSTNAME=`hostname -I`
export KAFKA_JMX_OPTS="-javaagent:/opt/kafka/libs/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=0.0.0.0
-Dcom.sun.management.jmxremote.password.file=/opt/kafka/config/jmxremote.p
assword -Dcom.sun.management.jmxremote.ssl=false
-Djava.rmi.server.hostname=$RMI_HOSTNAME
-Dcom.sun.management.jmxremote.rmi.port=$JMX_PORT"
```

Note that example above is using 'hostname -I' to setup the 'RMI_HOSTNAME' environment variable. In multiple IP machines, this will need to be tweaked to gather the IP you care about for RMI connections.

You can choose a different port for JMX (9999 above) and Jolokia (8778). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk.

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:
Kafka Broker	Cluster Namespace Broker	Node Name Node IP

Troubleshooting

Additional information may be found from the Support page.

Kibana Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Kibana.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Kibana.

Select the Operating System or Platform on which the Telegraf agent is installed.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.



Setup

Information may be found in the Kibana documentation.

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:	Datapoints:
Kibana	Namespace Address	Node IP Node Name Version Status	Concurrent Connections Heap Max Heap Used Requests per Second Response Time Average Response Time Max Uptime

Troubleshooting

Additional information may be found from the Support page.

Kubernetes Monitoring Operator Installation and Configuration

Data Infrastructure Insights offers the **Kubernetes Monitoring Operator** for Kubernetes collection. Navigate to **Kubernetes > Collectors > +Kubernetes Collector** to deploy a new operator.

Before installing the Kubernetes Monitoring Operator

See the Pre-requisites documentation before installing or upgrading the Kubernetes Monitoring Operator.

Installing the Kubernetes Monitoring Operator



Kubernetes

Deploy NetApp Monitoring Operator Quickly install and configure a Kubernetes Operator to send cluster information to Cloud Insights.

	the reaction of the date of the worke			
EY2024 (vw6NdM)		 + API Access Token 	Production Best Practices 🕜	
tallation Instruction	15			Need Help
ase review the pre-requisit	es for installing the NetApp Kubernetes	Monitoring Operator.		
update an existing operato	r installation please follow these steps.			
Dofino Kubornetas a	luster name and namespace			
Provide the Kubernetes C	uster name and pracify a namespace	for deploying the monitoring con	anananta	
Frovide the Ruberhetes	cluster name and specify a namespace	for deploying the monitoring con	iponents.	
Cluster	Namespace			
clustername	netapp-monitoring			
Devente of the ensure	La MANU GLAS			
Download the opera	tor YAML files			
Download the opera Execute the following do	tor YAML files wnload command in a <i>bash</i> prompt.			
Download the opera Execute the following do Copy Download Comm	tor YAML files wnload command in a <i>bash</i> prompt. nand Snippet			
Download the operat Execute the following do Copy Download Comr	tor YAML files wwnload command in a <i>bash</i> prompt. nand Snippet mmand Snippet			
Download the operat Execute the following do Copy Download Comp E Reveal Download Com	tor YAML files wonload command in a <i>bash</i> prompt. nand Snippet mmand Snippet			

Optional: Upload the operator images to your private repository

By default, the operator pulls container images from the Cloud Insights repository. To use a private repository, download the required images using the Image Pull command. Then upload them to your private repository maintaining the same tags and directory structure. Finally, update the image paths in operator-deployment.yaml and the docker repository settings in operator-config.yaml. For more information review the documentation.



Steps to install Kubernetes Monitoring Operator agent on Kubernetes:

- 1. Enter a unique cluster name and namespace. If you are upgrading from a previous Kubernetes Operator, use the same cluster name and namespace.
- 2. Once these are entered, you can copy the Download Command snippet to the clipboard.
- 3. Paste the snippet into a *bash* window and execute it. The Operator installation files will be downloaded. Note that the snippet has a unique key and is valid for 24 hours.
- 4. If you have a custom or private repository, copy the optional Image Pull snippet, paste it into a *bash* shell and execute it. Once the images have been pulled, copy them to your private repository. Be sure to maintain the same tags and folder structure. Update the paths in *operator-deployment.yaml* as well as the docker repository settings in *operator-config.yaml*.
- 5. If desired, review available configuration options such as proxy or private repository settings. You can read more about configuration options.
- When you are ready, deploy the Operator by copying the kubectl Apply snippet, downloading it, and executing it.
- 7. The installation proceeds automatically. When it is complete, click the Next button.
- 8. When installation is complete, click the *Next* button. Be sure to also delete or securely store the *operator-secrets.yaml* file.

If you are have a custom repository, read about using a custom/private docker repository.

Kubernetes Monitoring Components

Data Infrastructure Insights Kubernetes Monitoring is comprised of four monitoring components:

- Cluster Metrics
- Network Performance and Map (optional)
- Event Logs (optional)
- Change Analysis (optional)

The optional components above are enabled by default for each Kubernetes collector; if you decide you don't need a component for a particular collector, you can disable it by navigating to **Kubernetes > Collectors** and selecting *Modify Deployment* from the collector's "three dots" menu on the right of the screen.

NetApp / Observability / Colle	ctors				
		Data Collectors 👥 21 Acquisition Units 🖸	4 Kubernetes Collectors		
Kubernetes Collectors (13)		View Upgrad	le/Delete Documentation 🛛 🧧	+ Kubernetes Collector 🗧 Fi	lter
Cluster Name 🏌	Status	Operator Version	Network Performance Map	e and Change Analysis	
au-pod	🛕 Outdated	1 .1540.0	1 .347.0	1.162.0	
jks-troublemaker	Latest	1.1579.0	N/A	1.201.0	
oom-test	A Outdated	1 .1555.0	N/A		Modify Deployment

The screen shows the current state of each component and allows you to disable or enable components for that collector as needed.



Upgrading to the latest Kubernetes Monitoring Operator

DII Push-Button upgrades

You can upgrade the Kubernetes Monitoring Operator through the DII Kubernetes Collectors page. Click on the menu next to the cluster you would like to upgrade and select *Upgrade*. The operator will verify the image signatures, perform a snapshot of your current installation and perform the upgrade. Within a few minutes you should see the operator Status progress through Upgrade In Progress to Latest. If you encounter an error you can select the Error status for more details and refer to the Push-Button Upgrades Troubleshooting table below.

Push-Button upgrades with private repositories

If your operator is configured to use a private repository please ensure all images required to run the operator and their signatures are available in your repository. If you encounter an error during the upgrade process for missing images simply add them to your repository and retry the upgrade. To upload the image signatures to your repository please use the cosign tool as follows, making sure to upload signatures for all images specified under 3 Optional: Upload the operator images to your private repository > Image Pull Snippet

```
cosign copy example.com/src:v1 example.com/dest:v1
#Example
cosign copy <DII container registry>/netapp-monitoring:<image version>
<private repository>/netapp-monitoring:<image version>
```

Rolling back to a previously running version

If you upgraded using the push-button upgrades feature and encounter any difficulties with the current version of the operator within seven days of the upgrade, you can downgrade to the previously running version using the snapshot created during the upgrade process. Click the menu next to the cluster you would like to roll back and select *Roll back*.

Manual Upgrades

Determine whether an AgentConfiguration exists with the existing Operator (if your namespace is not the default *netapp-monitoring*, substitute the appropriate namespace):

```
kubectl -n netapp-monitoring get agentconfiguration netapp-ci-monitoring-
configuration
```

If an AgentConfiguration exists:

- Install the latest Operator over the existing Operator.
 - $\circ\,$ Ensure you are pulling the latest container images if you are using a custom repository.

If the AgentConfiguration does not exist:

• Make note of your cluster name as recognized by Data Infrastructure Insights (if your namespace is not the default netapp-monitoring, substitute the appropriate namespace):

```
kubectl -n netapp-monitoring get agent -o
jsonpath='{.items[0].spec.cluster-name}'
```

• Create a backup of the existing Operator (if your namespace is not the default netapp-monitoring, substitute the appropriate namespace):

kubectl -n netapp-monitoring get agent -o yaml > agent backup.yaml

- Uninstall the existing Operator.
- Install the latest Operator.
 - Use the same cluster name.
 - After downloading the latest Operator YAML files, port any customizations found in agent_backup.yaml to the downloaded operator-config.yaml before deploying.
 - Ensure you are pulling the latest container images if you are using a custom repository.

Stopping and Starting the Kubernetes Monitoring Operator

To stop the Kubernetes Monitoring Operator:

```
kubectl -n netapp-monitoring scale deploy monitoring-operator --replicas=0
```

To start the Kubernetes Monitoring Operator:

```
kubectl -n netapp-monitoring scale deploy monitoring-operator --replicas=1
```

Uninstalling

To remove the Kubernetes Monitoring Operator

Note that the default namespace for the Kubernetes Monitoring Operator is "netapp-monitoring". If you have set your own namespace, substitute that namespace in these and all subsequent commands and files.

Newer versions of the monitoring operator can be uninstalled with the following commands:

```
kubectl -n <NAMESPACE> delete agent -l installed-by=nkmo-<NAMESPACE>
kubectl -n <NAMESPACE> delete
clusterrole,clusterrolebinding,crd,svc,deploy,role,rolebinding,secret,sa
-l installed-by=nkmo-<NAMESPACE>
```

If the monitoring operator was deployed in its own dedicated namespace, delete the namespace:

kubectl delete ns <NAMESPACE>

Note: If the first command returns "No resources found", use the following instructions to uninstall older versions of the monitoring operator.

Execute each of the following commands in order. Depending on your current installation, some of these commands may return 'object not found' messages. These messages may be safely ignored.

```
kubectl -n <NAMESPACE> delete agent agent-monitoring-netapp
kubectl delete crd agents.monitoring.netapp.com
kubectl -n <NAMESPACE> delete role agent-leader-election-role
kubectl delete clusterrole agent-manager-role agent-proxy-role agent-
metrics-reader <NAMESPACE>-agent-manager-role <NAMESPACE>-agent-proxy-role
<NAMESPACE>-cluster-role-privileged
kubectl delete clusterrolebinding agent-manager-rolebinding agent-proxy-
rolebinding agent-cluster-admin-rolebinding <NAMESPACE>-agent-manager-
rolebinding <NAMESPACE>-agent-proxy-rolebinding <NAMESPACE>-cluster-role-
binding-privileged
kubectl delete <NAMESPACE>-psp-nkmo
kubectl delete ns <NAMESPACE>
```

If a Security Context Constraint was previously-created:

kubectl delete scc telegraf-hostaccess

About Kube-state-metrics

The NetApp Kubernetes Monitoring Operator installs its own kube-state-metrics to avoid conflict with any other instances.

For information about Kube-State-Metrics, see this page.

Configuring/Customizing the Operator

These sections contain information on customizing your operator configuration, working with proxy, using a custom or private docker repository, or working with OpenShift.

Configuration Options

Most commonly modified settings can be configured in the *AgentConfiguration* custom resource. You can edit this resource before deploying the operator by editing the *operator-config.yaml* file. This file includes commented-out examples of settings. See the list of available settings for the most recent version of the operator.

You can also edit this resource after the operator has been deployed by using the following command:

kubectl -n netapp-monitoring edit AgentConfiguration

To determine if your deployed version of the operator supports AgentConfiguration, run the following command:

kubectl get crd agentconfigurations.monitoring.netapp.com

If you see an "Error from server (NotFound)" message, your operator must be upgraded before you can use the AgentConfiguration.

Configuring Proxy Support

There are two places where you may use a proxy on your tenant in order to install the Kubernetes Monitoring Operator. These may be the same or separate proxy systems:

- Proxy needed during execution of the installation code snippet (using "curl") to connect the system where the snippet is executed to your Data Infrastructure Insights environment
- Proxy needed by the target Kubernetes cluster to communicate with your Data Infrastructure Insights environment

If you use a proxy for either or both of these, in order to install the Kubernetes Operating Monitor you must first ensure that your proxy is configured to allow good communication to your Data Infrastructure Insights environment. If you have a proxy and can access Data Infrastructure Insights from the server/VM from which you wish to install the Operator, then your proxy is likely configured properly.

For the proxy used to install the Kubernetes Operating Monitor, before installing the Operator, set the *http_proxy/https_proxy* environment variables. For some proxy environments, you may also need to set the *no_proxy* environment variable.

To set the variable(s), perform the following steps on your system **before** installing the Kubernetes Monitoring Operator:

- 1. Set the https_proxy and/or http_proxy environment variable(s) for the current user:
 - a. If the proxy being setup does not have Authentication (username/password), run the following command:

export https_proxy=<proxy_server>:<proxy_port>

b. If the proxy being setup does have Authentication (username/password), run this command:

```
<proxy_vserver>:<proxy_password>@<proxy_server>:<proxy_po
rt></proxy_vserver>:<proxy_po
```

For the proxy used for your Kubernetes cluster to communicate with your Data Infrastructure Insights environment, install the Kubernetes Monitoring Operator after reading all of these instructions.

Configure the proxy section of AgentConfiguration in operator-config.yaml before deploying the Kubernetes Monitoring Operator.

```
agent:
  . . .
 proxy:
   server: <server for proxy>
   port: <port for proxy>
   username: <username for proxy>
   password: <password for proxy>
   # In the noproxy section, enter a comma-separated list of
    # IP addresses and/or resolvable hostnames that should bypass
    # the proxy
   noproxy: <comma separated list>
   isTelegrafProxyEnabled: true
    isFluentbitProxyEnabled: <true or false> # true if Events Log enabled
   isCollectorsProxyEnabled: <true or false> # true if Network
Performance and Map enabled
   isAuProxyEnabled: <true or false> # true if AU enabled
  . . .
. . .
```

Using a custom or private docker repository

By default, the Kubernetes Monitoring Operator will pull container images from the Data Infrastructure Insights repository. If you have a Kubernetes cluster used as the target for monitoring, and that cluster is configured to only pull container images from a custom or private Docker repository or container registry, you must configure access to the containers needed by the Kubernetes Monitoring Operator.

Run the "Image Pull Snippet" from the NetApp Monitoring Operator install tile. This command will log into the Data Infrastructure Insights repository, pull all image dependencies for the operator, and log out of the Data Infrastructure Insights repository. When prompted, enter the provided repository temporary password. This command downloads all images used by the operator, including for optional features. See below for which features these images are used for.

Core Operator Functionality and Kubernetes Monitoring

- netapp-monitoring
- ci-kube-rbac-proxy
- ci-ksm
- ci-telegraf
- distroless-root-user

Events Log

- ci-fluent-bit
- ci-kubernetes-event-exporter

Network Performance and Map

• ci-net-observer

Push the operator docker image to your private/local/enterprise docker repository according to your corporate policies. Ensure that the image tags and directory paths to these images in your repository are consistent with those in the Data Infrastructure Insights repository.

Edit the monitoring-operator deployment in operator-deployment.yaml, and modify all image references to use your private Docker repository.

```
image: <docker repo of the enterprise/corp docker repo>/ci-kube-rbac-
proxy:<ci-kube-rbac-proxy version>
image: <docker repo of the enterprise/corp docker repo>/netapp-
monitoring:<version>
```

Edit the AgentConfiguration in operator-config.yaml to reflect the new docker repo location. Create a new imagePullSecret for your private repository, for more details see https://kubernetes.io/docs/tasks/configure-pod-container/pull-image-private-registry/

```
agent:
...
# An optional docker registry where you want docker images to be pulled
from as compared to CI's docker registry
# Please see documentation link here:
xref:{relative_path}task_config_telegraf_agent_k8s.html#using-a-custom-or-
private-docker-repository
dockerRepo: your.docker.repo/long/path/to/test
# Optional: A docker image pull secret that maybe needed for your
private docker registry
dockerImagePullSecret: docker-secret-name
```

OpenShift Instructions

If you are running on OpenShift 4.6 or higher, you must edit the AgentConfiguration in *operator-config.yaml* to enable the *runPrivileged* setting:

```
# Set runPrivileged to true SELinux is enabled on your kubernetes nodes
runPrivileged: true
```

Openshift may implement an added level of security that may block access to some Kubernetes components.

Tolerations and Taints

The *netapp-ci-telegraf-ds*, *netapp-ci-fluent-bit-ds*, and *netapp-ci-net-observer-l4-ds* DaemonSets must schedule a pod on every node in your cluster in order to correctly collect data on all nodes. The operator has been configured to tolerate some well known **taints**. If you have configured any custom taints on your nodes, thus preventing pods from running on every node, you can create a **toleration** for those taints in the *AgentConfiguration*. If you have applied custom taints to all nodes in your cluster, you must also add the necessary tolerations to the operator deployment to allow the operator pod to be scheduled and executed.

Learn More about Kubernetes Taints and Tolerations.

Return to the NetApp Kubernetes Monitoring Operator Installation page

A Note About Secrets

To remove permission for the Kubernetes Monitoring Operator to view secrets cluster-wide, delete the following resources from the *operator-setup.yaml* file before installing:

```
ClusterRole/netapp-ci<namespace>-agent-secret
ClusterRoleBinding/netapp-ci<namespace>-agent-secret
```

If this is an upgrade, also delete the resources from your cluster:

```
kubectl delete ClusterRole/netapp-ci-<namespace>-agent-secret-clusterrole
kubectl delete ClusterRoleBinding/netapp-ci-<namespace>-agent-secret-
clusterrolebinding
```

If Change Analysis is enabled, modify the *AgentConfiguration* or *operator-config.yaml* to uncomment the change-management section and include *kindsTolgnoreFromWatch: "secrets"* under the change-management section. Note the presence and position of single and double quotes in this line.

```
change-management:
...
# # A comma separated list of kinds to ignore from watching from the
default set of kinds watched by the collector
# # Each kind will have to be prefixed by its apigroup
# # Example: '"networking.k8s.io.networkpolicies,batch.jobs",
"authorization.k8s.io.subjectaccessreviews"'
kindsToIgnoreFromWatch: '"secrets"'
...
```

Verifying Kubernetes Monitoring Operator Image Signatures

The image for the operator and all related images it deploys are signed by NetApp. You can manually verify the images before installation using the cosign tool, or configure a Kubernetes admission controller. For more details please see the Kubernetes documentation.

The public key used to verify the image signatures is available in the Monitoring Operator install tile under *Optional: Upload the operator images to your private repository > Image Signature Public Key*

To manually verify an image signature, perform the following steps:

- 1. Copy and run the Image Pull Snippet
- 2. Copy and enter the Repository Password when prompted
- 3. Store the Image Signature Public Key (dii-image-signing.pub in the example)
- 4. Verify the images using cosign. Refer to the following example of cosign usage

```
$ cosign verify --key dii-image-signing.pub --insecure-ignore-sct
--insecure-ignore-tlog <repository>/<image>:<tag>
Verification for <repository>/<image>:<tag> --
The following checks were performed on each of these signatures:
      The cosign claims were validated
      The signatures were verified against the specified public key
[{"critical":{"identity":{"docker-
reference":"<repository>/<image>"},"image":{"docker-manifest-
digest":"sha256:<hash>"},"type":"cosign container image
signature"},"optional":null}]
```

Troubleshooting

Some things to try if you encounter problems setting up the Kubernetes Monitoring Operator:

Problem:	Try this:
I do not see a hyperlink/connection between my	Follow the steps to uninstall the existing Telegraf
Kubernetes Persistent Volume and the corresponding	agent, then re-install the latest Telegraf agent. You
back-end storage device. My Kubernetes Persistent	must be using Telegraf version 2.0 or later, and your
Volume is configured using the hostname of the	Kubernetes cluster storage must be actively
storage server.	monitored by Data Infrastructure Insights.

Problem:	Try this:
l'm seeing messages in the logs resembling the following: E0901 15:21:39.962145 1 reflector.go:178]	These messages may occur if you are running kube- state-metrics version 2.0.0 or above with Kubernetes versions below 1.20.
k8s.io/kube-state-metrics/internal/store/builder.go:352: Failed to list *v1.MutatingWebhookConfiguration: the	To get the Kubernetes version:
server could not find the requested resource E0901 15:21:43.168161 1 reflector.go:178]	kubectl version
Failed to list *v1.Lease: the server could not find the requested resource (get leases.coordination.k8s.io) etc.	To get the kube-state-metrics version:
	kubectl get deploy/kube-state-metrics -o jsonpath='{image}'
	To prevent these messages from happening, users can modify their kube-state-metrics deployment to disable the following Leases:
	mutatingwebhookconfigurations validatingwebhookconfigurations volumeattachments resources
	More specifically, they can use the following CLI argument:
	resources=certificatesigningrequests,configmaps,cron jobs,daemonsets, deployments,endpoints,horizontalpodautoscalers,ingr esses,jobs,limitranges, namespaces,networkpolicies,nodes,persistentvolume claims,persistentvolumes, poddisruptionbudgets,pods,replicasets,replicationcont rollers,resourcequotas, secrets,services,statefulsets,storageclasses
	The default resource list is:
	"certificatesigningrequests,configmaps,cronjobs,daem onsets,deployments, endpoints,horizontalpodautoscalers,ingresses,jobs,lea ses,limitranges, mutatingwebhookconfigurations,namespaces,network policies,nodes, persistentvolumeclaims,persistentvolumes,poddisrupti onbudgets,pods,replicasets, replicationcontrollers,resourcequotas,secrets,services, statefulsets,storageclasses, validatingwebhookconfigurations,volumeattachments"

Problem:	Try this:
I see error messages from Telegraf resembling the following, but Telegraf does start up and run: Oct 11 14:23:41 ip-172-31-39-47 systemd[1]: Started The plugin-driven server agent for reporting metrics into InfluxDB. Oct 11 14:23:41 ip-172-31-39-47 telegraf[1827]: time="2021-10-11T14:23:41Z" level=error msg="failed to create cache directory. /etc/telegraf/.cache/snowflake, err: mkdir /etc/telegraf/.ca che: permission denied. ignored\n" func="gosnowflake.(*defaultLogger).Errorf" file="log.go:120" Oct 11 14:23:41 ip-172-31-39-47 telegraf[1827]: time="2021-10-11T14:23:41Z" level=error msg="failed to open. lgnored. open /etc/telegraf/.cache/snowflake/ocsp_response_cache.j son: no such file or directory\n" func="gosnowflake.(*defaultLogger).Errorf" file="log.go:120" Oct 11 14:23:41 ip-172-31-39-47 telegraf[1827]: 2021- 10-11T14:23:412 I! Starting Telegraf 1.19.3	This is a known issue. Refer to This GitHub article for more details. As long as Telegraf is up and running, users can ignore these error messages.
On Kubernetes, my Telegraf pod(s) are reporting the following error: "Error in processing mountstats info: failed to open mountstats file: /hostfs/proc/1/mountstats, error: open /hostfs/proc/1/mountstats: permission denied"	If SELinux is enabled and enforcing, it is likely preventing the Telegraf pod(s) from accessing the /proc/1/mountstats file on the Kubernetes node. To overcome this restriction, edit the agentconfiguration, and enable the runPrivileged setting. For more details, refer to the OpenShift Instructions.
On Kubernetes, my Telegraf ReplicaSet pod is reporting the following error: [inputs.prometheus] Error in plugin: could not load keypair /etc/kubernetes/pki/etcd/server.crt:/etc/kubernetes/pki/ etcd/server.key: open /etc/kubernetes/pki/etcd/server.crt: no such file or directory	The Telegraf ReplicaSet pod is intended to run on a node designated as a master or for etcd. If the ReplicaSet pod is not running on one of these nodes, you will get these errors. Check to see if your master/etcd nodes have taints on them. If they do, add the necessary tolerations to the Telegraf ReplicaSet, telegraf-rs. For example, edit the ReplicaSet kubectl edit rs telegraf-rs and add the appropriate tolerations to the spec. Then, restart the ReplicaSet pod.

Problem:	Try this:
I have a PSP/PSA environment. Does this affect my monitoring operator?	If your Kubernetes cluster is running with Pod Security Policy (PSP) or Pod Security Admission (PSA) in place, you must upgrade to the latest Kubernetes Monitoring Operator. Follow these steps to upgrade to the current Operator with support for PSP/PSA: 1. Uninstall the previous monitoring operator:
	 netapp-monitoring kubectl delete ns netapp-monitoring kubectl delete crd agents.monitoring.netapp.com kubectl delete clusterrole agent-manager-role agent- proxy-role agent-metrics-reader kubectl delete clusterrolebinding agent-manager- rolebinding agent-proxy-rolebinding agent-cluster- admin-rolebinding
	2. Install the latest version of the monitoring operator.
I ran into issues trying to deploy the Operator, and I have PSP/PSA in use.	 Edit the agent using the following command: kubectl -n <name-space> edit agent</name-space> Mark 'security-policy-enabled' as 'false'. This will disable Pod Security Policies and Pod Security Admission and allow the Operator to deploy. Confirm by using the following commands: kubectl get psp (should show Pod Security Policy removed) kubectl get all -n <namespace> grep -i psp (should show that nothing is found)</namespace>
"ImagePullBackoff" errors seen	These errors may be seen if you have a custom or private docker repository and have not yet configured the Kubernetes Monitoring Operator to properly recognize it. Read more about configuring for custom/private repo.

Problem:	Try this:
I am having an issue with my monitoring-operator deployment, and the current documentation does not help me resolve it.	Capture or otherwise note the output from the following commands, and contact the Technical Support team.
	<pre>kubectl -n netapp-monitoring get all kubectl -n netapp-monitoring describe all kubectl -n netapp-monitoring logs <monitoring-operator-pod>all -containers=true kubectl -n netapp-monitoring logs <telegraf-pod>all -containers=true</telegraf-pod></monitoring-operator-pod></pre>
Pods are running in Operator namespace (default:	These pods correspond to Workload Map data collector for Network Observability. Try these: • Check the logs of one of the pods to confirm minimum kernel version. For example: {"ci-tenant-id":"your-tenant-id","collector-cluster":"your- k8s-cluster- name","environment":"prod","level":"error","msg":"faile d in validation. Reason: kernel version 3.10.0 is less than minimum kernel version of 4.18.0","time":"2022- 11-09T08:23:08Z"} • Net-observer pods requires the Linux kernel version to be at least 4.18.0. Check the kernel version using the command "uname -r" and ensure they are >= 4.18.0
Pods are running in Operator namespace (default: netapp-monitoring), but no data is shown in UI for workload map or Kubernetes metrics in Queries	Check the time setting on the nodes of the K8S cluster. For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Agent machine using Network Time Protocol (NTP) or Simple Network Time Protocol (SNTP).
Some of the net-observer pods in Operator namespace are in Pending state	 Net-observer is a DaemonSet and runs a pod in each Node of the k8s cluster. Note the pod which is in Pending state, and check if it is experiencing a resource issue for CPU or memory. Ensure the required memory and CPU is available in the node.

Problem:	Try this:
I'm seeing the following in my logs immediately after installing the Kubernetes Monitoring Operator: [inputs.prometheus] Error in plugin: error making HTTP request to http://kube-state- metrics. <namespace>.svc.cluster.local:8080/metrics: Get http://kube-state- metrics.<namespace>.svc.cluster.local:8080/metrics: dial tcp: lookup kube-state- metrics.<namespace>.svc.cluster.local: no such host</namespace></namespace></namespace>	This message is typically only seen when a new operator is installed and the <i>telegraf-rs</i> pod is up before the <i>ksm</i> pod is up. These messages should stop once all pods are running.
I do see not any metrics being collected for the Kubernetes CronJobs that exist in my cluster.	Verify your Kubernetes version (i.e. kubectl version). If it is v1.20.x or below, this is an expected limitation. The kube-state-metrics release deployed with the Kubernetes Monitoring Operator only supports v1.CronJob. With Kubernetes 1.20.x and below, the CronJob resource is at v1beta.CronJob. As a result, kube-state-metrics cannot find the CronJob resource.
After installing the operator, the telegraf-ds pods enter CrashLoopBackOff and the pod logs indicate "su: Authentication failure".	Edit the telegraf section in <i>AgentConfiguration</i> , and set <i>dockerMetricCollectionEnabled</i> to false. For more details, refer to the operator's configuration options. spec: telegraf: - name: docker run-mode: - DaemonSet substitutions: - key: DOCKER_UNIX_SOCK_PLACEHOLDER value: unix:///run/docker.sock
I see repeating error messages resembling the following in my Telegraf logs: E! [agent] Error writing to outputs.http: Post "https:// <tenant_url>/rest/v1/lake/ingest/influxdb": context deadline exceeded (Client.Timeout exceeded while awaiting headers)</tenant_url>	Edit the telegraf section in <i>AgentConfiguration</i> , and increase <i>outputTimeout</i> to 10s. For more details, refer to the operator's configuration options.
I'm missing <i>involvedobject</i> data for some Event Logs.	Be sure you have followed the steps in the Permissions section above.

Problem:	Try this:
Why am I seeing two monitoring operator pods running, one named netapp-ci-monitoring-operator- <pod> and the other named monitoring-operator- <pod>?</pod></pod>	As of October 12, 2023, Data Infrastructure Insights has refactored the operator to better serve our users; for those changes to be fully adopted, you must remove the old operator and install the new one.
My kubernetes events unexpectedly stopped reporting to Data Infrastructure Insights.	Retrieve the name of the event-exporter pod:
	<pre>`kubectl -n netapp-monitoring get pods grep event-exporter awk '{print \$1}' sed 's/event- exporter./event-exporter/'`</pre>
	It should be either "netapp-ci-event-exporter" or "event-exporter". Next, edit the monitoring agent kubectl -n netapp-monitoring edit agent, and set the value for LOG_FILE to reflect the appropriate event-exporter pod name found in the previous step. More specifically, LOG_FILE should be set to either "/var/log/containers/netapp-ci-event- exporter.log" or "/var/log/containers/event- exporter*.log"
	<pre>fluent-bit: - name: event-exporter-ci substitutions: - key: LOG_FILE values: - /var/log/containers/netapp- ci-event-exporter*.log </pre>
	Alternatively, one can also uninstall and reinstall the agent.
I'm seeing pod(s) deployed by the Kubernetes Monitoring Operator crash because of insufficient resources.	Refer to the Kubernetes Monitoring Operator configuration options to increase the CPU and/or memory limits as needed.
A missing image or invalid configuration caused the netapp-ci-kube-state-metrics pods to fail to startup or become ready. Now the StatefulSet is stuck and configuration changes are not being applied to the netapp-ci-kube-state-metrics pods.	The StatefulSet is in a broken state. After fixing any configuration problems bounce the netapp-ci-kube-state-metrics pods.
Problem:	Try this:
---	---
netapp-ci-kube-state-metrics pods fail to start after running a Kubernetes Operator upgrade, throwing ErrImagePull (failing to pull the image).	Try resetting the pods manually.
"Event discarded as being older then maxEventAgeSeconds" messages are being observed for my Kubernetes cluster under Log Analysis.	Modify the Operator <i>agentconfiguration</i> and increase the <i>event-exporter-maxEventAgeSeconds</i> (i.e. to 60s), <i>event-exporter-kubeQPS</i> (i.e. to 100), and <i>event-exporter-kubeBurst</i> (i.e. to 500). For more details on these configuration options, see the configuration options page.
Telegraf warns of, or crashes because of, insufficient lockable memory.	Try increasing the limit of lockable memory for Telegraf in the underlying operating system/node. If increasing the limit is not an option, modify the NKMO agentconfiguration and set <i>unprotected</i> to <i>true</i> . This will instruct Telegraf to no attempt to reserve locked memory pages. While this can pose a security risk as decrypted secrets might be swapped out to disk, it allows for execution in environments where reserving locked memory is not possible. For more details on the <i>unprotected</i> configuration options, refer to the configuration options page.
I see warning messages from Telegraf resembling the following: <i>W</i> ! [inputs.diskio] Unable to gather disk name for "vdc": error reading /dev/vdc: no such file or directory	For the Kubernetes Monitoring Operator, these warning message are benign and can be safely ignored. Alternatively, edit the telegraf section in AgentConfiguration, and set <i>runDsPrivileged</i> to true. For more details, refer to the operator's configuration options.

Problem:	Try this:
My fluent-bit pod is failing with the following errors: [2024/10/16 14:16:23] [error] [/src/fluent- bit/plugins/in_tail/tail_fs_inotify.c:360 errno=24] Too many open files [2024/10/16 14:16:23] [error] failed initialize input tail.0 [2024/10/16 14:16:23] [error] [engine] input initialization failed	Try to change your <i>fsnotify</i> settings in your cluster: sudo sysctl fs.inotify.max_user_instances (take note of setting) sudo sysctl fs.inotify.max_user_instances= <som ething larger than current setting> sudo sysctl fs.inotify.max_user_watches (take note of setting) sudo sysctl fs.inotify.max_user_watches=<somet hing larger than current setting> Restart Fluent-bit. Note: to make these settings persistent across node restarts, you need to put the following lines in /etc/sysctl.conf</somet </som
The telegraf DS pods are reporting errors pertaining to the kubernetes input plugin failing to make HTTP requests due to the inability to validate the TLS certificate. For example: E! [inputs.kubernetes] Error in plugin: error making HTTP request to "https:// <kubelet_ip>:10250/stats/summary": Get "https://<kubelet_ip>:10250/stats/summary": tls: failed to verify certificate: x509: cannot validate certificate for</kubelet_ip></kubelet_ip>	<pre>fs.inotify.max_user_instances=<so current="" larger="" mething="" setting="" than=""> fs.inotify.max_user_watches=<some current="" larger="" setting="" than="" thing=""> This will occur if the kubelet is using self-signed certificates, and/or the specified certificate does not include the <kubelet_ip> in the certificates Subject Alternative Name list. To resolve this, the user can modify the agent configuration, and set telegraf:insecureK8sSkipVerify to true. This will configure the telegraf input plugin to skip verification. Alternatively, the user can configure the kubelet for serverTLSBootstrap, which will trigger a certificate</kubelet_ip></some></so></pre>

Additional information may be found from the Support page or in the Data Collector Support Matrix.

Memcached Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Memcached.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Memcached.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.



Information may be found in the Memcached wiki.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
Memcached	Namespace Server	Node IP Node Name	Accepting Connections Handled Authentication Requests Failed Authentications Bytes Used Bytes Read (per sec) Bytes Written (per sec) CAS Badval CAS Hits CAS Misses Flush Reqs (per sec) Get Reqs (per sec) Get Reqs (per sec) Connection Yields (per sec) Connection Structures Open Connections Current Stored Items Decr Requests Hits (per sec) Delete Requests Misses (per sec) Delete Requests Misses (per sec) Delete Requests Misses (per sec) Delete Requests Misses (per sec) Items Evicted Valid Evictions Expired Items Get Hits (per sec) Get Misses (per sec) Used Hash Bytes Hash Is Expanding Hash Power Level Incr Requests Misses (per sec) Server Max Bytes Listen Disabled Num Reclaimed Worker Threads Count Total Opened Connections Total Items Stored Touch Hits Touch Misses Server Uptime

Additional information may be found from the Support page.

MongoDB Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from MongoDB.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose MongoDB.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

mongoDB	MongoDB Configuration Gathers MongoDB metrics.	
What Operating S	ystem or Platform Are You Using?	Need Help?
🗱 RHEL & CentOS	•	
Select existing Ag	ent Access Key or create a new one	
Default (405fb5ec-d4cb-440	04-977b-71fa931e1ad3) * + Agent Access Key	
*Please ensure that you h	ave a Telegraf Agent in you environment before configuring Show Instruction	ons
Fallow Configurat		Need Unio
Follow Configurat	ion Steps	Weed help:
 Open mongod.com Telegraf agent resid Copy the contents contents to the /et 	Ion Steps If. Locate the line beginning with "bindlp", and append the ad des. After saving the change, restart the MongoDB server. below into a new .conf file under the /etc/telegraf/telegraf.d/ c c/telegraf/telegraf.d/cloudinsights-mongodb.conf file.	dress of the node on which the directory. For example, copy the
 Open mongod.con Telegraf agent resid Copy the contents contents to the /et [[inputs.mon ## An arra ## "mong 	Ion Steps If. Locate the line beginning with "bindlp", and append the ad des. After saving the change, restart the MongoDB server. below into a new conf file under the /etc/telegraf/telegraf.d/ of c/telegraf/telegraf.d/cloudinsights-mongodb.conf file. hgodb]] ay of URLs of the form: godb://" [user ":" pass "@"] host [":" port]	dress of the node on which the directory. For example, copy the
 Open mongod.con Telegraf agent resid Copy the contents contents to the /et [[inputs.mon ## An arra ## For exe ## For exe ## For exe 	Ion Steps If. Locate the line beginning with "bindlp", and append the ad des. After saving the change, restart the MongoDB server. below into a new conf file under the /etc/telegraf/telegraf.d/c c/telegraf/telegraf.d/cloudinsights-mongodb.conf file. ngodb]] ay of URLs of the form: godb://" [user ":" pass "@"] host [":" port] imple: bdb://user:auth_key@10.10.3.30:27017, user is a common	dress of the node on which the directory. For example, copy the
 Open mongod.con Telegraf agent resid Copy the contents contents to the /et Copy the contents contents to the /et ([inputs.mongot ## An arra ## "mongot ## For exe ## mongot execution Replace <insert_i machine address, or</insert_i 	Ion Steps If. Locate the line beginning with "bindlp", and append the ad des. After saving the change, restart the MongoDB server. below into a new conf file under the /etc/telegraf/telegraf.d/ of c/telegraf/telegraf.d/cloudinsights-mongodb.conf file. reodb]] ay of URLs of the form: godb://" [user ":" pass "@"] host [":" port] ample: bdb://user:auth_key@10.10.3.38:27017, db://user:auth_key@10.10.3.38:27017, db://user:auth_key@10.10.3.38:27017, db://too.io.nonvince.	dress of the node on which the directory. For example, copy the
 Open mongod.con Telegraf agent resid Copy the contents contents to the /et Copy the contents contents to the /et Inputs.mon ## An arra ## "mongo ## For exe ## mongo Replace <insert_1< li=""> Replace <insert_1< li=""> </insert_1<></insert_1<>	Ion Steps If. Locate the line beginning with "bindlp", and append the ad des. After saving the change, restart the MongoDB server. below into a new conf file under the /etc/telegraf/telegraf.d/ of c/telegraf/telegraf.d/cloudinsights-mongodb.conf file. ngodb]] ay of URLs of the form: godb://" [user ":" pass "@"] host [":" port] ample: odb://user:auth_key@10.10.3.30:27017, db://user:auth_key@10.10.3.30:27017, db://too is in continued MONGODB_ADDRESS> with the applicable MongoDB server a and refrain from using a loopback address. MONGODB_PORT> with the applicable MongoDB port.	dress of the node on which the directory. For example, copy the
 Open mongod.con Telegraf agent resid Copy the contents contents to the /et Copy the contents contents to the /et (inputs.mon ## An arra ## "mongo ## For exe ## mongo Replace <insert_i machine address, a</insert_i Replace <insert_i< li=""> Restart the Telegram </insert_i<>	Ion Steps If. Locate the line beginning with "bindlp", and append the ad des. After saving the change, restart the MongoDB server. below into a new conf file under the /etc/telegraf/telegraf.d/c c/telegraf/telegraf.d/cloudinsights-mongodb.conf file. ngodb]] ay of URLs of the form: godb://" [user ":" pass "@"] host [":" port] mple: odb://user:auth_key@10.10.3.30:27017, vb://ws.concorrect MONGODB_ADDRESS> with the applicable MongoDB server a and refrain from using a loopback address. MONGODB_PORT> with the applicable MongoDB port. f service.	dress of the node on which the directory. For example, copy the c

Information may be found in the MongoDB documentation.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
MongoDB	Namespace Hostname		
MongoDB Database	Namespace Hostname Database name		

Information may be found from the Support page.

MySQL Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from MySQL.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose MySQL.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

MySQL	MySQL Configuration Gathers MySQL metrics.	
What Operating S	System or Platform Are You Using?	Need Help
Windows	Ψ	
Select existing Ag	gent Access Key or create a new one	
Default (405fb5ec-d4cb-44	i04-977b-71fa931e1ad3) + Agent Access Key	
*Please ensure that you	have a Telegraf Agent in you environment before configuring Show Instruc	tions
ollow Configura	tion Steps	Need Hel
copy the contents [[inputs.my ## USER-/ ## e.g. s ## Plesss (i.e. localhost	to the C:\Program Files\telegraf\telegraf.d\cloudinsights-mys sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1).	ql.conffile. IP(s), and port(s) a loopback address
copy the contents [[inputs.my ## USER-4 ## e.g. = ## Please (i.e. localhost 	to the C:\Program Files\telegraf\telegraf.d\cloudinsights-mys sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1).	ql.conf file. IP(s), and port(s) a loopback address
 copy the contents [[inputs.my ## USER-4 ## e.g. E ## Please (i.e. localhost Review and verify Replace <insert_< li=""> </insert_<>	to the C:\Program Files\telegraf\telegraf.d\cloudinsights-mys sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1). the contents of the configuration file. USERNAME> and <insert_password> with the applicable</insert_password>	<pre>qLconf file. IP(s), and port(s) a loopback address MySQL credentials.</pre>
 copy the contents [[inputs.my ## USER-4 ## USER-4 ## Please (i.e. localhost Review and verify Replace <insert_< li=""> Replace <insert_< li=""> </insert_<></insert_<>	to the C:\Program Files\telegraf.telegraf.d\cloudinsights-mys sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1). the contents of the configuration file. USERNAME> and <insert_password> with the applicable PROTOCOL> with the applicable MySQL connection protocol</insert_password>	<pre>qLconf file. IP(s), and port(s) a loopback address MySQL credentials. The typical protocol is tcp.</pre>
 copy the contents [[inputs.my ## USER-/ ## USER-/ ## e.g. i ## Please (i.e. localhost contents Review and verify Replace <insert_< li=""> Replace <insert_< li=""> Replace <insert_< li=""> </insert_<></insert_<></insert_<>	sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1). the contents of the configuration file. USERNAME> and <insert_password> with the applicable PROTOCOL> with the applicable MySQL connection protocol MYSQL_ADDRESS> with the applicable MySQL server address in from using a loopback address.</insert_password>	<pre>qLconf file. [IP(s), and port(s)] a loopback address [MySQL credentials.]. The typical protocol is tcp.] s. Please specify a real machine</pre>
 copy the contents [[inputs.my ## USER-/ ## USER-/ ## USER-/ ## Please (i.e. localhost 	sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3506)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1). the contents of the configuration file. USERNAME> and <insert_password> with the applicable PROTOCOL> with the applicable MySQL connection protocol MYSQL_ADDRESS> with the applicable MySQL server address in from using a loopback address.</insert_password>	ql.conf file. IP(s), and port(s) a loopback address MySQL credentials. The typical protocol is tcp. S. Please specify a real machine spical port is 3306.
 copy the contents copy the contents [[inputs.my ## USER-/ ## e.g. i ## Please (i.e. localhost contents Review and verify Replace <insert_< li=""> Replace <insert_ address, and refra</insert_ Replace <insert_< li=""> Replace <insert_< li=""> Modify the 'tls' pa </insert_<></insert_<></insert_<>	sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@tcp(127.0.0.1:3506)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1). CTIONERNAME> and <insert_password> with the applicable PROTOCOL> with the applicable MySQL connection protocol MYSQL_ADDRESS> with the applicable MySQL server address in from using a loopback address. MYSQL_PORT> with the applicable MySQL server port. The ty rameter in accordance to the MySQL server configuration.</insert_password>	<pre>gl.conf file. [] IP(s), and port(s) a loopback address a loopbac</pre>
 copy the contents copy the contents [[inputs.my ## USER-/ ## e.g. i ## Please (i.e. localhost corr Review and verify Replace <insert_< li=""> </insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<></insert_<>	sql]] CTION: Provide comma-separated list of mySQL credentials, ervers = ["user:passwd@top(127.0.0.1:3506)/?tls=false"] specify actual machine IP address, and refrain from using or 127.0.0.1). the contents of the configuration file. USERNAME> and <insert_password> with the applicable PROTOCOL> with the applicable MySQL connection protocol MYSQL_ADDRESS> with the applicable MySQL server address in from using a loopback address. MYSQL_PORT> with the applicable MySQL server port. The ty rameter in accordance to the MySQL server configuration. af service.</insert_password>	IP(s), and port(s) a loopback address

Information may be found in the MySQL documentation.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
MySQL	Namespace MySQL Server	Node IP Node Name	Aborted Clients (per sec) Aborted Connects (per sec) RX Bytes (per sec) Commands Admin (per sec) Commands Alter Event Commands Alter Function Commands Alter Function Commands Alter Instance Commands Alter Instance Commands Alter Server Commands Alter Server Commands Alter User Commands Alter User Commands Alter User Commands Assign To Keycache Commands Begin Commands Begin Commands Call Procedure Commands Change DB Commands Change DB Commands Change DB Commands Change Repl Filter Commands Check Commands Check Commands Checksum Commands Create DB Commands Create DB Commands Create Event Commands Create Index Commands Create Event Commands Create Event Commands Create Event Commands Create Index Commands Create Index

Additional information may be found from the Support page.

Netstat Data Collector

Data Infrastructure Insights uses this data collector to gather Netstat metrics.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Netstat.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

netstat	Netstat Configuration Gathers netstat metrics of the host where telegraf agent is installed.	
What Operating S	System or Platform Are You Using?	Need Help?
t Windows	▼	
Select existing Ag Default (405fb5ec-d4cb-44 *Please ensure that you h	104-977b-71fa931e1ad3) Agent Access Key have a Telegraf Agent in you environment before configuring Show Instructions	5
Follow Configurat	tion Steps	Need Help?
Copy the contents copy the contents	; below into a new .conf file under the C:\Program Files\telegraf\te to the C:\Program Files\telegraf\telegraf.d\cloudinsights-netstat.	elegraf.d\ folder. For example, conf file.
<pre># Read TCP [[inputs.ne</pre>	" metrics such as established, time wait and sockets counts. tstat]] iguration etstat.tags] sights = "true"	C
2 Restart the Telegra	af service.	
Stop-Service -N	ame telegraf -ErrorAction SilentlyContinue; Start-Service -Nam	me telegraf [

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:	Datapoints:
Netstat	Node UUID	Node IP Node Name	

Troubleshooting

Additional information may be found from the Support page.

Nginx Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Nginx.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Nginx.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

NGINX	Nginx Configuration Gathers Nginx metrics.		
What Operating S	stem or Platform Are	Voi Helero	Need Hele?
that operating 5	stem of Flation Are	You Using?	Iveep melb:
O Ubuntu & Debian	stem of Platform Are	* Using?	Need neep
Ubuntu & Debian Select existing Age	ent Access Key or creat	* a new one	Neo nep:

Follow Configuration Steps
1 If you already have a URL enabled to provide Nginx metrics, go directly to the plugin configuration.
Nginx metrics are available through a status page when the HTTP stub status module is enabled.Refer to the below link for verifying/enabling http_stub_status_module.
http://nginx.org/en/docs/http/ngx_http_stub_status_module.html
After verifying the module is enabled , modify the Nginx configuration to set up a locally-accessible URL for the status page.
<pre>server { listen <port number="">; Please specify actual machine IP address, and refrain from using a loopback address (i.e. localhost or 127.0.0.1) server_name <ip address="">; location /nginx_status { +</ip></port></pre>
Reload the configuration:
nginx -s reload
Copy the contents below into a new conf file under the /etc/telegraf/telegraf.d/ directory. For example, copy the contents to the /etc/telegraf/telegraf.d/cloudinsights-nginx.conf file.
<pre>[[inputs.nginx]] ## USER-ACTION: Provide Nginx status url ## Please specify actual machine IP address where nginx_status is enabled, and refrain from using a loopback address (i.e. localhost or 127.0.0.1). ## When configuring with multiple Nginx servers, enter them in the format ["url1", "url2", "</pre>
Replace <insert_nginx_address> with the applicable Nginx address. Please specify a real machine address, and refrain from using a loopback address.</insert_nginx_address>
Replace <insert_nginx_port> with the applicable Nginx port_</insert_nginx_port>
8 Restart the Telegraf service.
systemotl restart telegraf

Nginx metric collection requires that Nginx http_stub_status_module be enabled.

Additional information may be found in the Nginx documentation.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
Nginx	Namespace Server	Node IP Node Name Port	Accepts Active Handled Reading Requests Waiting Writing

Additional information may be found from the Support page.

PostgreSQL Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from PostgreSQL.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose PostgreSQL.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the + Agent Access Key button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

RostgreSQL	PostgreSQL Configuration Gathers PostgreSQL metrics.	
What Operating S	ystem or Platform Are You Using?	Need Help?
🐐 RHEL & CentOS		
Select existing Age	ent Access Key or create a new one	
Default (405fb5ec-d4cb-440	14-977b-71fa931e1ad3) - Agent Access Key	
*Please ensure that you h	ave a Telegraf Agent in you environment before configuring Show Instruction	15
Follow Configurat	ion Steps	Need Help?
Copy the contents contents to the /etc	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file.	rectory. For example, copy the
Copy the contents contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p <insert_postgres< p=""></insert_postgres<>	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. gresql]] W: Provide credentials for access, address of PostgreSQL ser er, one DB for access postgres:// <insert_username>:<insert_password>@<insert_postgr QL_PORT>/<insert_db>"</insert_db></insert_postgr </insert_password></insert_username>	rectory. For example, copy the rver, port for RESQL_ADDRESS>:
 Copy the contents contents to the /etc contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p <insert_postgres< li=""> Replace <insert_u< li=""> </insert_u<></insert_postgres<>	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. gresql]] W1: Provide credentials for access, address of PostgreSQL ser instgres://(INSERT_USERNAME>: <insert_password>@<insert_postg QL_PORT>/<insert_db>" JSERNAME> and <insert_password> with the applicable Po</insert_password></insert_db></insert_postg </insert_password>	rectory. For example, copy the rver, port for RESQL_ADDRESS>: ostgreSQL credentials.
 Copy the contents contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p <insert_postgres <="" li=""> Replace <insert_c machine address, a </insert_c </insert_postgres>	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. gresql]] W: Provide credentials for access, address of PostgreSQL ser er, one DB for access bostgres:// <insert_username>:<insert_password>@<insert_postg bQL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable Po POSTGRESQL_ADDRESS> with the applicable PostgreSQL addr and refrain from using a loopback address.</insert_password></insert_db></insert_postg </insert_password></insert_username>	rectory. For example, copy the rver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real
 Copy the contents contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p <insert_postgres <="" li=""> Replace <insert_c machine address, a </insert_c Replace <insert_c <="" li=""> </insert_c></insert_postgres>	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. resql]] W: Provide credentials for access, address of PostgreSQL ser r, one DB for access postgres:// <insert_username>:<insert_password>@<insert_postg PQL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL port.</insert_password></insert_db></insert_postg </insert_password></insert_username>	rectory. For example, copy the Tver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real
 Copy the contents contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p <insert_postgres <="" li=""> Replace <insert_c machine address, a Replace <insert_c Seplace <insert_c< td=""><td>below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. resql]] W: Provide credentials for access, address of PostgreSQL ser r, one DB for access postgres://<insert_username>:<insert_password>@<insert_postg PQL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable Post postGRESQL_ADDRESS> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL port. DB> with the applicable PostgreSQL database.</insert_password></insert_db></insert_postg </insert_password></insert_username></td><td>rectory. For example, copy the Inver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real</td></insert_c<></insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_postgres>	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. resql]] W: Provide credentials for access, address of PostgreSQL ser r, one DB for access postgres:// <insert_username>:<insert_password>@<insert_postg PQL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable Post postGRESQL_ADDRESS> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL port. DB> with the applicable PostgreSQL database.</insert_password></insert_db></insert_postg </insert_password></insert_username>	rectory. For example, copy the Inver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real
 Copy the contents contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p CINSERT_POSTGRES Replace <insert_c machine address, a </insert_c Replace <insert_c <="" li=""> Replace <insert_c <="" li=""> Replace <insert_c <="" li=""> Replace <insert_c <="" li=""> </insert_c></insert_c></insert_c></insert_c>	below into a new .conf file under the /etc/telegraf/telegraf.d/ di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. resql]] W: Provide credentials for access, address of PostgreSQL set er, one DB for access postgres:// <insert_username>:<insert_password>@<insert_postgr QL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable Po POSTGRESQL_ADDRESS> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL port. DB> with the applicable PostgreSQL database. e' if needed for server disambiguation (to avoid name clashes).</insert_password></insert_db></insert_postgr </insert_password></insert_username>	rectory. For example, copy the Tver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real
 Copy the contents contents to the /etc [[inputs.postg # USER-ACTIO PostgreSQL serve address = "p <insert_postgres <="" li=""> Replace <insert_c machine address, a Replace <insert_c Replace <insert_c< td=""><td>below into a new .conf file under the /etc/telegraf/telegraf.d/di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. resql]] W: Provide credentials for access, address of PostgreSQL ser r, one DB for access postgres://<insert_username>:<insert_password>@<insert_postg QL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL port. DB> with the applicable PostgreSQL database. e' if needed for server disambiguation (to avoid name clashes). f service.</insert_password></insert_db></insert_postg </insert_password></insert_username></td><td>rectory. For example, copy the Inver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real</td></insert_c<></insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_c </insert_postgres>	below into a new .conf file under the /etc/telegraf/telegraf.d/di c/telegraf/telegraf.d/cloudinsights-postgresql.conf file. resql]] W: Provide credentials for access, address of PostgreSQL ser r, one DB for access postgres:// <insert_username>:<insert_password>@<insert_postg QL_PORT>/<insert_db>" USERNAME> and <insert_password> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL addr and refrain from using a loopback address. POSTGRESQL_PORT> with the applicable PostgreSQL port. DB> with the applicable PostgreSQL database. e' if needed for server disambiguation (to avoid name clashes). f service.</insert_password></insert_db></insert_postg </insert_password></insert_username>	rectory. For example, copy the Inver, port for RESQL_ADDRESS>: ostgreSQL credentials. ess. Please specify a real

Information may be found in the PostgreSQL documentation.

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
PostgreSQL Server	Namespace Database Server	Node Name Node IP	Buffers Allocated Buffers Backend Buffers Backend File Sync Buffers Checkpoint Buffers Clean Checkpoints Sync Time Checkpoints Write Time Checkpoints Requests Checkpoints Timed Max Written Clean
PostgreSQL Database	Namespace Database Server	Database OID Node Name Node IP	Blocks Read Time Blocks Write Time Blocks Hits Blocks Reads Conflicts Deadlocks Client Number Temp Files Bytes Temp Files Number Rows Deleted Rows Fetched Rows Inserted Rows Inserted Rows Returned Rows Updated Transactions Committed Transactions Rollbacked

Additional information may be found from the Support page.

Puppet Agent Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Puppet Agent.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Puppet.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.



Information may be found in the Puppet documentation

Objects and Counters

Object:	Identifiers:	Attributes:	Datapoints:
---------	--------------	-------------	-------------

PupperAgent	Node UUID	Node Name Location Node IP Version Configstring Version Puppet	Events Failure Events Success Events Total Resources Changed Resources Failed Resources Failed To Restart Resources Outofsync Resources Restarted Resources Scheduled Resources Scheduled Resources Total Time Anchor Time Configretrieval Time Cron Time Exec Time File Time File Time File Dime File Time Schedule Time Schedule Time Schedule Time Schedule Time Schedule Time Total Time User
-------------	-----------	--	---

Additional information may be found from the Support page.

Redis Data Collector

Data Infrastructure Insights uses this data collector to gather metrics from Redis. Redis is an open source, in-memory data structure store used as a database, cache, and message broker, supporting the following data structures: strings, hashes, lists, sets, and more.

Installation

1. From **Observability > Collectors**, click **+Data Collector**. Choose Redis.

- 2. If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click *Show Instructions* to expand the Agent installation instructions.
- 3. Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the **+ Agent Access Key** button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- 4. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

🥞 redis	Redis Configuration Gathers Redis metrics.	
What Operating S	ystem or Platform Are You Using?	Need Help?
Windows	Ψ	
Select existing Age	ent Access Key or create a new one	
Default (405fb5ec-d4cb-440	4-977b-71fa931e1ad3) - Agent Access Key	
*Please ensure that you ha	eve a Telegraf Agent in you environment before configuring Show Instruct	ions
Follow Configurati	on Steps	Need Help?
	accept connections from the address of the node on which t	he Telegraf agent resides. Open
Vi /etc/redis.co	tion file.	B
Configure Redis to the Redis configura vi /etc/redis.co Locate the line that resides	ition file. nf begins with 'bind 127.0.0.1', and append the address of the r	rode on which the Telegraf agent
 Configure Redis to the Redis configuration of the Redis config	ition file. nf begins with 'bind 127.0.0.1', and append the address of the r	node on which the Telegraf agent
 Configure Redis to the Redis configure vi /etc/redis.co Locate the line that resides bind 127.0.0.1 < Copy the contents I copy the contents to the copy the contents to the content to the conte	ition file. nf begins with 'bind 127.0.0.1', and append the address of the r NODE_IP_ADDRESS> below into a new conf file under the C:\Program Files\telegra o the C:\Program Files\telegraf.\telegraf.d\cloudinsights-redis	node on which the Telegraf agent the felegraf.d\ folder. For example, conf file.
 Contigure Redis to the Redis configuration of the Redis configuration of the Redis configuration of the Redis configuration of the Redis control of the Redis contents of the contents of the contents of the Read metrics of the residue of the residue	tion file. nf begins with 'bind 127.0.0.1', and append the address of the r NODE_IP_ADDRESS> below into a new conf file under the C:\Program Files\telegra o the C:\Program Files\telegraf.telegraf.d\cloudinsights-redis from one or many redis servers]] ervers via a url matching: 1://][:password]@address[:port]	thode on which the Telegraf agent the Telegraf agent the telegraf.d folder. For example, conf file.
 Contigure Redis to the Redis configuration of the Redistance o	ation file. Inf begins with 'bind 127.0.0.1', and append the address of the r NODE_IP_ADDRESS> below into a new conf file under the C.\Program Files\telegra o the C.\Program Files\telegraf.d\cloudinsights-redis from one or many redis servers I] from one or many redis servers I] ervers via a url matching: 1://][:password]@address[:port] 	tode on which the Telegraf agent tode on the Telegraf agent
 Contigure Redis to the Redis configuration of the Redis contents of the contents of the contents of the Read metrics of the contents of the Read metrics of the residual of the Read metrics of the residual of the r	Ation file. Inf begins with 'bind 127.0.0.1', and append the address of the r NODE_IP_ADDRESS> below into a new conf file under the C.\Program Files\telegra o the C.\Program Files\telegraf.d\cloudinsights-redis from one or many redis servers if ervers via a url matching: 1://][:password]@address[:port] REDIS_ADDRESS> with the applicable Redis address. Please s ing a loopback address. REDIS_PORT> with the applicable Redis port.	tode on which the Telegraf agent the folder. For example, conf file. specify a real machine address,
 Configure Redis to the Redis configuration of the Redis contents of the Copy the contents of the Read metrics [[inputs.redis ## specify s ## [protoco ## e.g	ation file. Inf begins with 'bind 127.0.0.1', and append the address of the r NODE_IP_ADDRESS> below into a new conf file under the C.\Program Files\telegra o the C.\Program Files\telegraf.d\cloudinsights-redis from one or many redis servers]] ervers via a url matching: 1://][:password]@address[:port] :	node on which the Telegraf agent the telegraf.d\ folder. For example, conf file. specify a real machine address,

Information may be found in the Redis documentation.

Objects and Counters

The following objects and their counters are collected:

Object:	Identifiers:	Attributes:	Datapoints:
Redis	Namespace Server		

Troubleshooting

Additional information may be found from the Support page.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.