



Collecting Data

Data Infrastructure Insights

NetApp
February 03, 2026

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructure-insights/task_getting_started_with_cloud_insights.html on February 03, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Collecting Data	1
Getting started gathering data	1
Install an Acquisition Unit	1
Configure the Data Collector - Infrastructure	2
Configure the Data Collector - Operating Systems and Services	2
Add Dashboards	2
That's all there is to it	3
Useful definitions	3
Acquisition Unit Requirements	3
Requirements	3
Additional recommendations	5
Regarding Sizing	5
Configuring Acquisition Units	6
Adding a Linux Acquisition Unit	6
Adding a Windows Acquisition Unit	8
Uninstalling an Acquisition Unit	10
Reinstalling an Acquisition Unit	11
Viewing AU Details	12
Configuring an Agent to Collect Data (Windows/Linux)	13
Installing an Agent	13
Verifying Telegraf Package Checksums	21
Creating and Using API Access Tokens	23
Troubleshooting	23
Configuring Data Collectors	24
Determining data collector acquisition status	25
Managing configured data collectors	25
Controlling Data Collector polling	26
Editing data collector information	26
Cloning data collectors	26
Performing bulk actions on data collectors	27
Researching a failed data collector	27

Collecting Data

Getting started gathering data

After you have signed up for Data Infrastructure Insights and log in to your environment for the first time, you will be guided through the following steps in order to begin collecting and managing data.

Data collectors discover information from your data sources, such as storage devices, network switches, and virtual machines. The information gathered is used for analysis, validation, monitoring and troubleshooting.


Data Infrastructure Insights has available three types of data collectors:

- Infrastructure (storage devices, network switches, compute infrastructure)
- Operating Systems (such as VMware or Windows)
- Services (such as Kafka)

Select your first data collector from the supported vendors and models available. You can easily add additional data collectors later.

Install an Acquisition Unit

If you selected an *Infrastructure* data collector, an Acquisition Unit is required to inject data into Data Infrastructure Insights. You will need to download and install the Acquisition Unit software on a server or VM on the data center from which you will be collecting. A single Acquisition Unit can be used for multiple data collectors.



The screenshot shows the 'Install Acquisition Unit' section of the NetApp ONTAP Data Management Software. It includes a 'What Operating System or Platform Are You Using?' dropdown set to 'Linux', 'Linux Versions Supported' (with a help icon), 'Production Best Practices' (with a help icon), and a 'Need Help?' link. Below this is an 'Installation Instructions' section with numbered steps: 1. Copy Installer Snippet (with a note about a 24-hour key), 2. Paste the snippet into a bash shell to run the installer, and 3. Waiting for Acquisition Unit to connect... There is also a note about instructions for the Acquisition Unit software.

NetApp

ONTAP Data Management Software

Install Acquisition Unit

Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.

What Operating System or Platform Are You Using?

Linux ▼ Linux Versions Supported ⓘ Production Best Practices ⓘ

Need Help?

Installation Instructions

1 ⓘ Copy Installer Snippet
This snippet has a unique key valid for 24 hours for this Acquisition Unit only.

+ Reveal Installer Snippet

2 Paste the snippet into a bash shell to run the installer.

3 ⓘ Waiting for Acquisition Unit to connect...

- Follow the [instructions](#) displayed to install your Acquisition Unit. Once the Acquisition Unit software is installed, the Continue button is displayed and you can proceed to the next step.

You may set up additional acquisition units later if needed. For example, you may want different Acquisition Units collecting information from data centers in different regions.

Configure the Data Collector - Infrastructure

For *Infrastructure* data collectors, you will be asked to fill out the data collector fields presented:

- Give the data collector a unique and meaningful name.
- Enter the credentials (user name and password) to connect to the device, as appropriate.
- Fill in any other mandatory fields in *Configuration* and *Advanced Configuration* sections.
- Click **Add Collector** to save the data collector.

You will be able to configure additional data collectors later.

Configure the Data Collector - Operating Systems and Services

Operating System:

For *Operating System* data collectors, choose a platform (Linux, Windows) to install a Data Infrastructure Insights Agent.

You must have at least one agent to collect data from Services.

The agent also collects data from the host itself, for use in Data Infrastructure Insights. This data is categorized as "Node" data in widgets, etc.

- Open a terminal or command window on the agent host or VM, and paste the displayed command to install the agent.
- When installation is complete, click **Complete Setup**.

Services:

For *Service* data collectors, click on a tile to open the instructions page for that service.

- Choose a platform and an Agent Access Key.
- If you don't have an agent installed on that platform, follow the instructions to install the agent.
- Click **Continue** to open the data collector instruction page.
- Follow the instructions to configure the data collector.
- When configuration is complete, click **Complete Setup**.

Add Dashboards

Depending on the type of initial data collector you selected to configure (storage, switch, etc.), one or more relevant dashboards will be imported. For example, if you configured a storage data collector, a set of storage-related dashboards will be imported, and one will be set as your Data Infrastructure Insights Home Page. You can change the home page from the **Dashboards > Show All Dashboards** list.

You can import additional dashboards later, or [create your own](#).

That's all there is to it

After you complete the initial setup process, your environment will begin to collect data.

If your initial setup process is interrupted (for example, if you close the browser window), you will need to follow the steps manually:

- Choose a Data Collector
- Install an Agent or Acquisition Unit if prompted
- Configure the Data Collector

Useful definitions

The following definitions may be useful when talking about Data Infrastructure Insights data collectors or features:

- Collector life cycle: A collector will belong to one of the following states in its life cycle:
 - **Preview:** Available in a limited capacity or to a limited audience. [Preview features](#) and data collectors are expected to become GA following the preview period. Preview periods vary based on audience or functionality.
 - **GA:** A feature or data collector that is Generally Available to all customers, based on Edition or feature set.
 - **Deprecated:** Applies to data collectors that are, or are expected to become, no longer functionally sustainable. Deprecated data collectors are often replaced with newer, functionally-updated data collectors.
 - **Deleted:** A data collector that has been removed and is no longer available.
- Acquisition Unit: a computer dedicated to hosting data collectors, typically a Virtual Machine. This computer is typically located in the same data center / VPC as the monitored items.
- Data Source: a module for communicating with a hardware or software stack. It consists of a configuration and code that runs on the AU computer to communicate with the device.

Acquisition Unit Requirements

You must install an Acquisition Unit (AU) in order to acquire information from your infrastructure data collectors (storage, VM, port, EC2, etc.). Before you install the Acquisition Unit, you should ensure that your tenant meets operating system, CPU, memory, and disk space requirements.

Requirements

Component	Linux Requirement	Windows Requirement
-----------	-------------------	---------------------

Operating system	<p>A computer running a licensed version of one of the following:</p> <ul style="list-style-type: none"> * AlmaLinux 9.3 through 9.5 * Centos (64-bit): 7.2 through 7.9, 8.1 through 8.4, Stream 8, Stream 9 * Debian (64-bit): 9 through 12 * OpenSUSE Leap 15.1 through 15.6 * Oracle Enterprise Linux (64-bit): 7.5 through 7.9, 8.1 through 8.10, 9.3 through 9.6 * Red Hat (64-bit): 7.2 through 7.9, 8.1 through 8.10, 9.1 through 9.6 * Rocky 9.0 through 9.5 * SUSE Linux Enterprise Server 15, 15 SP2 through 15 SP6 * Ubuntu Server: 18.04, 20.04, 22.04, 24.04 LTS * SELinux on the above platforms <p>This computer should be running no other application-level software. A dedicated server is recommended.</p> <p>If you are running with SELinux, it is recommended to execute the following commands on the acquisition unit system:</p> <pre>sudo semanage fcontext -a -t usr_t "/opt/netapp/cloudinsights(.*)?" sudo restorecon -R /opt/netapp/cloudinsights</pre>	<p>A computer running a licensed version of one of the following:</p> <ul style="list-style-type: none"> * Microsoft Windows 11 * Microsoft Windows Server 2016 * Microsoft Windows Server 2019 * Microsoft Windows Server 2022 * Microsoft Windows Server 2025 <p>This computer should be running no other application-level software. A dedicated server is recommended.</p>
CPU	2 CPU cores	Same
Memory	8 GB RAM	Same
Available disk space	<p>50 GB (100 GB recommended)</p> <p>For Linux, disk space should be allocated in this manner:</p> <ul style="list-style-type: none"> /opt/netapp 10 GB (20 GB for large environments) /var/log/netapp 40 GB (80 GB for large environments) /tmp at least 1 GB available during installation 	50 GB

Network	<p>100 Mbps/1 Gbps Ethernet connection, static IP address, and port 80 or 443 (outbound) connectivity from Acquisition Unit to *.cloudinsights.netapp.com or your Data Infrastructure Insights tenant (i.e. <a href="https://<tenant_id>.c01.cloudinsights.netapp.com">https://<tenant_id>.c01.cloudinsights.netapp.com) is required. For requirements between Acquisition Unit and each Data Collector, please refer to instructions for the Data Collector.</p> <p>If your organization requires proxy usage for internet access, you may need to understand your organization's proxy behavior and seek certain exceptions for Data Infrastructure Insights to work. For example, does your organization block access by default, and only allow access to specific web sites/domains by exception? If so, you will need to get the following domain added to the exception list:</p> <p>*.cloudinsights.netapp.com</p> <p>For more information, read about Proxies here (Linux) or here (Windows).</p>	Same
Permissions	Sudo permissions on the Acquisition Unit server. /tmp must be mounted with exec capabilities.	Administrator permissions on the Acquisition Unit server
Virus Scan		During installation, you must completely disable all virus scanners. Following installation, the paths used by the Acquisition Unit software must be excluded from virus scanning.

Additional recommendations

- For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Acquisition Unit machine using **Network Time Protocol (NTP)** or **Simple Network Time Protocol (SNTP)**.

Regarding Sizing

You can get started with a Data Infrastructure Insights Acquisition Unit with just 8GB memory and 50GB of disk

space, however, for larger environments you should ask yourself the following questions:

Do you expect to:

- Discover more than 2500 virtual machines or 10 large (> 2 node) ONTAP clusters, Symmetrix, or HDS/HPE VSP/XP arrays on this Acquisition Unit?
- Deploy 75 or more total data collectors on this Acquisition Unit?

For each "Yes" answer above, it is recommended to add 8 GB of memory and 50 GB of disk space to the AU. So for example if you answered "Yes" to both, you should deploy a 24GB memory system with 150GB or more of disk space. On Linux, the disk space to be added to the log location.

For additional sizing questions, contact NetApp Support.

Configuring Acquisition Units

Data Infrastructure Insights collects device data using one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Data Infrastructure Insights for analysis.

This topic describes how to add Acquisition Units and describes additional steps required when your environment uses a proxy.



For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Acquisition Unit machine using **Network Time Protocol (NTP)** or **Simple Network Time Protocol (SNTP)**.

Read about Data Infrastructure Insights security [here](#).

Adding a Linux Acquisition Unit

Before you begin

- If your system is using a proxy, you must set the proxy environment variables before the acquisition unit is installed. For more information, see [Setting proxy environment variables](#).

Steps for Linux Acquisition Unit Installation

1. Log in as Administrator or Account Owner to your Data Infrastructure Insights environment.
2. Click **Observability > Collectors > Acquisition Units > +Acquisition Unit**

The system displays the *Install Acquisition Unit* dialog. Choose Linux.

Install Acquisition Unit

Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.

What Operating System or Platform Are You Using?

Linux



Linux Versions Supported [i](#)

Production Best Practices [i](#)

[Need Help?](#)

Installation Instructions

1

[Copy Installer Snippet](#)

This snippet has a unique key valid for 24 hours for this Acquisition Unit only.

[Reveal Installer Snippet](#)

2

Paste the snippet into a bash shell to run the installer.

3

Waiting for Acquisition Unit to connect...

1. Verify that the server or VM hosting the Acquisition Unit meets the recommended system requirements.
2. Verify that the server is running a supported version of Linux. Click *OS Versions Supported (i)* for a list of supported versions.
3. Copy the Installation command snippet in the dialog into a terminal window on the server or VM that will host the Acquisition unit.
4. Paste and execute the command in the Bash shell.

After you finish

- Click **Observability > Collectors > Acquisition units** to check the status of Acquisition Units.
- You can access the Acquisition Unit logs at `/var/log/netapp/cloudinsights/acq/acq.log`
- Use the following script to control the Acquisition Unit:
 - `cloudinsights-service.sh` (stop, start, restart, check the status)
- Use the following script to uninstall the Acquisition Unit:
 - `cloudinsights-uninstall.sh`

Setting proxy environment variables

For environments that use a proxy, you must set the proxy environment variables before you add the Acquisition Unit. The instructions for configuring the proxy are provided on the *Add Acquisition Unit* dialog.

1. Click `+` in *Have a Proxy Server?*
2. Copy the commands to a text editor and set your proxy variables as needed.

Note: Be aware of restrictions on special characters in proxy username and password fields: '%' and '!' are allowed in the username field. '.', '%', and '!' are allowed in the password field.

3. Run the edited command in a terminal using the Bash shell.
4. Install the Acquisition Unit software.

Proxy Configuration

The Acquisition Unit uses 2-way/mutual authentication to connect to the Data Infrastructure Insights server. The client certificate must be passed to the Data Infrastructure Insights server to be authenticated. To accomplish this, the proxy must be set up to forward the https request to the Data Infrastructure Insights server without decrypting the data.

The simplest way to do this is to specify wildcard configuration in your proxy/firewall to communicate with Data Infrastructure Insights, for example:

```
*.cloudinsights.netapp.com
```



The use of an asterisk (*) for wildcard is common, but your proxy/firewall configuration may use a different format. Check with your proxy documentation to ensure correct wildcard specification in your environment.

More information on proxy configuration can be found in the NetApp [Knowledgbase](#).

Viewing Proxy URLs

You can view your proxy endpoint URLs by clicking the **Proxy Settings** link when choosing a data collector during onboarding, or the link under *Proxy Settings* on the **Help > Support** page. A table like the following is displayed.

Proxy Settings				
If your organization requires proxy usage for internet access, you need to understand your organization's proxy behavior and seek certain exceptions for Cloud Insights to work. The simplest way is to add the following domains to the exception list:				
Hostname	Port	Protocol	Methods	Endpoint URL Purpose
qtrjkso.proxyserver.cloudinsights-dev.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Tenant
00b1100.1234.abcd.12bc.a1b2c3ef56a7.proxyserver.cloudinsights-dev.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Acquisition Unit Ingestion
aulogin.proxyserver.cloudinsights-dev.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Acquisition Unit Authentication
portal.proxy.cloud.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Gateway

[Close](#)

If you have Workload Security on your tenant, the configured endpoint URLs will also be displayed in this list.

Adding a Windows Acquisition Unit

Steps for Windows Acquisition Unit Installation

1. Log in to the Acquisition Unit server/VM as a user with Administrator permissions.
2. On that server, open a browser window and log in to your Data Infrastructure Insights environment as Administrator or Account Owner.
3. Click **Observability > Collectors > Acquisition Units > +Acquisition Unit**.

The system displays the *Install Acquisition Unit* dialog. Choose Windows.

Install Acquisition Unit

Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.

What Operating System or Platform Are You Using?

Windows ▼ Windows Versions Supported ⓘ Production Best Practices ⓘ

Installation Instructions

[Need Help?](#)

1 [Download Installer \(Windows 64-bit\)](#)

2 [Copy Access Key](#)

This access key is a unique key valid for 24 hours for this Acquisition Unit only.

[Reveal Access Key](#)

3 Paste access key into installer when prompted.

4 Please ensure you have copied and pasted the access key into the installer.

[Have a Proxy Server?](#)

- 1 Verify that the server or VM hosting the Acquisition Unit meets the recommended system requirements.
- 2 Verify that the server is running a supported version of Windows. Click *OS Versions Supported (i)* for a list of supported versions.
- 3 Click the **Download Installer (Windows 64-bit)** button.
- 4 Copy the Access Key. You will need this during the Installation.
- 5 On the Acquisition Unit server/VM, execute the downloaded installer.
- 6 Paste the Access Key into the installation wizard when prompted.
- 7 During installation, you will be presented with the opportunity to provide your proxy server settings.

After you finish

- Click * > Observability > Collectors > Acquisition units* to check the status of Acquisition Units.
- You can access the Acquisition Unit log in <install dir>\Cloud Insights\Acquisition Unit\log\acq.log
- Use the following script to stop, start, restart, or check the status of the Acquisition Unit:

```
cloudinsights-service.sh
```

Proxy Configuration

The Acquisition Unit uses 2-way/mutual authentication to connect to the Data Infrastructure Insights server. The client certificate must be passed to the Data Infrastructure Insights server to be authenticated. To accomplish this, the proxy must be set up to forward the https request to the Data Infrastructure Insights server without decrypting the data.

The simplest way to do this is to specify wildcard configuration in your proxy/firewall to communicate with Data

Infrastructure Insights, for example:

`*.cloudinsights.netapp.com`



The use of an asterisk (*) for wildcard is common, but your proxy/firewall configuration may use a different format. Check with your proxy documentation to ensure correct wildcard specification in your environment.

More information on proxy configuration can be found in the NetApp [Knowledgbase](#).

Viewing Proxy URLs

You can view your proxy endpoint URLs by clicking the **Proxy Settings** link when choosing a data collector during onboarding, or the link under *Proxy Settings* on the **Help > Support** page. A table like the following is displayed.

Proxy Settings					X
If your organization requires proxy usage for internet access, you need to understand your organization's proxy behavior and seek certain exceptions for Cloud Insights to work. The simplest way is to add the following domains to the exception list:					Close
Hostname	Port	Protocol	Methods	Endpoint URL Purpose	
qtrjkso.proxyserver.cloudinsights-dev.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Tenant	
00b1100.1234.abcd.12bc.a1b2c3ef56a7.proxyserver.cloudinsights-dev.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Acquisition Unit Ingestion	
aulogin.proxyserver.cloudinsights-dev.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Acquisition Unit Authentication	
portal.proxy.cloud.netapp.com	443	https	GET, POST, PATCH, PUT, DELETE	Gateway	

If you have Workload Security on your tenant, the configured endpoint URLs will also be displayed in this list.

Uninstalling an Acquisition Unit

To uninstall the Acquisition Unit software, do the following:

Windows:

If you are uninstalling a **Windows** acquisition unit:

1. On the Acquisition Unit server/VM, open Control Panel and choose **Uninstall a Program**. Select the Data Infrastructure Insights Acquisition Unit program for removal.
2. Click Uninstall and follow the prompts.

Linux:

If you are uninstalling a **Linux** acquisition unit:

1. On the Acquisition Unit server/VM, run the following command:

```
sudo cloudinsights-uninstall.sh -p
```

2. For help with uninstall, run:

```
sudo cloudinsights-uninstall.sh --help
```

Windows and Linux:

After uninstalling the AU:

1. In Data Infrastructure Insights, go to **Observability > Collectors** and select the ***Acquisition Units** tab.
2. Click the Options button to the right of the Acquisition Unit you wish to uninstall, and select *Delete*. You can delete an Acquisition Unit only if there are no data collectors assigned to it.



You cannot delete an Acquisition Unit (AU) that has data collectors connected to it. Move all of the AU's data collectors to another AU (edit the collector and simply select a different AU) before deleting the original AU.

An Acquisition unit with a star next to it is being used for device resolution. Before removing this AU, you must select another AU to use for Device Resolution. Hover over a different AU and open the "three dots" menu to select "Use for Device Resolution".



Reinstalling an Acquisition Unit

To re-install an Acquisition Unit on the same server/VM, you must follow these steps:

Before you begin

You must have a temporary Acquisition Unit configured on a separate server/VM before re-installing an Acquisition Unit.

Steps

1. Log in to the Acquisition Unit server/VM and uninstall the AU software.
2. Log into your Data Infrastructure Insights environment and go to **Observability > Collectors**.
3. For each data collector, click the Options menu on the right and select *Edit*. Assign the data collector to the temporary Acquisition Unit and click **Save**.

You can also select multiple data collectors of the same type and click the **Bulk Actions** button. Choose *Edit* and assign the data collectors to the temporary Acquisition Unit.

4. After all of the data collectors have been moved to the temporary Acquisition Unit, go to **Observability > Collectors** and select the **Acquisition Units** tab.

5. Click the Options button to the right of the Acquisition Unit you wish to re-install, and select **Delete**. You can delete an Acquisition Unit only if there are no data collectors assigned to it.
6. You can now re-install the Acquisition Unit software on the original server/VM. Click **+Acquisition Unit** and follow the instructions above to install the Acquisition Unit.
7. Once the Acquisition Unit has been re-installed, assign your data collectors back to the Acquisition Unit.

Viewing AU Details

The Acquisition Unit (AU) detail page provides useful detail for an AU as well as information to help with troubleshooting. The AU detail page contains the following sections:

- A **summary** section showing the following:
 - **Name** and **IP** of the Acquisition Unit
 - Current connection **Status** of the AU
 - **Last Reported** successful data collector poll time
 - The **Operating System** of the AU machine
 - Any current **Note** for the AU. Use this field to enter a comment for the AU. The field displays the most recently added note.
- A table of the AU's **Data Collectors** showing, for each data collector:
 - **Name** - Click this link to drill down into the data collector's detail page with additional information
 - **Status** - Success or error information
 - **Type** - Vendor/model
 - **IP** address of the data collector
 - Current **Impact** level
 - **Last Acquired** time - when the data collector was last successfully polled

Acquisition Unit Summary								
Name xp-linux	Connection Status OK - Need Help?	Operating System Linux	Note					
IP 10.197.120.145	Last Reported 2 minutes ago							
Data Collectors (3)								
<input type="checkbox"/> Name ↑ + Data Collector Bulk Actions Filter...								
	Status	Type	IP	Impact	Last Acquired			
<input type="checkbox"/> foo	Inventory failed	NetApp Data ONTAP 7-Mode	foo	Low	Never			
<input type="checkbox"/> xp-cisco	All successful	Cisco MDS Fabric Switches	10.197.136.66		2 minutes ago			
<input type="checkbox"/> xpcdot26	All successful	NetApp ONTAP Data Management Software	10.197.136.26		8 minutes ago			

For each data collector, you can click on the "three dots" menu to Clone, Edit, Poll, or Delete the data collector. You can also select multiple data collectors in this list to perform bulk actions on them.

To restart the Acquisition Unit, click the **Restart** button at the top of the page. Drop down this button to attempt to **Restore Connection** to the AU in the event of a connection problem.

Configuring an Agent to Collect Data (Windows/Linux)

Data Infrastructure Insights uses [Telegraf](#) as its agent for collection of integration data. Telegraf is a plugin-driven server agent that can be used to collect and report metrics, events, and logs. Input plugins are used to collect the desired information into the agent by accessing the system/OS directly, by calling third-party APIs, or by listening to configured streams (i.e. Kafka, statsD, etc). Output plugins are used to send the collected metrics, events, and logs from the agent to Data Infrastructure Insights.

For information on installing on Kubernetes, see the [NetApp Kubernetes Monitoring Operator](#) page.



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)**.



If you want to verify the installation files before installing the Agent, see the section below on [Verifying Telegraf Package Checksums](#).

Installing an Agent

If you are installing a Service data collector and have not yet configured an Agent, you are prompted to first install an Agent for the appropriate Operating System. This topic provides instructions for installing the Telegraf agent on the following Operating Systems:

- [Windows](#)
- [RHEL and CentOS](#)
- [Ubuntu and Debian](#)

To install an agent, regardless of the platform you are using, you must first do the following:

1. Log into the host you will use for your agent.
2. Log in to your Data Infrastructure Insights environment and navigate to **Observability > Collectors**.
3. Click on **+Data Collector** and choose a data collector to install.
4. Choose the appropriate platform for your host (Windows, Linux)
5. Follow the remaining steps for each platform.



Once you have installed an agent on a host, you do not need to install an agent again on that host.



Once you have installed an agent on a server/VM, Data Infrastructure Insights collects metrics from that system in addition to collecting from any data collectors you configure. These metrics are gathered as ["Node" metrics](#).



If you are using a proxy, read the proxy instructions for your platform before installing the Telegraf agent.

Log Locations

Telegraf log messages are redirected from stdout to the following log files by default:

- RHEL/CentOS: /var/log/telegraf/telegraf.log
- Ubuntu/Debian: /var/log/telegraf/telegraf.log
- Windows: C:\Program Files\telegraf\telegraf.log

Windows

Pre-requisites:

- PowerShell must be installed
- If you are behind a proxy, you must follow the instructions in the **Configuring Proxy Support for Windows** section.

Configuring Proxy Support for Windows



If your environment uses a proxy, read this section before you install.



The steps below outline the actions needed to set the *http_proxy/https_proxy* environment variables. For some proxy environments, users may also need to set the *no_proxy environment* variable.

For systems residing behind a proxy, perform the following to set the *https_proxy* and/or *http_proxy* environment variable(s) **PRIOR** to installing the Telegraf agent:

```
[System.Environment]:SetEnvironmentVariable("https_proxy",  
"<proxy_server>:<proxy_port>", [System.EnvironmentVariableTarget]:Machine)
```

Installing the agent



Install Agent

Quickly setup an agent in your environment and immediately start monitoring data

Select existing API Access Token or create a new one

KEY1 (...Zqlk0c)

+ API Access Token

Installation Instructions

Need Help?

1

[Copy Agent Installer Snippet](#)

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)

[Reveal Agent Installer Snippet](#)

2

Open a PowerShell window as administrator and paste the snippet

3

[Complete Setup](#)

Steps to install agent on Windows:

1. Choose an Agent Access Key.
2. Copy the command block from the agent installation dialog. You can click the clipboard icon to quickly copy the command to the clipboard.
3. Open a PowerShell window
4. Paste the command into the PowerShell window and press Enter.
5. The command will download the appropriate agent installer, install it, and set a default configuration. When finished, it will restart the agent service. The command has a unique key and is valid for 24 hours.
6. Click **Finish** or **Continue**

After the agent is installed, you can use the following commands to start/stop the service:

Start-Service telegraf

Stop-Service telegraf

Uninstalling the Agent

To uninstall the agent on Windows, do the following in a PowerShell window:

1. Stop and delete the Telegraf service:

Stop-Service telegraf

sc.exe delete telegraf

2. Remove the certificate from the trustore:

```
cd Cert:\CurrentUser\Root
//rm E5FB7B68C08B1CA902708584C274F8EFC7BE8ABC
rm 1A918038E8E127BB5C87A202DF173B97A05B4996
```

3. Delete the `C:\Program Files\telegraf` folder to remove the binary, logs, and configuration files
4. Remove the `SYSTEM\CurrentControlSet\Services\EventLog\Application\telegraf` key from the registry

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop and delete the telegraf service:

```
Stop-Service telegraf
sc.exe delete telegraf
```

2. Delete the `SYSTEM\CurrentControlSet\Services\EventLog\Application\telegraf` key from the registry
3. Delete `C:\Program Files\telegraf\telegraf.conf`
4. Delete `C:\Program Files\telegraf\telegraf.exe`
5. [Install the new agent.](#)

RHEL and CentOS

Pre-requisites:

- The following commands must be available: curl, sudo, ping, sha256sum, openssl, and dmidecode
- If you are behind a proxy, you must follow the instructions in the [Configuring Proxy Support for RHEL/CentOS](#) section.

Configuring Proxy Support for RHEL/CentOS



If your environment uses a proxy, read this section before you install.



The steps below outline the actions needed to set the `http_proxy/https_proxy` environment variables. For some proxy environments, users may also need to set the `no_proxy` environment variable.

For systems residing behind a proxy, perform the following steps **PRIOR** to installing the Telegraf agent:


1. Set the `https_proxy` and/or `http_proxy` environment variable(s) for the current user:

```
export https_proxy=<proxy_server>:<proxy_port>
```

2. Create `/etc/default/telegraf`, and insert definitions for the `https_proxy` and/or `http_proxy` variable(s):

```
https_proxy=<proxy_server>:<proxy_port>
```

Installing the agent



Install Agent
Quickly setup an agent in your environment and immediately start monitoring data

Select existing API Access Token or create a new one

default_ingestion_api_key1 (...xEKVyK) ▼ **+ API Access Token** Production Best Practices ?

Installation Instructions Need Help?

- 1 For environments operating behind a proxy server, follow the instructions to [configure proxy support to install and run Telegraf](#).
- 2 **Copy Agent Installer Snippet**
This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)
 Reveal Agent Installer Snippet
- 3 Open a terminal window and paste the snippet in a Bash shell (requires curl, sudo, ping, sha256sum, and dmidecode).
- 4 **Complete Setup**

Steps to install agent on RHEL/CentOS:

1. Choose an Agent Access Key.
2. Copy the command block from the agent installation dialog. You can click the clipboard icon to quickly copy the command to the clipboard.
3. Open a Bash window
4. Paste the command into the Bash window and press Enter.
5. The command will download the appropriate agent installer, install it, and set a default configuration. When finished, it will restart the agent service. The command has a unique key and is valid for 24 hours.
6. Click **Finish** or **Continue**

After the agent is installed, you can use the following commands to start/stop the service:

If your operating system is using systemd (CentOS 7+ and RHEL 7+):

```
sudo systemctl start telegraf
sudo systemctl stop telegraf
```

If your operating system is not using systemd (CentOS 7+ and RHEL 7+):

```
sudo service telegraf start  
sudo service telegraf stop
```

Uninstalling the Agent

To uninstall the agent on RHEL/CentOS, in a Bash terminal, do the following:

1. Stop the Telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd  
(CentOS 7+ and RHEL 7+)  
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the Telegraf agent:

```
yum remove telegraf
```

3. Remove any configuration or log files that may be left behind:

```
rm -rf /etc/telegraf*  
rm -rf /var/log/telegraf*
```

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop the telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd  
(CentOS 7+ and RHEL 7+)  
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the previous telegraf agent:

```
yum remove telegraf
```

3. [Install the new agent.](#)

Ubuntu and Debian

Pre-requisites:

- The following commands must be available: curl, sudo, ping, sha256sum, openssl, and dmidecode
- If you are behind a proxy, you must follow the instructions in the **Configuring Proxy Support for Ubuntu/Debian** section.

Configuring Proxy Support for Ubuntu/Debian



If your environment uses a proxy, read this section before you install.



The steps below outline the actions needed to set the *http_proxy/https_proxy* environment variables. For some proxy environments, users may also need to set the *no_proxy environment* variable.

For systems residing behind a proxy, perform the following steps **PRIOR** to installing the Telegraf agent:

1. Set the *https_proxy* and/or *http_proxy* environment variable(s) for the current user:

```
export https_proxy=<proxy_server>:<proxy_port>
```

2. Create /etc/default/telegraf, and insert definitions for the *https_proxy* and/or *http_proxy* variable(s):

```
https_proxy=<proxy_server>:<proxy_port>
```

Installing the agent



Ubuntu & Debian

Install Agent

Quickly setup an agent in your environment and immediately start monitoring data

Select existing API Access Token or create a new one

default_ingestion_api_key1 (...xEKVyK)

+ API Access Token

Production Best Practices

Installation Instructions

Need Help?

- 1 For environments operating behind a proxy server, follow the instructions to [configure proxy support to install and run Telegraf](#).

- 2 [Copy Agent Installer Snippet](#)

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)

[Reveal Agent Installer Snippet](#)

- 3 Open a terminal window and paste the snippet in a Bash shell (requires curl, sudo, ping, sha256sum, and dmidecode).

- 4 [Complete Setup](#)

Steps to install agent on Debian or Ubuntu:

1. Choose an Agent Access Key.
2. Copy the command block from the agent installation dialog. You can click the clipboard icon to quickly copy the command to the clipboard.
3. Open a Bash window
4. Paste the command into the Bash window and press Enter.
5. The command will download the appropriate agent installer, install it, and set a default configuration. When finished, it will restart the agent service. The command has a unique key and is valid for 24 hours.
6. Click **Finish or Continue**

After the agent is installed, you can use the following commands to start/stop the service:

If your operating system is using systemd:

```
sudo systemctl start telegraf
sudo systemctl stop telegraf
```

If your operating system is not using systemd:

```
sudo service telegraf start
sudo service telegraf stop
```

Uninstalling the Agent

To uninstall the agent on Ubuntu/Debian, in a Bash terminal, run the following:

1. Stop the Telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd)
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the Telegraf agent:

```
dpkg -r telegraf
```

3. Remove any configuration or log files that may be left behind:

```
rm -rf /etc/telegraf*
rm -rf /var/log/telegraf*
```

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop the telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd)  
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the previous telegraf agent:

```
dpkg -r telegraf
```

3. [Install the new agent.](#)

Verifying Telegraf Package Checksums

The Data Infrastructure Insights agent installer performs integrity checks, but some users may want to perform their own verifications before installing the downloaded Telegraf binary. This can be done by downloading the installer and generating a checksum for the downloaded package, then comparing the checksum to the value shown in the install instructions.

Download the installer package without installing

To perform a download-only operation (as opposed to the default download-and-install), users can edit the agent installation command obtained from the UI and remove the “install” option.

Follow these steps:

1. Copy the Agent Installer snippet as directed.
2. Instead of pasting the snippet into a command window, paste it into a text editor.
3. Remove the trailing “--install” (Linux) or “-install” (Windows) from the command.
4. Copy the entire command from the text editor.
5. Now paste it into your command window (in a working directory) and run it.

Non-Windows (these examples are for Kubernetes; actual script names may vary):

- Download and install (default):

```
installerName=cloudinsights-ubuntu_debian.sh ... && ./$installerName  
--download --verify && sudo -E -H ./$installerName --install
```

- Download-only:

```
installerName=cloudinsights-ubuntu_debian.sh ... && ./$installerName  
--download --verify
```

Windows:

- Download and install (default):

```
!$(installerName=".\\cloudinsights-windows.ps1") ... -and $(if(((Get-  
FileHash $installerName).Hash).ToLower() -eq "INSTALLER_CHECKSUM ") {  
&$installerName -download -verify -install } else { Write-Host "Install  
script checksum does not match"})"
```

- Download-only:

```
!$(installerName=".\\cloudinsights-windows.ps1") ... -and $(if(((Get-  
FileHash $installerName).Hash).ToLower() -eq "INSTALLER_CHECKSUM ") {  
&$installerName -download -verify } else { Write-Host "Install script  
checksum does not match"})"
```

The download-only command will download all required artifacts from Data Infrastructure Insights to the working directory. The artifacts include, but may not be limited to:

- an installation script
- an environment file
- a Telegraf binary
- a signature for the Telegraf binary
- a public certificate to verify the binary signature

The install snippet downloaded and copied from DII automatically checksums the installation script, and the telegraf binary's signature is verified by the installation script.

Verify checksum value

To generate the checksum value, perform the following command for your appropriate platform:

- RHEL/Ubuntu:

```
sha256sum <package_name>
```

- Windows:

```
Get-FileHash telegraf.zip -Algorithm SHA256 | Format-List
```

Install the downloaded package

Once all of the artifacts have been satisfactorily verified, the agent installation can be initiated by running:

Non-Windows:

```
sudo -E -H ./<installation_script_name> --install
```

Windows:

```
.\cloudinsights-windows.ps1 -install
```

Creating and Using API Access Tokens

To create an API access token for Telegraf data ingestion please do one of the following:

Create through the Data Collector Installation Page

1. Navigate to the Data Collector install page for the platform you wish to use (Windows, Linux).
2. Create a token with the + API Access Token button.
3. Input a name and click save.
4. The token name should now be selected in the dropdown, and will be used when installing the collector.

Manually create an API access token

1. Navigate to Admin>API Access.
2. Click + API Access Token.
3. Input a name and optionally a description.
4. Under “What type of APIs will this token be used to call?” select “Data Ingestion” only and then unselect “Acquisition Unit”.
5. Under “Permissions” select Read/Write.
6. Unselect “Automatically rotate tokens for Kubernetes”.

To use your newly created API access token, select it from the “Select existing API Access Token or create a new one” dropdown on the installer page.

Please note that only tokens with the following properties can be used:

- API Type: “Data Ingestion” only
- Permissions: Read/Write
- Kubernetes Auto Rotation: Off

Troubleshooting

Some things to try if you encounter problems setting up an agent:

Problem:	Try this:
<p>After configuring a new plugin and restarting Telegraf, Telegraf fails to start up. The logs indicate that an error resembling the following:</p> <pre data-bbox="137 312 816 466">"[telegraf] Error running agent: Error loading config file /etc/telegraf/telegraf.d/cloudinsights-default.conf: plugin outputs.http: line <linenumber>: configuration specified the fields ["use_system_proxy"], but they weren't used"</pre>	<p>The installed Telegraf version is outdated. Follow the steps on this page to Upgrade the Agent for your appropriate platform.</p>
<p>I ran the installer script on an old installation and now the agent isn't sending data</p>	<p>Uninstall the telegraf agent and then re-run the installation script. Follow the Upgrade the Agent steps on this page for your appropriate platform.</p>
<p>I already installed an agent using Data Infrastructure Insights</p>	<p>If you have already installed an agent on your host/VM, you do not need to install the agent again. In this case, simply choose the appropriate Platform and Key in the Agent Installation screen, and click on Continue or Finish.</p>
<p>I already have an agent installed but not by using the Data Infrastructure Insights installer</p>	<p>Remove the previous agent and run the Data Infrastructure Insights Agent installation, to ensure proper default configuration file settings. When complete, click on Continue or Finish.</p>

Additional information may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

Configuring Data Collectors

You configure Data Collectors in your Data Infrastructure Insights environment to collect data from devices in the data center.

Before you begin

- You must have configured an Acquisition Unit before you can start collecting data.
- You need credentials for the devices from which you are collecting Data.
- Device network addresses, account information, and passwords are required for all devices you are collecting data from.

Steps

1. From the Data Infrastructure Insights menu, click **Observability > Collectors**

The system displays the available Data Collectors arranged by vendor.

2. Click **+ Collector** and select the data collector to configure.

In the dialog box you can configure the data collector and add an Acquisition Unit.

3. Enter a name for the data collector.
4. Click **Advanced Configuration** to add additional configuration fields. (Not all data collectors require advanced configuration.)

5. Click **Test Configuration** to verify that the data collector is properly configured.
6. Click **Add Collector** to save the configuration and add the data collector to your Data Infrastructure Insights tenant.

It may take up to two poll periods before data from the service is displayed in dashboards or available for querying.

- 1st inventory poll: immediately
- 1st performance data poll to establish a baseline: immediately after inventory poll
- 2nd performance poll: within 15 seconds after completion of 1st performance poll

Polling then proceeds according to the configured inventory and performance poll intervals.

Determining data collector acquisition status

Because data collectors are the primary source of information for Data Infrastructure Insights, it is imperative that you ensure that they remain in a running state.

Data collector status is displayed in the upper right corner of any asset page as the message "Acquired N minutes ago", where N indicates the most recent acquisition time of the asset's data collector(s). The acquisition time/date is also displayed.

Clicking on the message displays a table with data collector name, status, and last successful acquisition time. If you are signed in as an Administrator, clicking on the data collector name link in the table takes you to detail page for that data collector.

Managing configured data collectors

The Installed Data Collectors page provides access to the data collectors that have been configured for Data Infrastructure Insights. You can use this page to modify existing data collectors.

Steps

1. In the Data Infrastructure Insights menu, click **Observability > Collectors**

The Available Data Collectors screen is displayed.

2. Click **Installed Data Collectors**

A list of all of the installed Data Collectors is displayed. The list provides collector name, status, the IP address the collector is accessing, and when data was last acquired from the a device. Action that can be performed on this screen include:

- Control polling
- Change data collector credentials
- Clone data collectors

Controlling Data Collector polling

After making a change to a data collector, you might want it to poll immediately to check your changes, or you might want to postpone the data collection on a data collector for one, three, or five days while you work on a problem.

Steps

1. In the Data Infrastructure Insights menu, click **Observability > Collectors**
2. Click **Installed Data Collectors**
3. Select the check box to the left of the Data Collector you want to change
4. Click **Bulk Actions** and select the polling action you want to take.

Bulk actions can be performed simultaneously on multiple Data Collectors. Select the data collectors, and chose the action to perform from the **Bulk Action** menu.

Editing data collector information

You can edit existing data collector setup information.

To edit a single data collector:

1. In the Data Infrastructure Insights menu, click **Observability > Collectors** to open the list of installed Data Collectors.
2. In the options menu to the right of the data collector you want to modify, click **Edit**.

The Edit Collector dialog is opened.

3. Enter the changes and click **Test Configuration** to test the new configuration or click **Save** to save the configuration.

You can also edit multiple data collectors:

1. Select the check box to the left of each data collector you want to change.
2. Click the **Bulk Actions** button and choose **Edit** to open the Edit data Collector dialog.
3. Modify the fields as above.



The data collectors selected must be the same vendor and model, and reside on the same Acquisition Unit.

When editing multiple data collectors, the Data Collector Name field shows “Mixed” and cannot be edited. Other fields such as user name and password show “Mixed” and can be edited. Fields that share the same value across the selected data collectors show the current values and can be edited.

When editing multiple data collectors, the **Test Configuration** button is not available.

Cloning data collectors

Using the clone facility, you can quickly add a data source that has the same credentials and attributes as another data source. Cloning allows you to easily configure multiple instances of the same device type.

Steps

1. In the Data Infrastructure Insights menu, click **Observability > Collectors**.
2. Click **Installed Data Collectors**.
3. Click the check box to the left of the data collector you want to copy.
4. In the options menu to the right of the selected data collector, click **Clone**.

The Clone Data Collector dialog is displayed.

5. Enter new information in the required fields.
6. Click **Save**.

After you finish

The clone operation copies all other attributes and settings to create the new data collector.

Performing bulk actions on data collectors

You can simultaneously edit some information for multiple data collectors. This feature allows you to initiate a poll, postpone polling, and resume polling on multiple data collectors. In addition, you can delete multiple data collectors.

Steps

1. In the Data Infrastructure Insights menu, click **Observability > Collectors**
2. Click **Installed Data Collectors**
3. Click the check box to the left of the data collectors you want to modify.
4. In the options menu to the right, click the option you want to perform.

After you finish

The operation you selected is performed on the data collectors. When you chose to delete data collectors, a dialog is displayed requiring you to conform the action.

Researching a failed data collector

If a data collector has failure message and a High or Medium Impact, you need to research this problem using the data collector summary page with its linked information.

Use the following steps to determine the cause of failed data collectors. Data collector failure messages are displayed on the **Admin** menu and on the **Installed Data Collectors** page.

Steps

1. Click **Admin > Data Collectors > Installed Data Collectors**.
2. Click the linked Name of the failing data collector to open the Summary page.
3. On the Summary page, check the Comments area to read any notes that might have been left by another engineer who might also be investigating this failure.
4. Note any performance messages.
5. Move your mouse pointer over the segments of the Event Timeline graph to display additional information.
6. Select an error message for a Device and displayed below the Event Timeline and click the Error details icon that displays to the right of the message.

The Error details include the text of the error message, most likely causes, information in use, and suggestions of what can be tried to correct the problem.

7. In the Devices Reported By This Data Collector area, you might filter the list to display only devices of interest, and you can click the linked **Name** of a device to display the asset page for that device.
8. When you return to the data collector summary page, check the **Show Recent Changes** area at the bottom of the page to see if recent changes could have caused the problem.

Copyright information

Copyright © 2026 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.