



**Splunk Cloud Platform™**

**Splunk Cloud Platform Admin Manual 9.2.2406**

Generated: 10/16/2024 1:00 pm

# Table of Contents

<b>Get Started Managing Splunk Cloud Platform</b> .....	<b>1</b>
Welcome to the Splunk Cloud Platform Admin Manual.....	1
Splunk Cloud Platform deployment types.....	4
Splunk Cloud Platform Quick Start.....	6
Add a global banner to your Splunk Cloud Platform deployment.....	8
Determine your Splunk Cloud Platform Experience.....	11
Change the UI theme of Splunk Cloud Platform.....	12
<b>Get Data Into Splunk Cloud Platform</b> .....	<b>13</b>
Introduction to Getting Data In.....	13
Get Amazon Web Services (AWS) data into Splunk Cloud Platform.....	17
Get Microsoft Azure data into Splunk Cloud Platform.....	27
Get *nix data into Splunk Cloud Platform.....	32
Get Windows Data into Splunk Cloud Platform.....	37
Forward data from files and directories to Splunk Cloud Platform.....	47
Upgrade your Forwarders.....	48
<b>Configure your Splunk Cloud Platform Deployment</b> .....	<b>53</b>
Configure IP allow lists using Splunk Web.....	53
Configure Dashboards Trusted Domains List.....	55
Manage custom bookmarks.....	58
Configure webhook allow list using Splunk Web.....	59
Configure limits using Splunk Web.....	60
Manage HTTP Event Collector (HEC) tokens in Splunk Cloud Platform.....	63
<b>Manage Splunk Cloud Platform Users and Roles</b> .....	<b>64</b>
Manage Splunk Cloud Platform users and roles.....	64
Configure Splunk Cloud to use SAML for authentication tokens.....	64
<b>Monitor your Splunk Cloud Platform Deployment</b> .....	<b>65</b>
Introduction to the Cloud Monitoring Console.....	65
Use the Overview dashboard.....	66
Use the Health dashboard.....	69
Use the Maintenance dashboard.....	73
Use the Alerts dashboard.....	76
Use the Indexing dashboards.....	79
Use the Search dashboards.....	87
Use the Usage dashboards.....	94
Use the License Usage dashboards.....	96
Use the Forwarder dashboards.....	112
Use the Workload Management Monitoring dashboard.....	116
Monitor your deployment with the splunkd health report.....	117
How Splunk monitors Splunk Cloud Platform.....	123
Manage your Splunk Cloud Platform capacity.....	125

# Table of Contents

<b>Optimize indexing and search processes.....</b>	<b>126</b>
Optimize indexing and search processes.....	126
<b>Manage your Indexes and Data in Splunk Cloud Platform.....</b>	<b>128</b>
Manage Splunk Cloud Platform indexes.....	128
Store expired Splunk Cloud Platform data in your private archive.....	134
Store expired Splunk Cloud Platform data in a Splunk-managed archive.....	148
Manage indexes on Splunk Cloud Platform Classic Experience.....	156
<b>Manage Apps and Add-ons in Splunk Cloud Platform.....</b>	<b>161</b>
Install apps on your Splunk Cloud Platform deployment.....	161
Manage private apps on your Splunk Cloud Platform deployment.....	163
Manage the Splunk Product Guidance app on your Splunk Cloud Platform deployment.....	173
Manage a rolling restart in Splunk Cloud Platform.....	174
<b>Configure Search Settings in Splunk Cloud Platform.....</b>	<b>178</b>
Configure hybrid search.....	178
Set limits for concurrent scheduled searches.....	180
<b>Manage Search Workloads in Splunk Cloud Platform.....</b>	<b>182</b>
Workload Management overview.....	182
Configure workload rules.....	184
Configure admission rules to prefilter searches.....	188
Manually assign searches to workload pools.....	193
Workload Management examples.....	196
<b>Enable Automatic UI Updates.....</b>	<b>199</b>
Enable automatic UI updates.....	199

# Get Started Managing Splunk Cloud Platform

## Welcome to the Splunk Cloud Platform Admin Manual

This manual contains information to help you administer your Splunk Cloud Platform deployment.

Splunk Cloud Platform delivers the benefits of-winning Splunk® Enterprise as a cloud-based service. Using Splunk Cloud Platform, you gain the functionality of the Splunk Enterprise platform for collecting, searching, monitoring, reporting, and analyzing all of your real-time and historical machine data using a cloud service that is centrally and uniformly delivered by Splunk to its large number of cloud customers, from Fortune 100 companies to small and medium-size businesses. Splunk manages and updates the Splunk Cloud Platform service uniformly, so all customers of Splunk Cloud Platform receive the most current features and functionality.

If you're a new Splunk Cloud Platform administrator, or a Splunk Enterprise administrator planning a move to Splunk Cloud Platform, visit the Splunk Cloud Platform Migration section of the Splunk Lantern Resource Hub for helpful guidance from Splunk experts.

If you're interested in using Splunk Cloud Platform with Splunk Observability Cloud, see Splunk Observability Cloud and the Splunk platform.

### Splunk Cloud Platform features

Splunk Cloud Platform is available in multiple regions. For details please refer to the Splunk Cloud Platform Service Description. The following table lists major features of Splunk Cloud Platform.

Feature	Description
Data Collection	<p>You can send data to Splunk Cloud Platform as follows:</p> <p><b>Using Splunk forwarders:</b> There are two types of forwarder software: universal forwarder and heavy forwarder. In most situations, the universal forwarder is the best forwarder for Splunk Cloud Platform since it includes the essential components that it needs to forward data, uses significantly fewer hardware resources and is inherently scalable. For certain use cases when data needs to be parsed prior to forwarding or data needs to be forwarded based on criteria such as source or type of event, a heavy forwarder is required.</p> <p><b>Using HTTP Event Collector (HEC):</b> HEC lets you send data and application events using a token-based authentication mode to Splunk Cloud Platform over the Secure HTTP (HTTPS) protocol. You can generate a token and then configure a logging library or HTTPS client with the token to send data to HEC in a specific format.</p> <p><b>Using AWS Kinesis Data Firehose:</b> AWS Kinesis Data Firehose is a fully managed, scalable, and serverless option for streaming data from various AWS services directly into Splunk Cloud Platform.</p>
Ingestion	Splunk Cloud Platform indexes incoming data so you can search it. During indexing, data is partitioned into logical indexes, which you can configure to facilitate searching and control users' access to data in specific indexes.

Feature	Description
Inputs Data Manager	<p>The <b>Inputs Data Manager (IDM)</b> is a hosted solution for Splunk Cloud Platform that supports scripted and modular inputs for customers on the Classic Experience. To use scripted or modular inputs, you must package them in a private app and request that Support uploads the app to your IDM. The app must be vetted like any other private app.</p> <p>The IDM is not a one-to-one replacement for a heavy forwarder. You must still use a heavy forwarder if you need to perform parsing or activities other than standard scripted and modular data inputs. As a best practice, cloud-based add-ons should be installed on an IDM, and on-premise-based add-ons should be installed on a forwarder or heavy forwarder. <b>Note:</b> If the add-on is tightly integrated with an Enterprise Security search head, you should not use the IDM.</p> <p>For more information on using the IDM with the Classic Experience, see <a href="#">Work with Inputs Data Manager</a>.</p> <p>For more information about the Splunk Cloud Platform Classic and Victoria Experiences, see the following:</p> <ul style="list-style-type: none"> <li>• <a href="#">Determine your Splunk Cloud Platform Experience</a></li> <li>• Experience designations in the <i>Splunk Cloud Platform Service Description</i></li> </ul>
Retention	<p>When you send data to Splunk Cloud Platform, it is stored in indexes and you can self-manage your Splunk Cloud Platform indexes settings using the Indexes page in Splunk Web. Splunk Cloud Platform retains data based on index settings that enable you to specify when data is to be deleted. To configure different data retention settings for different sources of data, store the data in separate indexes according to the desired retention policy. You can configure different data retention policies for individual indexes according to your auditing and compliance requirements. Each index allows you to specify the maximum age of events in the index (specified in the Retention (days) field) on the Indexes page uses to determine when to delete data. When the index reaches the specified maximum age, the oldest data is deleted.</p>
Search	<p>Splunk Cloud Platform allows you to search and navigate all of the machine data that you ingest into the service. Searches can be done using the Splunk Search Processing Language (SPL), or using alternative ways to display and analyze data graphically without composing SPL queries. Searches can be ad-hoc and scheduled, with results in the form of visualizations, reports, and alerts.</p>
Reports	<p>Reports are saved searches and pivots. You can run reports on an ad hoc basis, schedule them to run on a regular interval, or set scheduled reports to generate alerts when the results of their runs meet particular conditions. You can add reports to dashboards as dashboard panels. (More information)</p>
Dashboards	<p>Dashboards are made up of panels that contain modules such as search boxes, fields, charts, tables, forms, and so on. Dashboard panels are usually hooked up to saved searches or pivots. They can display the results of completed searches as well as data from backgrounded real-time searches. (More information)</p>
Administration	<p>You can use Splunk Cloud Platform to perform the following administrative tasks:</p> <ul style="list-style-type: none"> <li>• Manage indexes. See <a href="#">Manage Splunk Cloud Platform indexes</a>.</li> <li>• Manage <b>Splunk apps</b> and add-ons (some apps require you to make a Support request to install). See <a href="#">Install apps in your Splunk Cloud Platform deployment</a>.</li> <li>• Manage users and their roles. See <a href="#">Manage Splunk Cloud Platform users and roles</a>.</li> </ul> <p>In Splunk Cloud Platform, you usually use Splunk Web to perform administrative tasks. Unlike Splunk Enterprise, you do not have access to the command line or file system of your Splunk Cloud Platform deployment, so you cannot use CLI commands or manually edit .conf files. If there is a task that you need to perform, but cannot do so from the Splunk Web interface, you can file a ticket using the Support Portal.</p>
REST API access	<p>Some administrative tasks can be done using the Splunk REST API. Splunk Cloud Platform supports a subset of the REST API endpoints available in Splunk Enterprise. For more information on supported REST endpoints, see the <i>REST API Reference Manual</i>. To use the REST API, you must have a paid subscription to Splunk Cloud Platform.</p>

Feature	Description
	<p>To enable the Splunk REST API and SDKs:</p> <ol style="list-style-type: none"> <li>1. Submit a support case on the Support Portal to request access. You can specify a range of IP addresses to control who can access the REST API.</li> <li>2. After you have gained access, use the following URL:  <a href="https://&lt;deployment-name&gt;.splunkcloud.com:8089">https://&lt;deployment-name&gt;.splunkcloud.com:8089</a></li> </ol>

## Supported browsers

Splunk Cloud Platform supports the following browsers:

- Chrome (latest)
- Firefox (latest)
- Safari (latest)
- Microsoft Edge (latest)

## Use IPv6 support with dual-stack configuration

Splunk Forwarder now supports IPv6 as a dual-stack IPv4/IPv6 configuration as an Early Access feature. If you are using dual-stack IPv6 support in your configuration, your Splunk Support engineer can help you configure Splunk Cloud Platform to use IPv6. See Use dual-stack IPv6 support for more information.

Note that during Early Access releases, Splunk products may have limitations on customer access, features, maturity, and regional availability. For additional information on Early Access please contact your Splunk representative.

## Third Party Documentation

As a convenience, this document includes instructions for using non-Splunk software to get data from varying platforms into Splunk Cloud Platform. Splunk does not warrant the performance of non-Splunk software based on the instructions in this documentation. Please review the product documentation provided by the other software providers before following these instructions. The screenshots and instructions in this documentation are updated on a best-effort basis.

## Splunk Technical Support

Splunk Standard Support is included in every Splunk Cloud Platform subscription. For details about the levels of technical support provided, read Support Programs. Only authorized support contacts from your company can open cases. Your Splunk support agreement specifies who your authorized contacts are. Your Support contract specifies a number of authorized contacts, and an expiration date. One of your contacts is a Support portal administrator, who can update the list. Only an authorized contact can open a case and track its status. An authorized contact can file a case in one of two ways:

- Log in to splunk.com and navigate to the **Support Portal**.
- In Splunk Cloud Platform, click **About** and select **File a Bug**.

## Splunk Support portal

Splunk Cloud Platform users who are registered with Splunk Support as a Splunk Cloud entitlement contact for their account can manage operational contacts for their account and file support cases using the Support portal. Operational

contacts are the people in your organization who are notified when their Splunk Cloud Platform environment undergoes maintenance or experiences an event that affects performance.

For more information on working with Splunk Support and understanding the difference between entitlement and operational contacts, visit the Splunk website and log into the Support Portal.

#### **To manage operational contacts:**

1. Go to **My Operational Contacts** in the Support Portal.
2. Follow the instructions on the page to add, edit, and remove operational contacts for your Splunk Cloud Platform environment.

#### **To file a case on the Support Portal:**

1. In the Support portal select **Get started** and then **Create Case**.
2. Follow the process to open a support case for the **Splunk Cloud** product.

Splunk Support replies to the case creator by email. You can update the case by replying to the email (be sure to keep the tracking ID in the email subject line). You can also update the case, check on its status, or close a case using the Support Portal.

## **Splunk community**

The Splunk user community is a great resource. Check out Splunk Answers, where you can ask and answer questions about the product. There are also a number of other ways to get involved in the Splunk community, such as user groups or the Splunk Trust. For more information about getting involved with the Splunk community, see the Community portal.

## **Learn more about Splunk products**

For detailed information about the Splunk platform, see the following resources:

- Splunk Docs is Splunk user documentation.
- Splunk Training offers courses on-site, off-site, and on the Web.
- Splunk Lantern provides clear and actionable guidance for many use cases.
- Splunk Videos offer training and demos on a variety of topics.

### ***About apps and add-ons***

Apps and add-ons extend the power of Splunk products to help you get value from your data faster. To browse Splunk apps and add-ons, see Splunkbase. If you develop your own app, read Splunk Developer Guidance. For an example of a properly constructed app, see the Splunk Reference App.

## **Splunk Cloud Platform deployment types**

Splunk Cloud Platform version 8.x and higher offers two deployment types:

- **Purchased:** Paid subscription to the Splunk Cloud Platform service.
- **Free trial:** Limited duration free trial initiated from the Splunk website.

The following table shows the fundamental differences between the free trial and purchased deployment types.

If the URL you use to access Splunk Cloud Platform begins with `https://prd-p-*` you are using a free trial. If you are using the Splunk Cloud Platform free trial with other Splunk Cloud Platform products, services, or offerings, additional limitations or restrictions may apply. For more information on supported functionality for Splunk Cloud free trials, see the installation and configuration documentation for that product, service, or offering.

Feature	Free trial	Purchased
<b>Access your Splunk Cloud Platform deployment</b>	Log into your Splunk account on the Splunk web site ( <a href="http://www.splunk.com">www.splunk.com</a> ) and go to the <b>Instances</b> page.	Using the URL you specified when you purchased Splunk Cloud Platform from Splunk Sales. This URL is included in the Welcome email you received when your deployment was enabled for you.
<b>Create inputs</b> (including file uploads, HTML event ingestion, and app-related inputs)	Configure forwarders on-premises or use Splunk Web	Configure forwarders on-premises or contact Splunk Support
<b>Data ingestion</b> (daily maximum)	5 GB per day for 15 days	Per your contract
<b>Data export, retention, and deletion</b>	Data that you add to your Splunk Cloud Platform free trial cannot be exported from the trial instance. When your free trial expires, the instance and all data in it is deleted. Currently, trials cannot be converted to paid accounts.	Splunk Cloud Platform retains data based on index settings that enable you to specify when data is to be deleted. For more information, see the Storage section in the <i>Splunk Cloud Platform Service Description</i> .
<b>Send data securely using Splunk Universal Forwarder</b>	Yes	Yes
<b>Active Directory/Single sign-on integration</b>	Not supported	Supported
<b>Splunk Cloud Platform regions</b>	US East	Multiple including GovCloud
<b>Allow list and deny list IP addresses</b>	No	Yes
<b>Premium apps available</b>	No	Yes
<b>Concurrent search limit</b>	10 max	Depends on the topology you purchase
<b>REST API URL</b>	Free Trial users do not have access to the REST API. Contact your Splunk sales representative for possible alternative methods during your trial period.	Use the following URL for purchased deployments. If necessary, submit a support case to open port 8089 on your deployment <code>https://&lt;deployment-name&gt;.splunkcloud.com:8089</code>
<b>HTTP Event Collector (HEC) URL</b>	<code>https://inputs.&lt;host&gt;.splunkcloud.com:8088/&lt;endpoint&gt;</code>	<code>https://http-inputs-&lt;host&gt;:443/&lt;endpoint&gt;</code> or <code>https://http-inputs.&lt;host&gt;.splunkcloud.com:443/&lt;endpoint&gt;</code> for Splunk Cloud Platform on Google Cloud
	Not supported	Supported



Feature	Free trial	Purchased
Inputs Data Manager (IDM)		

## Splunk Cloud Platform Quick Start

This topic shows you the basic steps required to start using your Splunk Cloud Platform deployment, and provides a simple quick start tutorial to help you get up and running quickly.

To get started with your Splunk Cloud Platform deployment, follow these high-level steps:

- Log in
- Get data in
- Search and manage your data

### ***Log in to Splunk Cloud Platform***

To log in to your Splunk Cloud Platform deployment, you must use the dedicated Splunk Cloud Platform URL and log in credentials provided to you in the "Welcome to Splunk Cloud Platform" email you received when you opened your account.

### ***Get data into Splunk Cloud Platform***

To get data into Splunk Cloud Platform, the most common approach is to install the Splunk Universal Forwarder on the machines where your source data resides, and configure them to send data to Splunk Cloud Platform. You can also upload files, or monitor files and inputs. For more information on the options available for getting data into Splunk Cloud Platform, see [Introduction to Getting Data In](#).

### ***Search and manage your data***

After you get your data into Splunk Cloud Platform, you can search the data to create reports, display the results using dashboards and visualizations, and set alerts that trigger when specific conditions are met. For detailed information, see the following manuals.

- Search Manual
- Reporting Manual
- Alerting Manual
- Dashboards and visualizations

## **Quick start tutorial**

If you are new to Splunk Cloud Platform and want to get started quickly, follow the steps in this brief tutorial to get some data into your Splunk Cloud Platform deployment and start searching it.

### ***What you need***

- Your Splunk Cloud Platform URL and log in credentials. See [Log in to Splunk Cloud Platform](#).
- A standard log file to use as sample data for this exercise, such as a `/var/log/messages` file on a Unix machine, or a text file in `C:\Windows\System32\LogFiles` on a Windows computer.

## Step 1. Log in to Splunk Cloud Platform

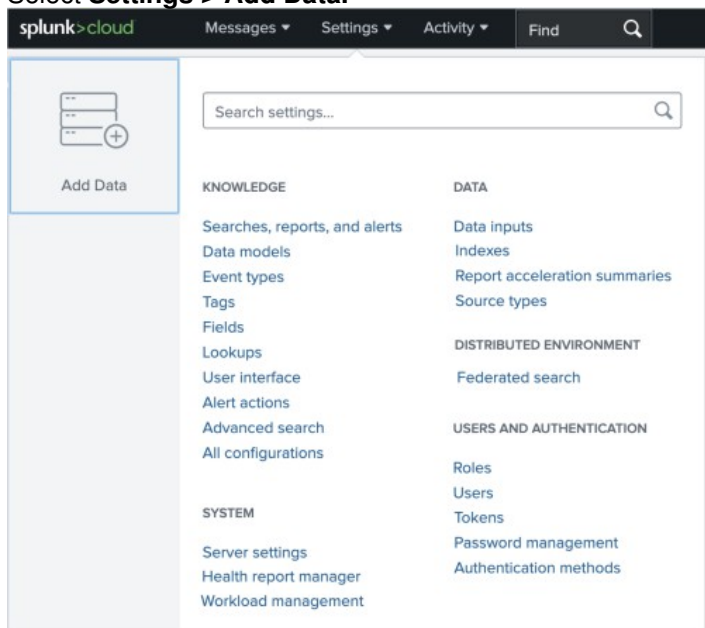
To log in to Splunk Cloud Platform:

1. In your web browser, navigate to your Splunk Cloud Platform URL. For example, <https://mycompany.splunkcloud.com> **OR** <https://prd-p-njqblk23gjdj.cloud.splunk.com>
2. Enter the credentials provided to you when you opened your account.  
The Splunk Web UI appears. You can now interact with your Splunk Cloud Platform deployment.

## Step 2. Upload a file

In Splunk Web, follow these steps:

1. To create a test index where you can store test data, click **Settings > Indexes**.
2. Click **New Indexes** and assign the index a name. To minimize resource consumption, specify a small size and retention period.
3. Select **Settings > Add Data**.



4. Click **Upload**.
5. Click **Select File**, browse to a log file on your computer, and click **Open**. The file is uploaded. Click **Next**.
6. On the **Set Source Type** page, select the correct source type for the file you uploaded, or, if none is appropriate, specify a name for the new source type and click **Next**.
7. On the **Input Settings** page, select your test index.
8. Click **Review** and verify your settings.
9. Click **Submit**.

After your data is uploaded, Splunk Web displays a "Success" message. You can now start searching your data.

### **Step 3. Search your data**

On the "Success" screen, click **Start searching**. Splunk Web displays the data from the log file that you just uploaded, parsed into time-stamped events. If you do not see search results, verify that the time range displayed to the right of the search bar corresponds to the time range of the events in the file that you uploaded.



### **Step 4. (optional) Forward data**

To feed data continually to your Splunk Cloud Platform deployment, you can install and configure the Splunk **universal forwarder** on the machine where the data resides. For information on how to install and configure forwarders, see the following platform-specific documentation:

- [Get Windows Data into Splunk Cloud Platform](#)
- [Get \\*nix data into Splunk Cloud Platform](#)
- [Forward data from files and directories to Splunk Cloud Platform](#)

As with the data you uploaded, you can isolate your test data from any production data by forwarding it to a test index.

### **Next steps**

- Send data directly to your Splunk Cloud deployment using HTTP protocol. For details, see [Set up and use HTTP Event Collector](#).
- Create users and administer their access to your Splunk Cloud Platform deployment. For more information, see [Manage Splunk Cloud Platform users and roles](#).

## **Add a global banner to your Splunk Cloud Platform deployment**

Splunk Cloud Platform lets you display a global banner that remains visible to all users on all UI pages across the product. The global banner feature gives organizations with strict security concerns the ability to display a site classification message that is required to run Splunk software in some environments. For example, you can display a global banner that tells users they are using a secure or classified site.

While the primary use case for the global banner is the persistent display of a site classification message, you can also use it to display any type of notification that requires a persistent, highly visible message. For example, you can use the global banner to notify users about:

- New features
- Software version upgrades
- Scheduled maintenance or downtime
- Data outages

Splunk Web bulletin messages are suitable for most in-product notifications that do not require a persistent global message. For more information on bulletin messages, see [Customize Splunk Web messages](#).

Splunk Cloud Platform supports the display of a single global banner only. The global banner does not appear on the Splunk Cloud Platform login page and users cannot dismiss the banner inside the product.

## Customize the global banner

To customize the global banner a role must have the `edit_global_banner` capability. This capability is provided to the `sc_admin` role by default.

You can enable and customize the global banner using Splunk Web, as follows:

1. In Splunk Web, click **Settings > Server Settings > Global Banner**.
  2. Toggle the **Banner Visibility** switch to On.
  3. Select a background color for the global banner.
  4. Enter your message text.
  5. (optional) Specify an absolute URL to generate a hyperlink to additional information, such as your organization's best practices documentation.
  6. (optional) Enter the text of the hyperlink. For example, "Learn about best practices".
- Your customized global banner now appears on all UI pages in Splunk Cloud Platform.

## Customize global banner

[Server settings](#) » Customize global banner

The global banner allows admins to configure and communicate a single persistent banner message at the top of every Splunk Web page to all

Banner Visibility  On

Background Color

- Blue
- Green
- Yellow
- Orange
- Red

Message

Secure Site

Banner text is limited to one line, text is truncated afterward.

Hyperlink

[http://www.myorg.com/docs/best\\_practices](http://www.myorg.com/docs/best_practices)

Links must start with http:// or https://. Links are appended to the end of the message.

Hyperlink Text

Learn about best practices

Cancel

Save

## Determine your Splunk Cloud Platform Experience

Your Splunk Cloud Platform deployment has one of two possible Experience designations: Classic or Victoria. For information on the differing capabilities of each Splunk Cloud Platform Experience, see [Differences between Classic Experience and Victoria Experience](#).

To determine your Splunk Cloud Platform Experience:

1. In Splunk Web, click **Support & Services > About**.
2. In the About panel, under Splunk Cloud, find your Experience: Classic or Victoria.

### ***Differences between Classic Experience and Victoria Experience***

Classic Experience and Victoria Experience provide nearly identical capabilities, with the following exceptions:

- **Self-service app installation:** Splunk Cloud Platform deployments on Victoria Experience support self-service app installation of private apps and most public apps available on Splunkbase. For more information on self-service app installation, see [Install apps on your Splunk Cloud Platform deployment](#).

When you install an app using self-service app installation on Victoria Experience, the app is automatically installed on all search heads and search head cluster members across your deployment, including premium search heads running premium apps, such as Splunk IT Service Intelligence (ITSI) and Splunk Enterprise Security (ES).

- **Inputs Data Manager (IDM):** Splunk Cloud Platform deployments on Victoria Experience do not require an IDM. You can now run modular and scripted inputs directly on a Victoria Experience search head without the overhead of a separate IDM instance. Upon upgrade from Classic Experience to Victoria Experience, Splunk migrates all apps, customizations, and configs from the IDM to the primary search head or search head cluster, and removes the IDM from your deployment. Note that users with `sc_admin` and `apps` roles can see all configurations on the search tier, including the newly migrated apps from the IDM. For more information on IDM, see [Splunk Cloud Platform features](#).

If your IDM requires special IP allowlisting, you must move that allowlisting to the search head IP address, as all inputs now run on the search tier. If your search tier is a search head cluster, you must include all SHC member IP addresses in the allowlist.

- **Hybrid Search:** Splunk Cloud Platform deployments on Victoria Experience do not support hybrid search. If your deployment is on Victoria Experience, you must use federated search instead. For more information on federated search, see [Migrate from hybrid search to federated search in the Splunk Cloud Search Manual](#).
- **HTTP Event Collector (HEC):** Splunk Cloud Platform deployments on Victoria Experience support HEC token management using Splunk Web and Admin Configuration Service (ACS) endpoints. Deployments on Classic Experience support HEC token management using Splunk Web and Splunk Cloud Classic endpoints. For more information on using ACS and Splunk Cloud Classic endpoints for HEC token management, see [Manage HEC tokens in Splunk Cloud Platform](#).

For more information on the differences between Classic Experience and Victoria Experience, see Experience designations in the [Splunk Cloud Platform Service Description](#).

Splunk assigns your deployment to an appropriate Splunk Cloud Platform Experience for you. You cannot independently change your current Experience. For further details, contact your Splunk Cloud Platform representative.

## Change the UI theme of Splunk Cloud Platform

Splunk Cloud Platform UI allows you to choose dark theme or light theme, or keep the default system settings for the UI theme. UI theme is set on a per-user basis.

1. On the Splunk Home page, click **Search & Reporting** in the Apps Panel to open the Search app.
2. From the Splunk bar, navigate to **Administrator > Preferences > Theme**. Instead of **Administrator**, the Splunk bar might display your account name.
3. Select **Default System Theme**, **Light**, or **Dark**.
4. Select **Accept** and refresh the page to see your theme.

Default system theme is enabled by default in Splunk Cloud Platform 9.0.2208 and higher. This option uses the UI theme settings specified in your browser. If the browser allows, it uses light theme during the day and dark theme at night.

# Get Data Into Splunk Cloud Platform

## Introduction to Getting Data In

Splunk Cloud Platform administrators can add data to their Splunk Cloud Platform deployment using a variety of methods. This topic provides an overview of those methods.

## Fundamental Splunk and Splunk Cloud Platform concepts

Before attempting to get data into your Splunk Cloud Platform deployment, you should have a solid understanding of certain Splunk and Splunk Cloud Platform concepts. The table lists these concepts. You should also review the Splunk Cloud Platform information in the *Getting Data In* manual.

Concept	Description
deployment server	A deployment server is a Splunk Enterprise instance that acts as a centralized configuration manager for any number of forwarders, called "deployment clients". The deployment server is hosted on your premises or your Cloud environment (such as AWS or Azure). For a more detailed description of the components of a deployment server, see <a href="#">Deployment Server Architecture</a> .
indexes	The index is the repository for your data. When the Splunk platform indexes raw data, it transforms the data into searchable events. For more information about indexes, see <a href="#">Manage Splunk Cloud Platform Indexes</a> .
Inputs Data Manager	The Inputs Data Manager (IDM) is a component of your Splunk Cloud Platform environment optimized for data ingestion. It is intended for use with cloud data sources or when using add-ons that require inputs on the search tier. <b>Note:</b> Splunk Cloud Platform deployments on Victoria Experience do not require IDM. If your deployment is on Victoria Experience you can run add-ons that contain scripted and modular inputs directly on the search head. To determine if your deployment has the Classic or Victoria experience, see <a href="#">Determine your Splunk Cloud Platform Experience</a> .
search head search head cluster	For more information, see <a href="#">search head and search head cluster in the Splaxicon</a> .
source types	A source type is one of the critical default fields that Splunk software assigns to all incoming data. It tells Splunk software what kind of data you have, so that it can format the data intelligently during indexing. For more information, see <a href="#">Why source types matter</a> .
Splunk applications and add-ons	A Splunk app is an application that runs on the Splunk platform and typically addresses several use cases. Add-ons support and extend the functionality of the Splunk platform and the apps that run on it, usually by providing inputs for a specific technology or vendor. For more information about add-ons, see <a href="#">About Splunk add-ons</a> .
universal forwarder	To forward data to Splunk Cloud Platform, you typically use the Splunk universal forwarder. The universal forwarder is a dedicated, streamlined version of Splunk Enterprise that contains only the essential components needed to forward data. The universal forwarder does not support Python and does not expose a UI. In most situations, the universal forwarder is the best way to forward data to indexes. Its main limitation is that it forwards unparsed data, except in certain cases, such as structured data. For more information, see <a href="#">Work with forwarders</a> .

## Types of data that Splunk Cloud Platform accepts

Splunk Cloud accepts a wide variety of data, and can also monitor relational databases and third-party infrastructures. For more information, see the following sections in the *Getting Data In* manual:

- [What data can I index?](#)
- [Monitor files and directories](#)
- [Get data from TCP and UDP ports](#)



- Monitor Windows data with the Splunk platform
- Share HEC Data
- Monitor First In, First Out (FIFO) queues

See also the *Data Manager for Splunk Cloud Platform User Manual*.

## Tools to get data into Splunk Cloud Platform

This section is designed to help you make decisions about the best way to get data into your Splunk Cloud Platform instance. There are a few different ways to get data into Splunk Cloud Platform: forwarders, HTTP Event Collector (HEC), apps and add-ons, or the Inputs Data Manager (IDM). The best way to get data in depends on the source of the data and what you intend to do with it.

### **Work with forwarders**

Usually, to get data from your customer site to Splunk Cloud Platform, you use a **forwarder**.

A forwarder is a version of Splunk Enterprise optimized to send data. A universal forwarder is a purpose-built data collection mechanism with very minimal resource requirements, whereas a heavy forwarder is full Splunk Enterprise deployment configured to act as a forwarder with indexing disabled.

Splunk forwarders send data from a datasource to your Splunk Cloud Platform deployment for indexing, which makes the data searchable. Forwarders are lightweight processes, so they can usually run on the machines where the data originates. For more information, see the following topics:

- [Forward data from files and directories to Splunk Cloud Platform](#)
- [Upgrade your Forwarders](#)
- Use forwarders to get data into Splunk Cloud Platform in the *Getting Data In* manual
- The *Forwarding Data* manual

### **Work with HTTP Event Collector**

The HTTP Event Collector (HEC) uses a token-based authentication model so you can securely send data and application events to a Splunk deployment over the HTTP and Secure HTTP (HTTPS) protocols.

For more information, see the following sections in Set up and use HTTP Event Collector in Splunk Web in the *Getting Data In* manual:

- HEC and Splunk Cloud Platform
- Configure HTTP Event Collector on Splunk Cloud Platform
- For general and specific information on sending data: Send data to HTTP Event Collector and Send data to HTTP Event Collector on Splunk Cloud Platform

### **Work with Apps and Add-ons**

Splunk **apps** and **add-ons** extend the capability and simplify the process of getting data into your Splunk platform deployment.

For more information, see Use apps and add-ons to get data in in the *Getting Data In* manual.

## Splunk Cloud Platform considerations

Apps and add-ons that contain a data collection component should be installed on forwarders, IDMs, or your Splunk Cloud Platform instance search head for their data collection functions (modular or scripted inputs). For more information, see [Work with forwarders](#) and [Work with Inputs Data Manager \(IDM\)](#).

If your Splunk Cloud Platform deployment is on Classic Experience, you must install apps and add-ons that contain modular or scripted inputs on a separate IDM instance. If your deployment is on Victoria Experience, you can install apps or add-ons that contain modular or scripted inputs directly on your Splunk Cloud Platform instance search head. To determine if your deployment is on Classic Experience or Victoria Experience, see [Determine your Splunk Cloud Platform Experience](#).

Regardless of your Splunk Cloud Platform Experience designation, you can deploy apps and add-ons that perform data collection functions, including those that contain modular or scripted inputs, to forwarders hosted locally in your environment.

Splunk Cloud Platform supports self-service installation of both public apps and add-ons from Splunkbase and private apps that you can create based on your unique requirements. For more information on public apps, see [Install apps on your Splunk Cloud Platform deployment](#). For more information on private apps, see [Manage private apps on your Splunk Cloud Platform deployment](#).

### *Work with Inputs Data Manager*

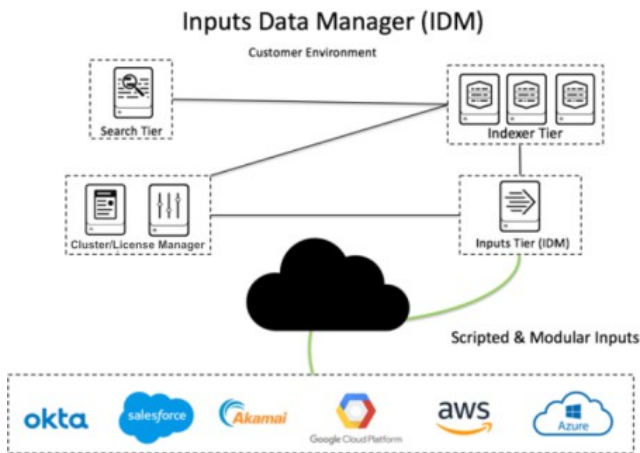
The **Inputs Data Manager** (IDM) is a hosted solution for Splunk Cloud Platform that supports scripted and modular inputs and cloud-based inputs that you want to send directly to Splunk Cloud Platform. In most cases, an IDM eliminates the need for customer-managed infrastructure.

IDM is required to run modular and scripted inputs on Splunk Cloud Platform deployments on Classic Experience only. If your deployment is on Victoria Experience, you can run modular and scripted inputs directly on the search head. To determine if your deployment is on Classic Experience or Victoria Experience, see [Determine your Splunk Cloud Platform Experience](#).

As a best practice, use an IDM in the following cases:

- You have scripted or modular inputs that you want to send to Splunk Cloud Platform. For example, you can poll a cloud-based database, web service, or API for specific data and process the results.
- You have cloud-based inputs such as Microsoft Azure or AWS that you want to send directly to Splunk Cloud Platform without the intermediary step of sending data to an on-premise forwarder. You can send these inputs directly to an IDM rather than routing them through a forwarder to get the data into Splunk Cloud Platform.

The following graphic shows the typical architecture of IDM. Note that the search tier and index tier are not hosted on the IDM. The IDM is not intended to store data or perform searches.



IDM is not supported on the Splunk Cloud Platform Free Trial.

### Ports opened for IDM

The following port access applies to inbound and outbound IDM ports:

- Inbound access to ports 443 and 8089 are controlled by an access list. Contact Support if you need to modify the access list.
- Outbound access to port 443 is open by default. Contact Support if you need to open additional outbound ports.

When you contact Support, provide a list of public IP addresses and subnets with this request. For example, you might want to open port 8089, the port for the REST API. Note that opening a specific outbound port opens the same port for all tiers in your Splunk Cloud environment.

### Apps supported on IDM

If the app contains modular inputs and is Splunk Cloud Platform certified, it is compatible with Splunk Cloud Platform IDM. Generally, apps that are cloud-based are well-suited to IDM. Many cloud-based apps are supported.

To verify if your app is supported on IDM, check Splunkbase.

### Limitations when working with IDM

The IDM is intended to function specifically as a forwarder for modular and scripted inputs, or to obviate the need to route cloud-based inputs through an on-premise forwarder. The following functions are not intended to be performed on the IDM:

- Search capabilities are capped for users on IDM. The IDM is not intended to function as a search head.
- IDM does not currently support Self-Service App Installations. To get modular and scripted input onto the IDM, you need to create a private app and request that Support upload it.
- If an add-on is packaged with or related to an Enterprise Security search head, do not use IDM. Instead, run the add-on on the Enterprise Security search head.
- HEC inputs are not supported with IDM.
- IDM isn't a syslog sink, nor can it receive unencrypted TCP streams.

- IDM isn't a one-to-one replacement for a heavy forwarder. You must still use a heavy forwarder if you need to perform parsing or activities other than standard scripted and modular data inputs.

**Use IDM with scripted and modular inputs**

To use scripted or modular inputs, you must package them in a private app. To do this, complete the following high-level steps:

1. Create your modular or scripted inputs. For instructions on creating these inputs, see *Get data from APIs and other remote data interfaces through scripted inputs* in the *Getting Data In* manual.
2. Package the script or modular input in a private app. For instructions on building a private app for Splunk Cloud Platform, see *Overview of developing a private Splunk Cloud Platform app*.
3. Submit the private app for Splunk Cloud Platform vetting.
4. Request that Support upload the app to your IDM.

**Use IDM with cloud-based add-ons**

When you work with IDM and Cloud-based add-ons, complete the following high-level steps to get data in:

1. Create a support request to install the Add-on.
2. Configure an index on your Splunk Cloud Platform instance. This index is going to be associated with your cloud input.
3. Perform any configurations needed on the cloud-based source that enables you to get data in.
4. Configure the Splunk Add-on on your Inputs Data Manager (IDM).
5. You will also need to configure inputs on the IDM. The IDM is responsible for data ingestion.
6. Verify that data is flowing to your Splunk Cloud Platform environment.

As a best practice, install cloud-based add-ons on an IDM, and install on-premises-based add-ons on a universal forwarder or heavy forwarder.

**See also**

For more information about	See
Getting AWS data in using IDM	<a href="#">Get Amazon Web Services (AWS) data into Splunk Cloud Platform</a>
Getting Microsoft Azure data in using IDM	<a href="#">Get Microsoft Azure data into Splunk Cloud Platform</a>

## Get Amazon Web Services (AWS) data into Splunk Cloud Platform

This topic guides you through the steps to use the Splunk Add-on for AWS to get AWS data into Splunk Cloud Platform.

The best practice in most use cases for onboarding this data type is to now use the Data Manager app. Data Manager comes installed on all Splunk Cloud instances. For more information on which method best fits your use cases, see: [Use cases for the AWS data input in Data Manager](#) [Use cases for the Splunk Add-on for AWS](#)

### Administrator requirements

Splunk Cloud Platform administrators must meet the following prerequisites to get AWS data into Splunk Cloud Platform:

- You need a valid AWS account with administrative permissions to configure the AWS services that provide your data.
- You need permission to create IAM roles and users. This lets you set up accounts or EC2 IAM roles with the ability to collect data from your AWS services.

If you do not have permissions to perform all the actions yourself, work closely with your organization's AWS administrator to complete all steps, including creating the accounts or EC2 IAM roles with the permissions that the Splunk Add-on for AWS uses to connect. If you need to review the AWS documentation, go to <https://docs.aws.amazon.com/index.html>

Customers are responsible for the setup, configuration, and maintenance of third-party services and resources, which includes payment. See Network connectivity and data transfer in the Splunk Cloud Platform Service Description.

## AWS region limitations

The Splunk Add-on for AWS supports all regions offered by AWS.

In the AWS China region and the AWS GovCloud regions, the add-on only supports the services that AWS supports in those regions. There are limitations in both of these regions, so ensure that you understand the limitations before you begin. For more information, see the following topics:

- For an up-to-date list of what products and services are supported in this region, see [amazonaws.cn/en/products/](http://amazonaws.cn/en/products/) or [aws.amazon.com/about-aws/global-infrastructure/regional-product-services/](http://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/).
- For an up-to-date list of what services and endpoints are supported in the GovCloud region, see the AWS documentation: [docs.aws.amazon.com/govcloud-us/latest/UserGuide/using-services.html](https://docs.aws.amazon.com/govcloud-us/latest/UserGuide/using-services.html).

## Before you begin

To get AWS data into Splunk Cloud Platform, you need a solid understanding of Splunk and AWS concepts. The table lists these concepts and provides links to more information.

Product	Concept	See
Splunk and Splunk Cloud	<b>indexes</b> <b>search head</b> <b>search head cluster</b> <b>source types</b> <b>Inputs Data Manager</b> <b>Splunk apps and add-ons</b>	<a href="#">Fundamental Splunk and Splunk Platform Cloud concepts</a>
AWS	<b>AWS CloudTrail</b>	This is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command-line tools, and other AWS services.

## Overview of getting your AWS CloudTrail data into Splunk Cloud Platform

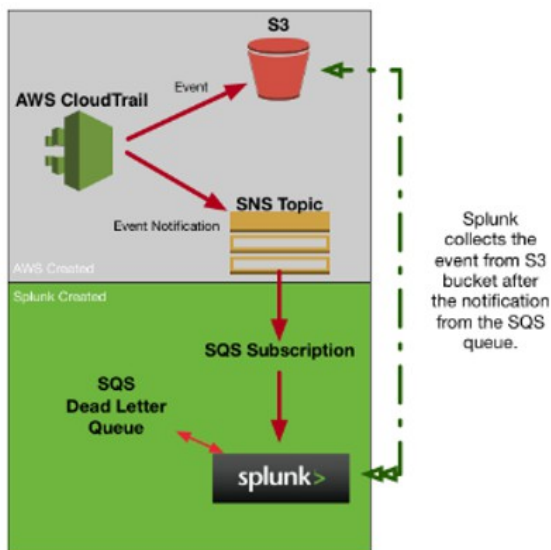
In this procedure, you'll set up your Splunk Cloud Platform instance to get data from AWS CloudTrail. There are many useful AWS services that you can configure later, but starting with CloudTrail provides a very comprehensive view of AWS activity.

Before performing this procedure, check if your Splunk Cloud Platform environment uses the Victoria Experience or the Classic Experience; see Determine your Splunk Cloud Platform Experience. If your Splunk Cloud Platform environment uses the Victoria Experience, perform steps 8-9 on your Splunk Cloud Platform search head or search head cluster member instance. If your Splunk Cloud Platform environment uses the Classic Experience, perform steps 8-9 on your Splunk Cloud Platform Inputs Data Manager (IDM) instance.

To get AWS CloudTrail data into Splunk Cloud Platform, complete the following high-level steps:

1. [Set up your Splunk Cloud Platform environment.](#)
2. [Configure an access policy for Splunk Access in AWS.](#)
3. [Create a Splunk Access user.](#)
4. [Create a group for Splunk Access Users.](#)
5. [Enable the AWS CloudTrail Service.](#)
6. [Create an SQS subscription.](#)
7. [Create an SQS subscription for the Dead Letter Queue.](#)
8. [Configure the Splunk Add-on for AWS.](#)
9. [Configure CloudTrail inputs.](#)

The following graphic shows the configuration of AWS and Splunk Cloud Platform that enables you to get AWS data into Splunk Cloud Platform:



When you finish the configuration steps, AWS CloudTrail populates the following Splunk App for AWS dashboards:

- Overview
- Topology
- Security Overview
- IAM Activity
- VPC Activity
- Security Groups
- Key Pairs Activity
- Network ACLs
- User Activity
- Insights Overview
- Security Anomaly Insights
- Timeline

## Step 1: Set up your Splunk Cloud Platform environment

Before you can get AWS data into Splunk Cloud Platform, you must ensure the following:

- Confirm that you are assigned the `sc_admin` role on your Splunk Cloud Platform instance.
- Install the following, and ensure you allow adequate time for these tasks to be completed before you attempt to get data in:
  - ◆ For Classic Experience customers, request that Splunk Support install the Splunk Add-on for AWS on your Inputs Data Manager and the Splunk App for AWS on your Splunk Cloud Platform instance.
  - ◆ Victoria Experience customers can use the self-service app install procedure described in [Install an app to install the Splunk Add-on for AWS on your search head or search head cluster member instance, and the Splunk App for AWS on your Splunk Cloud Platform instance.](#)
- Verify that port 443 is available. The Splunk Add-on for AWS makes REST API calls using HTTPS on this port.
- Create a test index in your Splunk Cloud Platform instance so that you can test your installation before going into production. Follow these instructions to create an index: [Create a Splunk Cloud Platform Index.](#)

## Step 2: Configure an access policy for Splunk Access in AWS

Complete the following steps to set up an access account in AWS so that Splunk Cloud Platform can communicate with AWS without having root access. The `SplunkAccess` account is used to collect data from AWS.

1. In AWS, open **IAM > Policies > Create New Policy**.
2. Click **Create New Policy**.
3. Select JSON.
4. In the JSON visual editor, paste the following JSON code:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sqs:GetQueueAttributes",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:GetQueueUrl",
        "sqs:SendMessage",
        "sqs:DeleteMessage",

```

```

"s3:ListBucket",
"s3:GetObject",
"s3:GetBucketLocation",
"s3:ListAllMyBuckets",
"s3:GetBucketTagging",
"s3:GetAccelerateConfiguration",
"s3:GetBucketLogging",
"s3:GetLifecycleConfiguration",
"s3:GetBucketCORS",
"config:DeliverConfigSnapshot",
"config:DescribeConfigRules",
"config:DescribeConfigRuleEvaluationStatus",
"config:GetComplianceDetailsByConfigRule",
"config:GetComplianceSummaryByConfigRule",
"iam:GetUser",
"iam:ListUsers",
"iam:GetAccountPasswordPolicy",
"iam:ListAccessKeys",
"iam:GetAccessKeyLastUsed",
"autoscaling:Describe*",
"cloudwatch:Describe*",
"cloudwatch:Get*",
"cloudwatch:List*",
"sns:Get*",
"sns:List*",
"sns:Publish",
"logs:DescribeLogGroups",
"logs:DescribeLogStreams",
"logs:GetLogEvents",
"ec2:DescribeInstances",
"ec2:DescribeReservedInstances",
"ec2:DescribeSnapshots",
"ec2:DescribeRegions",
"ec2:DescribeKeyPairs",
"ec2:DescribeNetworkAcls",
"ec2:DescribeSecurityGroups",
"ec2:DescribeSubnets",
"ec2:DescribeVolumes",
"ec2:DescribeVpcs",
"ec2:DescribeImages",
"ec2:DescribeAddresses",
"lambda:ListFunctions",
"rds:DescribeDBInstances",
"cloudfront:ListDistributions",
"elasticloadbalancing:DescribeLoadBalancers",
"elasticloadbalancing:DescribeInstanceHealth",
"elasticloadbalancing:DescribeTags",
"elasticloadbalancing:DescribeTargetGroups",
"elasticloadbalancing:DescribeTargetHealth",
"elasticloadbalancing:DescribeListeners",
"inspector:Describe*",
"inspector:List*",
"kinesis:Get*",
"kinesis:DescribeStream",
"kinesis:ListStreams",
"kms:Decrypt",
"sts:AssumeRole"
],
"Resource": [
  "*"
]
}

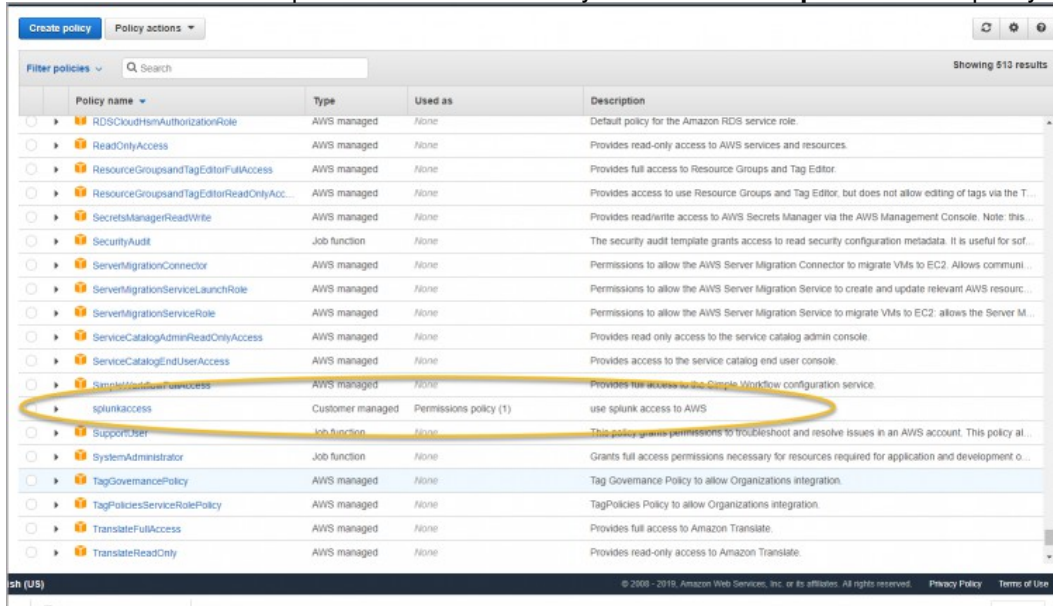
```



- ```

]
}

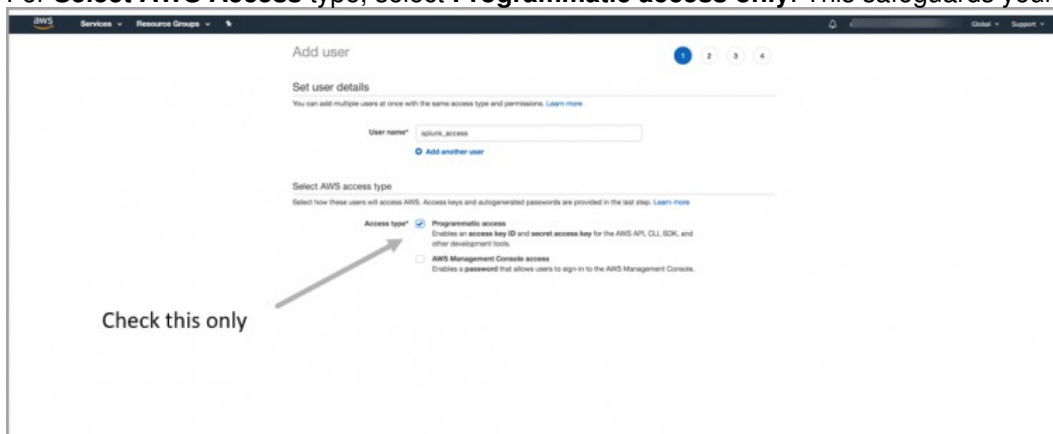
```
5. Click Review Policy.
  6. Name the policy **SplunkAccess**.
  7. Give it a description.
  8. Click **Create Policy**.
  9. Return to the list of IAM policies and ensure that you see the new **SplunkAccess** policy in the list of policies.



### Step 3: Create a Splunk Access user

Complete the following steps to create the Splunk Access user.

1. In AWS, from the IAM Users list, click **Users**.
2. Click **Add User**.
3. Add a user name. In this example, use the name **Splunk\_Access**.
4. For **Select AWS Access type**, select **Programmatic access only**. This safeguards your account.

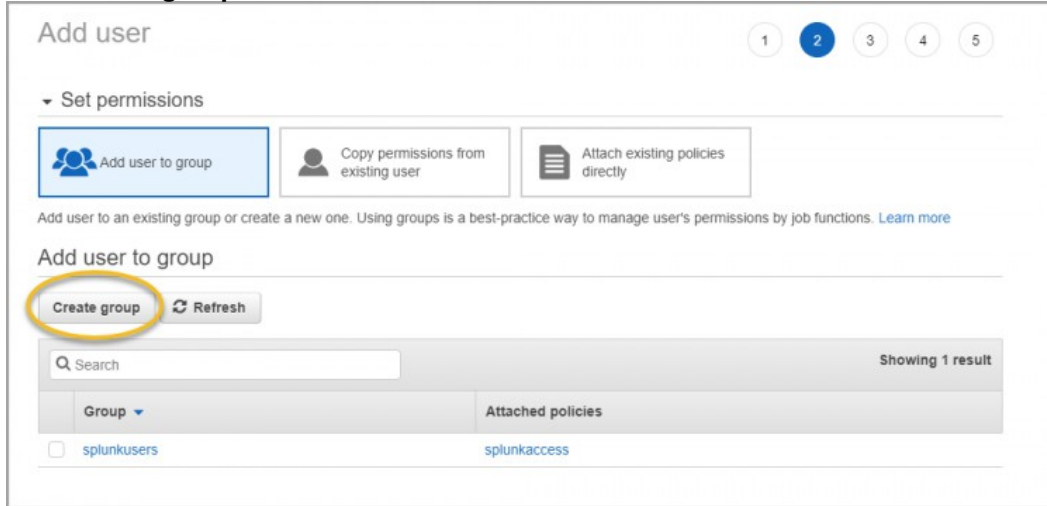


5. Click **Next: Permissions** to add the new user to a group.

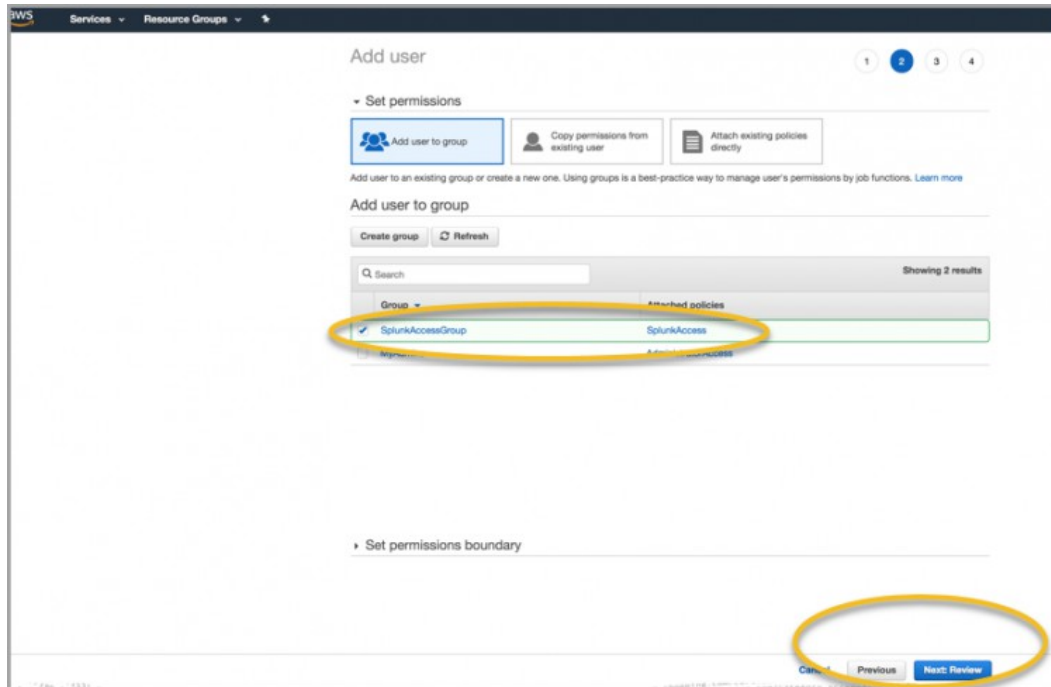
## Step 4: Create a group for Splunk Access users

Complete the following steps to assign the user to a group.

1. Click **Create group**:



2. In **Group Name**, enter **SplunkAccessGroup**.
3. Apply the IAM policy you created. To do this, search for the IAM policy by entering **splunk** in the **Filter Policies** field. The IAM access policy you created displays.
4. Select the checkbox next to the SplunkAccess IAM policy, and click **Create Group**.
5. Ensure your new group is selected and click **Next Review**:



6. Now that you have created a user and associated that user with a group and an IAM policy, click **Create User**.
7. After the user is created, Amazon provides access to a CSV file with the user credentials. Download this file and store it securely because you will need it later.

## Step 5: Enable the AWS CloudTrail Service

Complete the following steps to enable CloudTrail to capture AWS data and send it to Splunk Cloud Platform. AWS CloudTrail writes events to a Simple Notification Service (SNS) topic, and you can then create a Simple Queue Service (SQS) subscription. When SQS notifies Splunk Cloud Platform of an event, Splunk Cloud can collect the events from the S3 bucket.

1. In AWS Management Console, click **All Services**, and select **CloudTrail**.
2. Click **Get Started Now**.
3. Name your trail. For this example, use the name *'cloudtrail'*.
4. Apply the trail to all regions.
5. In the **Read/Write events** field, select **All**.
6. Create a storage location for CloudTrail. In this case, create a new S3 bucket. To do this, select **Create a new S3 bucket**. As a best practice, use the naming convention of *cloudtrail-<AWSAccountID>*.
7. Click **Advanced**.
8. In the **Send SNS notification for every log file delivery** field, click **Yes**.
9. In the **Create a new SNS topic** field, click **Yes**.
10. Enter a name for your SNS topic. For this example, name the topic **cloudtrail**.

## Step 6: Create an SQS subscription

Complete the following steps to set up an SQS to subscribe to the SNS topic created by AWS CloudTrail.

1. From the **AWS Management Console > Application Integration**, select **Simple Queue Service**.

2. Click **Create New Queue**.
3. Enter a name for the queue. For example, call it **cloudtrail**.
4. Select **Standard Queue**.
5. Click **Quick-Create Queue**. Your new queue is displayed.
6. Select your queue and select **Queue actions**.
7. Select **Subscribe Queue to SNS topic**.
8. In the **Choose a Topic** field, select the SNS topic you created in [step 5](#):

9. The Topic ARN auto-populates.
10. Select **Subscribe**.

## Step 7: Create an SQS subscription for the Dead Letter Queue

Complete the following steps to create a CloudTrail Dead Letter Queue.

This queue is required for the new Splunk SQS-based S3 input. To keep the inputs stateless, the Dead Letter Queue notifies Splunk Cloud Platform where the last input left off and where to continue collecting events from AWS.

1. From the **AWS Management Console > Application Integration**, select **Simple Queue Service**.
2. Click **Create New Queue**.
3. Enter a name for the queue. For this example, call it **cloudtrail-dlq**.
4. Select **Standard Queue**.
5. Click **Quick-Create Queue**. Your new queue is displayed.
6. From the list of queues, select your cloudtrail queue, and click **Queue Actions**.
7. Click **Configure Queue**.
8. In the Dead Letter Queue Settings, select **Use Redrive Policy**.
9. In the Dead Letter Queue field, enter **cloudtrail-dlq**.
10. In the Maximum receives field, enter **3**.
11. Click **Save Changes**.

## Step 8: Configure the Splunk Add-on for AWS

Complete the following steps to configure the Splunk Add-on for AWS to get data from your AWS account.

1. Open the Splunk Add-on for AWS from the list of available apps.  
 Victoria Experience customers: Open the app from your search head or search head cluster member instance.  
 Classic Experience customers: Open the app from your IDM. You can log into your IDM at

`https://idm-<cloudname>.splunkcloud.com` where `<cloudname>` represents your Splunk Cloud Platform name.

2. Click **Configuration > Create an account > Add**.
3. In the **Add Account** field, enter a name for the account. For this example, name it **splunk\_access**.
4. Open the credentials file you previously downloaded. Enter the credentials in the **Key ID** and **Secret Key** fields.
5. Leave the Region Category field as *Global* unless you are using GovCloud or AWS China.
6. Click **Add**.

## Step 9: Configure Cloudtrail inputs

Complete the following steps to configure Cloudtrail inputs for your specific Splunk Cloud Platform Experience, as follows:

- Victoria Experience customers: Perform these steps on your search head or search head cluster member instance.
- Classic Experience customers: Perform these steps on your IDM.

1. In the Inputs tab, click **Create New Input > Cloudtrail > SQS-Based S3**.
2. Enter a name for the CloudTrail input. For example, name it **cloudtrail**.
3. From the drop-down list, select your AWS account.
4. Select your AWS region.
5. In the SQS queue field, select your cloudtrail queue.
6. Use the default values for other settings, and click **Save**.

Now you can open your Splunk App for AWS dashboards and see that the dashboards for Cloudtrail are getting populated with data.

## What's next?

Now that you have configured the Splunk AWS Add-on for your Splunk Cloud Platform instance, you can set up your Splunk Cloud Platform instance to get data from various AWS services. The table lists these AWS services, and also additional topics of interest.

### See also

| For more information about                                                                                                                                                                                                                                                                            | See                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| <b>Amazon CloudWatch:</b> A monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources.                          | Splunk Add-on for AWS CloudWatch documentation                                                        |
| <b>AWS Config:</b> A service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. | Splunk Add-on for AWS Config documentation                                                            |
| <b>AWS Inspector:</b> An automated security assessment scanner that can evaluate security loopholes and deviation from the best practices for applications hosted on AWS. AWS Inspector communicates with EC2 instances using agents installed on it and generates reports.                           | Splunk Add-on for AWS Inspector documentation                                                         |
| Testing and troubleshooting data input                                                                                                                                                                                                                                                                | The Improve the data input process section in the Splunk Cloud Platform <i>Getting Data In</i> manual |

# Get Microsoft Azure data into Splunk Cloud Platform

This topic guides you through the steps to get Microsoft Azure data into Splunk Cloud Platform.

## Administrator requirements

Splunk Cloud Platform administrators must meet the following prerequisites to get Microsoft Azure data into Splunk Cloud Platform:

- Permissions necessary to make changes in your Microsoft Azure environment.
- Activity logs and subscriptions in your Microsoft Azure environment.

If you don't know this information or have these permissions, work closely with your organization's Microsoft Azure administrator to complete these steps.

Customers are responsible for the setup, configuration, and maintenance of third-party services and resources, which includes payment. See Network connectivity and data transfer in the Splunk Cloud Platform Service Description.

## Before you begin

To get Microsoft Azure data into Splunk Cloud Platform, you need a solid understanding of Splunk concepts. The table lists these concepts and provides links to more information.

| Product                          | Concept                                                                                                                                                   | See                                                                   |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| Splunk and Splunk Cloud Platform | <b>indexes</b><br><b>search head</b><br><b>search head cluster</b><br><b>source types</b><br><b>Inputs Data Manager</b><br><b>Splunk apps and add-ons</b> | <a href="#">Fundamental Splunk and Splunk Cloud Platform concepts</a> |

## Overview

In this procedure, you'll get your Microsoft Azure activity data into your Splunk Cloud Platform instance. The activity logs contain information on events and users' actions and when those actions occurred.

You'll first create an application registration, which is similar to a service account that you can use to authenticate to Microsoft Azure. The application registration has an application ID (similar to a user name) and an application key or secret (which is similar to a password). This allows your Splunk Cloud Platform instance to authenticate to Microsoft Azure and get the activity log data in. You'll then configure the Splunk Add-on for Microsoft Cloud Services on your search head, search head cluster member instance, or the Inputs Data Manager and on Splunk Cloud Platform to make it easy to get the data into Splunk Cloud Platform.

There are many other types of data you may want to get into your Splunk Cloud Platform instance, and this topic is not intended to be a comprehensive guide for getting all your Microsoft Azure data into Splunk Cloud Platform. Instead, use this process as a template that you can repeat until you've included all of the relevant source types.

Before performing this procedure, check if your Splunk Cloud Platform environment uses the Victoria Experience or the Classic Experience; see [Determine your Splunk Cloud Platform Experience](#). If your Splunk Cloud Platform environment uses the Victoria Experience, perform step 4 on your Splunk Cloud Platform search head or search head cluster member instance. If your Splunk Cloud Platform environment uses the Classic Experience, perform step 4 on your Splunk Cloud Platform Inputs Data Manager (IDM) instance.

To get Microsoft Azure data into Splunk Cloud Platform, complete the following high-level steps:

1. [Set up your Splunk Cloud Platform environment](#).
2. [Configure an index on your Splunk Cloud Platform instance](#).  
You create an index for the Microsoft Azure data you want to bring into your Splunk Cloud Platform deployment.
3. [Configure Microsoft Azure to authenticate and ingest data from Splunk Cloud Platform](#).  
While you are configuring Microsoft Azure, you'll also need to record some information that you'll use to connect Splunk Cloud Platform with Microsoft Azure.
4. [Configure the Splunk Add-on for Microsoft Cloud Services](#).  
To ensure data ingestion, you will also need to configure the add-on and the inputs. If you're a Victoria Experience customer, you'll do this configuration on the search head or search head cluster member instance. If you're a Classic Experience customer, you'll do this configuration on the Inputs Data Manager (IDM).
5. [Confirm data is flowing to your Splunk Cloud Platform instance](#).  
After configuring Microsoft Azure settings and add-on settings, verify that data is flowing to your Splunk Cloud Platform instance.

## Step 1: Set up your Splunk Cloud Platform environment

Before you can get Microsoft Azure data into your Splunk Cloud Platform instance, you must ensure the following:

- Confirm that you are assigned the `sc_admin` role on your Splunk Cloud Platform instance.
- Install the following, and ensure you allow adequate time for these tasks to be completed before you attempt to get data in:
  - ◆ Classic Experience customers must request that Splunk Support install the Splunk Add-on for Microsoft Cloud Services on your Inputs Data Manager and your Splunk Cloud Platform instance.
  - ◆ Victoria Experience customers can use the self-service app install procedure described in [Install an app to install Splunk Add-on for Microsoft Cloud Services to your search head or search head cluster member instance](#), and on your Splunk Cloud Platform instance.
- Create a test index in your Splunk Cloud instance so that you can test your installation before going into production. Follow these instructions to create an index: [Manage Splunk Cloud Platform Indexes](#).

## Step 2: Configure a new index on your Splunk Cloud Platform instance

Complete the following steps to create indexes to store the events sent from your Microsoft Azure instance.

It's a best practice to create separate indexes for different types of data. For this initial index, you'll create an index to store Microsoft Azure activity data.

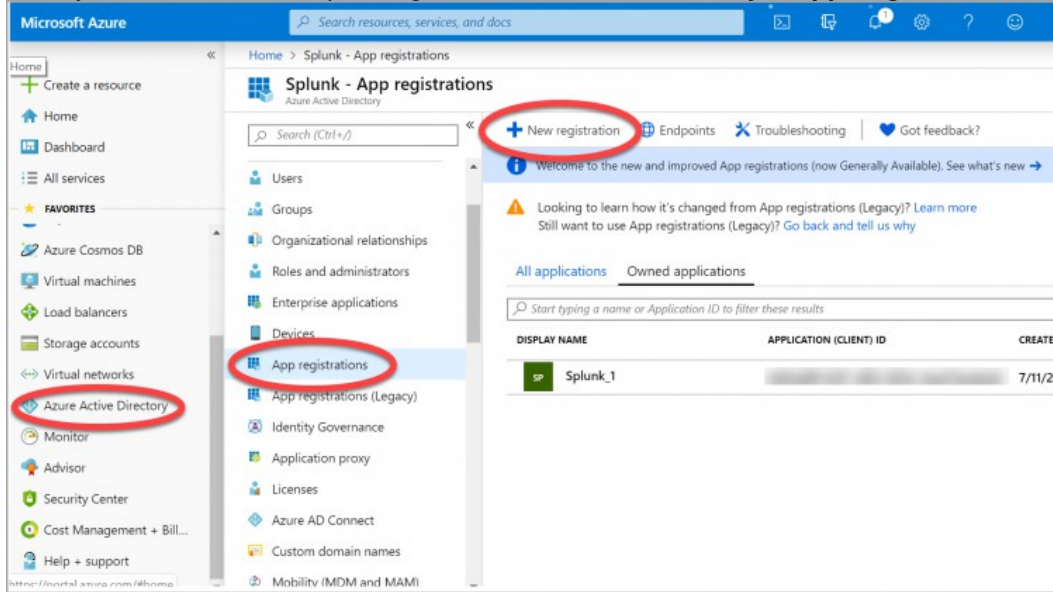
1. From your Splunk Cloud Platform instance, go to **Settings > Indexes**.
2. Click **New Index**.
3. In the Index name field, as an example, enter *azure-activity*. Alternatively, you can select a name that is consistent with your company's index naming convention.
4. For **Index Data Type**, select **Events**.
5. For **Searchable time (days)**, enter the number of days you want data to be searchable. As an example, enter *30*.
6. Click **No Additional Storage**, and click **Save**.



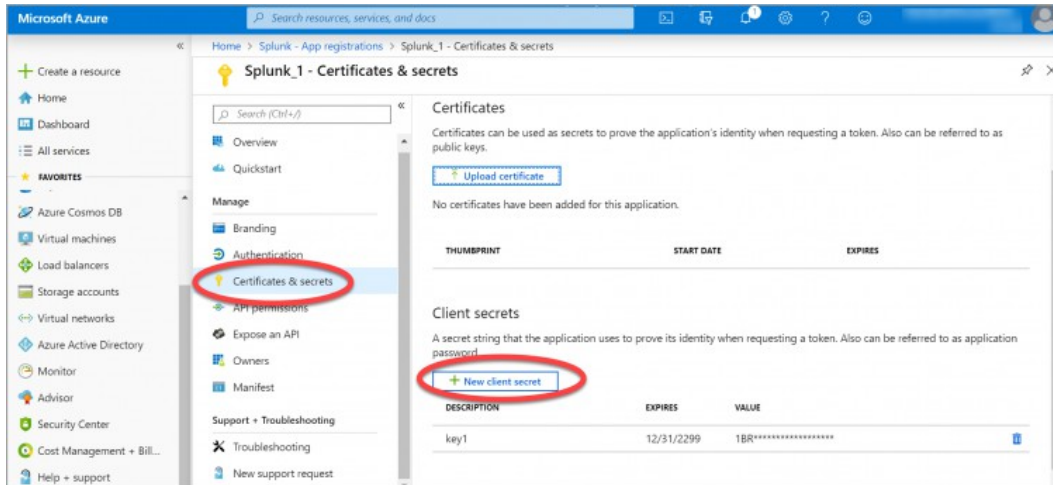
### Step 3: Configure Microsoft Azure to authenticate and ingest activity data

Complete the following steps to configure an Application Registration on Microsoft Azure and give it read access to resources in your subscription.

1. From your Microsoft Azure portal, go to **Azure Active Directory > App Registrations > New Registration**:

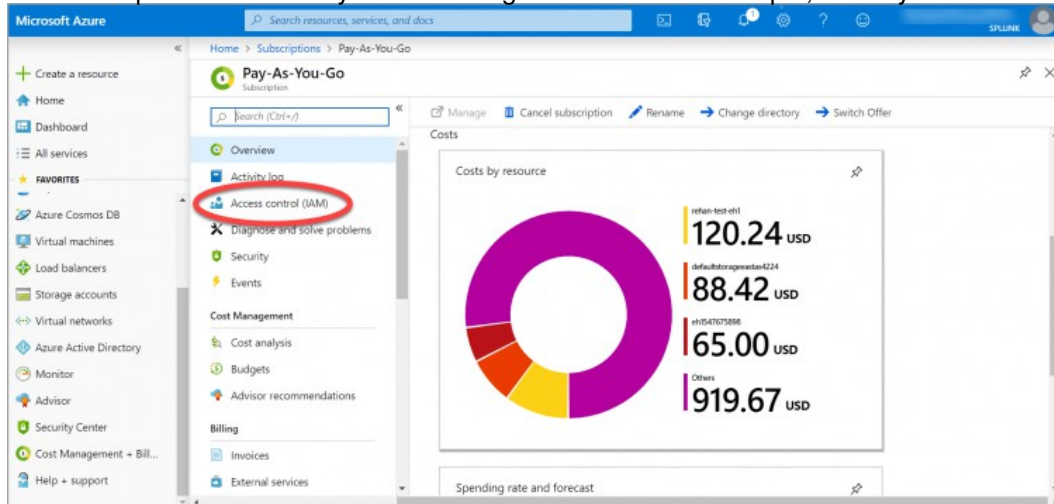


2. In the **New Registration** field, enter a name for the Application Registration.
3. Leave the **Application type** at the default value (Web app/API).
4. Leave the **Supported account type** as the default value.
5. Click the **Register** button.
6. In a separate location, note the application ID value. You can think of the application ID as a user ID. The application ID value maps to the Client ID field when you configure Microsoft Azure in Splunk Cloud Platform via the Microsoft Cloud Services add-on.
7. Create an application secret. This is comparable to a password or key. To do this, go to **Certificates & secrets**, and click **New client secret**:

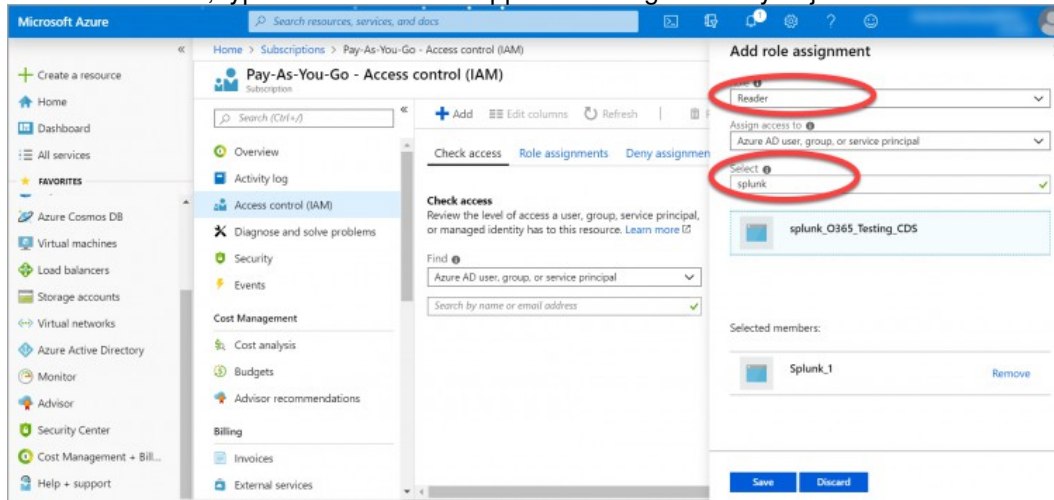




8. Enter a name for the secret in the **Description** field, and select the **Never** radio button under the **Expires** field. Click **Add**.
9. In a separate location, copy the value for the secret key from the clipboard as this is the only time it will display. You will need this value later when you configure the add-on to connect to Microsoft Azure. The secret key maps to the **Key (Client Secret)** value when you configure Microsoft Azure in Splunk via the Microsoft Cloud Services add-on. Note that you can create a new secret if you lose this value.
10. Grant the application registration read access to resources in your subscription. Click **Subscriptions**, and choose the subscription from which you want to ingest data. In this example, the Pay-As-You-Go subscription is selected:



11. Click **Access control (IAM) > Add > Add role assignment**.
12. From the Role dropdown menu, select **Reader**.
13. In the **Assign Access to** field, leave the default value.
14. In the select field, type the name of the Application Registration you just created:



15. Click **Save**
16. If you have multiple subscriptions, you can continue to add access to each of the subscriptions you want to include. When you have added all of the subscriptions that you want to include, your Microsoft Azure configuration should be complete.

## Step 4: Configure the Splunk Add-on for Microsoft Cloud Services

Complete the following steps to configure the Splunk Add-on for Microsoft Cloud Services and also the necessary inputs.

1. Open the Splunk Add-on for Microsoft Cloud Services from the list of available apps.  
Victoria Experience customers: Open the app from your search head or search head cluster member instance.  
Classic Experience customers: Open the app from your IDM. Log into your IDM instance at <https://idm-<cloudname>.splunkcloud.com> where <cloudname> represents your Splunk Cloud Platform name.
2. Go to **Apps > Splunk Add-on for Microsoft Cloud Services**.
3. From the Add-on, select **Configuration > Add Azure App Account**.
4. In the **Name** field, enter a name for the Microsoft Azure App account.
5. In the **Client ID** field, enter the application ID value that you saved earlier.
6. In the **Key (Client Secret)** field, enter the value for the client secret you saved earlier.
7. In the **Tenant ID** field, enter the Microsoft Azure Directory ID.  
You can find the Directory ID in the Azure Portal, by going to **Azure Active Directory > Properties**.
8. Click **Add**.
9. Click the **Inputs** tab to configure the inputs.
10. Click **Create New Input > Azure Audit**.
11. Enter a name for the input.
12. From the dropdown menu, select the account you created earlier.
13. In the **Subscription ID** field, enter the Azure Subscription ID. You can find this in the Azure Portal by going to **All services > Subscriptions**.  
Copy the value of the Subscription ID field.
14. Leave the start time value as the default value and the interval at 3600.
15. For the index, select the *azure-activity* index you created earlier.
16. Click **Add**.

## Step 5: Confirm data is flowing to your Splunk Cloud Platform instance

Complete the following steps to see if data is flowing to your Splunk Cloud Platform instance.

1. From your Splunk Cloud Platform instance, go to **Apps > Search and Reporting**.
2. In the search field, enter *index= azure-activity*
3. For the time range, select **Presets > Last 30 days**.
4. Click the search icon. Events from your Microsoft Azure environment should display.

## What's next?

The table lists additional topics of interest in the Splunk Cloud Platform *Getting Data In* manual.

### See also

| For more information about                     | See                              |
|------------------------------------------------|----------------------------------|
| Getting, managing, and monitoring Windows data | The The Get Windows data section |

| For more information about             | See                                        |
|----------------------------------------|--------------------------------------------|
| Testing and troubleshooting data input | The Improve the data input process section |

## Get \*nix data into Splunk Cloud Platform

This topic guides you through the steps to get \*nix data into Splunk Cloud Platform.

### Administrator requirements

Splunk Cloud Platform administrators must meet the following prerequisites to get \*nix data into Splunk Cloud Platform:

- On your \*nix server, you need root access if you plan to collect from system or root-owned files and directories. However, note the following precautions if you run as a non-root user: see [Run Splunk as a different or non-root user](#).
- Correct permissions to open the following port in your firewalls: 9997
- Knowledge of the location of your source files and which sourcetypes you want Splunk Cloud Platform to recognize.

If you don't know this information or have these permissions, work closely with your organization's \*nix Administrator to complete these steps.

Customers are responsible for the setup, configuration, and maintenance of third-party services and resources, which includes payment. See [Network connectivity and data transfer in the Splunk Cloud Platform Service Description](#).

### Before you begin

To get \*nix data into Splunk Cloud Platform, you need a solid understanding of Splunk concepts. The table lists these concepts and provides links to more information.

| Product                          | Concept                                                                                                                                   | See                                                                   |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| Splunk and Splunk Cloud Platform | <a href="#">indexes</a><br><a href="#">source types</a><br><a href="#">Inputs Data Manager</a><br><a href="#">Splunk apps and add-ons</a> | <a href="#">Fundamental Splunk and Splunk Cloud Platform concepts</a> |

### Overview

This topic takes you through the steps to get \*nix data into Splunk Cloud Platform using Linux commands.

The specific commands and syntax for these examples are run on Amazon Linux 2 AMI; however, the syntax for other \*nix systems may be slightly different. If you are using a different \*nix system, use the equivalent syntax to follow the steps.

To get \*nix data into Splunk Cloud Platform, complete the following high-level steps:

1. [Set up your Splunk Cloud Platform environment](#).
2. [Install and configure a Universal Forwarder on your host system](#).

On your \*nix server, you need to install a Universal Forwarder that will forward data on to your Splunk Cloud

Platform instance.

3. [Download and install the credentials for the Universal Forwarder.](#)

You will need to download and install the Splunk Cloud Platform credentials on the forwarder to allow it to send data to your Splunk Cloud instance.

4. [Install and configure the Splunk Add-on for Unix and Linux on your Universal Forwarder.](#)

On your Universal Forwarder, you will install an add-on to simplify the process of getting \*nix data into Splunk, and you'll configure some source types to ensure your Splunk Cloud Platform instance can recognize the types of sources you need to analyze.

5. [Verify that you can receive data from your \\*nix platform.](#)

Test your configuration to ensure that it's working properly.

## Step 1: Set up your Splunk Cloud Platform environment

Before you can get \*nix data into your Splunk Cloud Platform instance, you must ensure the following:

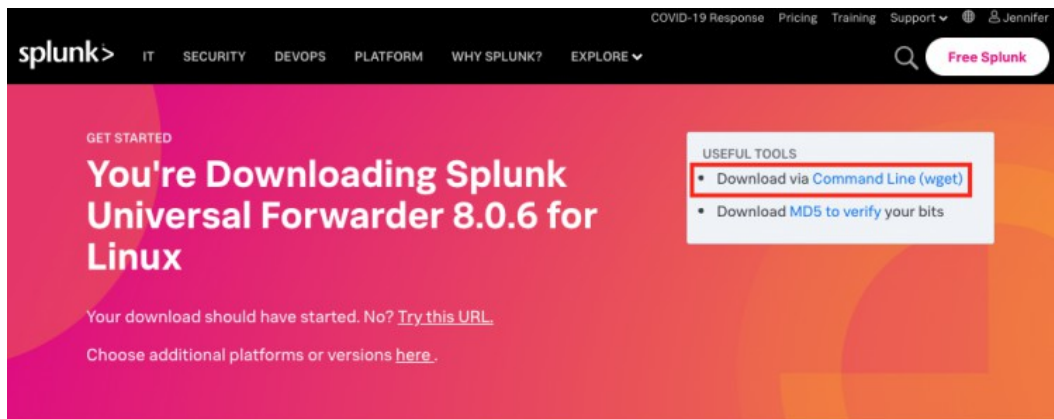
- Confirm that you are assigned the `sc_admin` role on your Splunk Cloud Platform instance.
- Request that Splunk Support install the following on your Splunk Cloud Platform instance, and ensure you allow adequate time to complete this task before you attempt to get data in:
  - ◆ Splunk Add-on for Unix and Linux
  - ◆ Splunk IT Service Intelligence or Splunk IT Essentials Work
  - ◆ Splunk App for Content Packs, which contains the necessary content packs for Unix and Linux
- Create a test index in your Splunk Cloud Platform instance so that you can test your installation before going into production. Follow these instructions to create an index: [Create a Splunk Cloud Platform Index.](#)

## Step 2: Install the Universal Forwarder in your \*nix environment

If you have already installed a Universal Forwarder, you can skip this step.

Complete the following steps to install and configure the Universal Forwarder to send the data to Splunk Cloud Platform.

1. Connect to your \*nix machine, and log in as the root user so you can install a package.
2. Go to [Splunk.com](#) and download the Universal Forwarder to a temporary directory (meaning, `/tmp`).
3. Use the `wget` command to download the forwarder to your Linux environment.  
Copy the code from the [thank-you-universalforwarder.html](#) page that appears after the download has started:



4. To ensure you use the rpm as root, enter the following command:

```
sudo -i
```

5. Use the rpm program to install RPM files.

To install the Splunk RPM in the default directory /opt/splunkforwarder, enter the following command:

```
rpm -i splunkforwarder-<...>-linux-2.6-x86_64.rpm
```

6. Log in as the Splunk user by entering the following command:

```
su - splunk
```

7. Go to the bin directory by entering the following command:

```
cd bin
```

8. Start your forwarder by entering the following command:

```
./splunk start
```

9. Enter a user name and password.

You should see the installation performing the steps of checking prerequisites, creating certs, checking conf files, and validating files against a hash. If the installation is successful, a message similar to the following displays:

```
All installed files intact.
Done
All preliminary checks passed.
Starting splunk server daemon (splunkd)...
Done
```

### Step 3: Download the credentials file and install it on your Universal Forwarder

You can skip this step if you have already downloaded and installed the credentials package.

Complete the following steps to install the credential files on the Universal Forwarder.

1. From your Splunk Cloud Platform instance, go to **Apps > Universal Forwarder**.
2. Click **Download Universal Forwarder Credentials**.
3. Note the location where the credentials file was downloaded. The credentials file is named `splunkclouduf.spl`.
4. Copy the file to your `/tmp` folder.
5. Install the following app by entering the following command:  

```
/opt/splunkforwarder/bin/splunk install app /tmp/splunkclouduf.spl
```
6. When you are prompted for a user name and password, enter the user name and password for the Universal Forwarder. The following message displays if the installation is successful:  

```
App '/tmp/splunkclouduf.spl' installed
```
7. Restart the forwarder to enable the changes by entering the following command.  

```
./splunk restart
```

## Step 4: Install and configure the Splunk Add-on for Unix and Linux on your Universal Forwarder

Complete the following steps to download the Splunk Add-on for Unix and Linux from Splunkbase, install it on your forwarder, and enable the inputs.

1. Go to Splunkbase, and download the Splunk Add-on for Unix and Linux.  
 Make sure you download and install the Splunk Add-on for Unix and Linux, and not the Splunk App for Unix and Linux. In Splunkbase, the icon for the add-on is yellow and black, while the icon for the app is green and black.
2. Copy the file (`splunk-add-on-for-unix-and-linux_602.tgz`) from the download location to the `/tmp` directory on the forwarder. You can use `scp` or another similar program.
3. Ssh/login to the forwarder instance.
4. Sudo or su to the splunk user.
5. Ensure that you are logged in as the splunk user using the following command: `whoami`
6. Untar the file using the following command:  

```
tar xfvz splunk-add-on-for-unix-and-linux_602.tgz
```
7. Move the files to the `Splunk_TA_nix` directory using the following command:  

```
mv Splunk_TA_nix/ /opt/splunkforwarder/etc/apps/
```
8. Go to the `apps` directory by entering the following command:  

```
cd /opt/splunkforwarder/etc/apps
```
9. Add the following directory:  

```
Splunk_TA_nix.
```

 You can see the list of directories by using the following command:  

```
ls -F
```
10. Go to the `Splunk_TA_nix` directory using the following command:  

```
cd Splunk_TA_nix
```
11. Create a local directory using the following command:  

```
mkdir local
```
12. Verify that the directory was created by entering:  

```
ls -F
```
13. Copy the `inputs.conf` file from the default directory to your local directory by entering the following command:  

```
cp default/inputs.conf local
```
14. Go to your local directory by entering:  

```
cd local
```
15. Open the file using your preferred text editor. In this case, we used `nano` by entering the command:  

```
nano inputs.conf
```
16. When you open the file for editing, you can see the inputs related to the \*nix operating system. Note that each of the inputs is disabled and each input displays as `disabled = 1`.
17. Change the inputs to read: `disabled=0`. This enables the inputs. You may later decide to disable some of these inputs when you become more familiar with them.

18. Enter the proper command for your chosen text editor to save the changes. In nano, use the following command:`ctrl + o`
19. Restart your forwarder to enable the changes. Go to the bin directory by entering:`cd bin`. Then enter `./splunk restart`.
20. The forwarder will notify you when it has restarted.

## Step 5: Verify that you can receive data from your \*nix platform

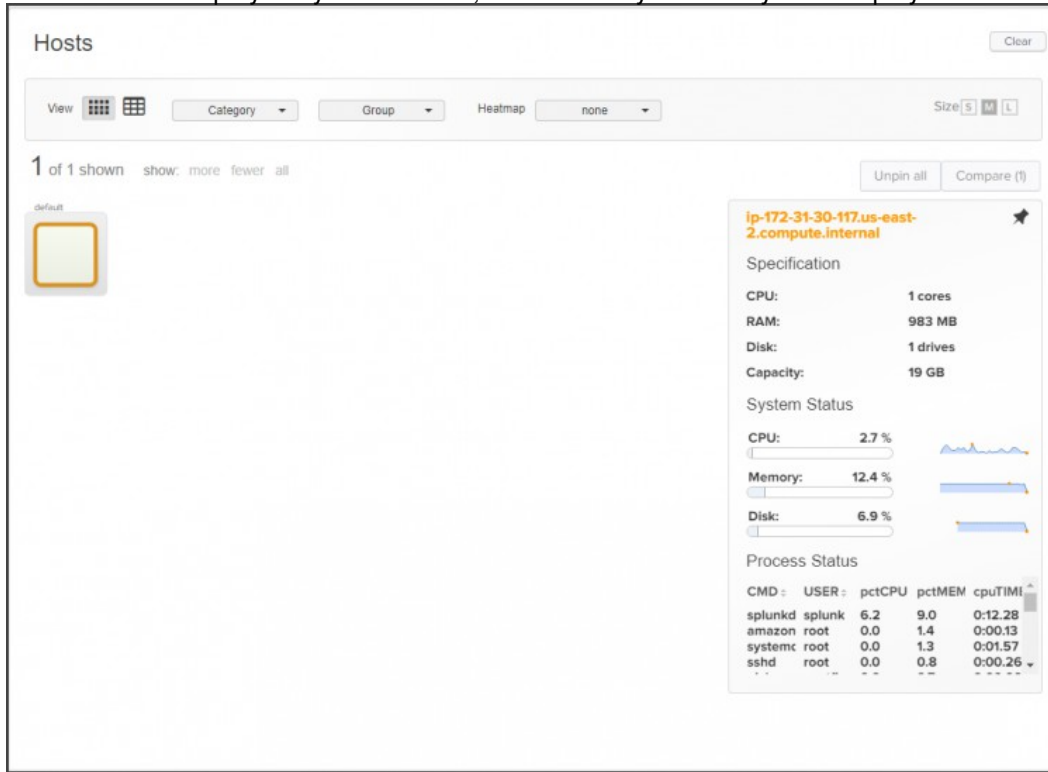
Certain steps in this procedure use functionality available in the Splunk Content Pack for Unix Dashboards and Reports, one of the content packs that was installed in Step 1 by Splunk Support. See the Content Pack for Unix Dashboards and Reports manual.

Complete the following steps to verify you can receive data from your \*nix platform.

1. Open your Splunk Cloud Platform instance.
2. Access the **Settings - Unix** dashboard to configure the correct index.  
See Configure the Content Pack for Unix and Dashboards and Reports.
3. If you configured a test index, set the index value to your test index. Otherwise, enter `index = main`.
4. Click **Save**.
5. From **Apps > Search and Reporting**, enter the search term `index=*` to do a search of incoming data.
6. In **Selected fields > hosts** field, select the host that corresponds to your \*nix operating system.
7. From the Selected fields, choose **sourcetype**. A list of \*nix sourcetypes like the following displays:

| Top 10 Values | Count  | %       |
|---------------|--------|---------|
| lsyf          | 10,127 | 54.385% |
| ps            | 5,373  | 28.854% |
| top           | 2,791  | 14.988% |
| cpu           | 120    | 0.644%  |
| bandwidth     | 30     | 0.161%  |
| interfaces    | 30     | 0.161%  |
| iostat        | 30     | 0.161%  |
| netstat       | 30     | 0.161%  |
| protocol      | 30     | 0.161%  |
| vmstat        | 30     | 0.161%  |

8. Access the **Hosts** dashboard.  
See Use the Hosts dashboard.
9. Your \*nix host displays. If you click on it, statistics for your \*nix system display:



## What's next?

Now that you have configured your Splunk Cloud Platform instance to get data from your \*nix system, you may want to use a deployment server to propagate the settings across multiple forwarders. The table lists additional topics of interest.

### See also

| For more information about                                                                                                                                  | See                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| <b>deployment server:</b> A tool for distributing configurations, apps, and content updates to groups of Splunk Enterprise instances, including forwarders. | About deployment server and forwarder management                                                      |
| Testing and troubleshooting data input                                                                                                                      | The Improve the data input process section in the Splunk Cloud Platform <i>Getting Data In</i> manual |

## Get Windows Data into Splunk Cloud Platform

This topic guides you through the steps to get Windows data into Splunk Cloud Platform.

### Administrator requirements

Splunk Cloud Platform administrators must meet the following prerequisites to get Windows data into Splunk Cloud Platform:

- Local Admin access on your Windows machines to install the Splunk Universal Forwarder.



- Permissions to open the Windows ports tcp/8089 inbound and tcp/9997 outbound, required for Step 1.

If you don't know this information or have these permissions, work closely with your organization's Windows administrator to complete these steps.

Customers are responsible for the setup, configuration, and maintenance of third-party services and resources, which includes payment. See Network connectivity and data transfer in the Splunk Cloud Platform Service Description.

## Before you begin

To get Windows data into Splunk Cloud Platform, you need a solid understanding of various Splunk concepts. The table lists these concepts and provides links to more information.

| Product                          | Concept                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | See                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Splunk and Splunk Cloud Platform | <b>indexes</b><br><b>source types</b><br><b>deployment server</b><br><b>universal forwarder</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <a href="#">Fundamental Splunk and Splunk Cloud Platform concepts</a>                                                                                                                                                                                                                                                                                                                                                            |
| Splunk Applications and Add-ons  | <p>In this configuration, you use the Splunk Universal Forwarder app to get data in, and the Splunk Add-on for Windows to simplify the process of getting data in. A Splunk app is an application that runs on the Splunk platform and typically addresses several use cases. Add-ons support and extend the functionality of the Splunk platform and the apps that run on it, usually by providing inputs for a specific technology or vendor. The Splunk Add-on for Windows allows a Splunk software administrator to collect:</p> <ul style="list-style-type: none"> <li>• CPU, disk, I/O, memory, log, configuration, and user data with data inputs.</li> <li>• Active Directory and Domain Name Server debug logs from Windows hosts that act as domain controllers for a supported version of a Windows Server. In some cases, you may need to configure Active Directory audit policy since Active Directory does not log certain events by default.</li> <li>• Domain Name Server debug logs from Windows hosts that run a Windows DNS Server. Windows DNS Server does not log certain events by default, and you must enable debug logging. Generally, you need to install the app on your Splunk Cloud Platform instance, and the add-on on your forwarder and Splunk Cloud Platform instance.</li> </ul> | <p>For more information about:</p> <ul style="list-style-type: none"> <li>• Add-ons: <a href="#">About Splunk add-ons</a></li> <li>• Splunk Add-on for Windows: <a href="#">About the Splunk Add-on for Windows</a></li> <li>• Best practices in Windows logging, see <a href="http://www.malwarearchaeology.com/cheat-sheets/">Malware Archeology's cheat sheet: http://www.malwarearchaeology.com/cheat-sheets/</a></li> </ul> |

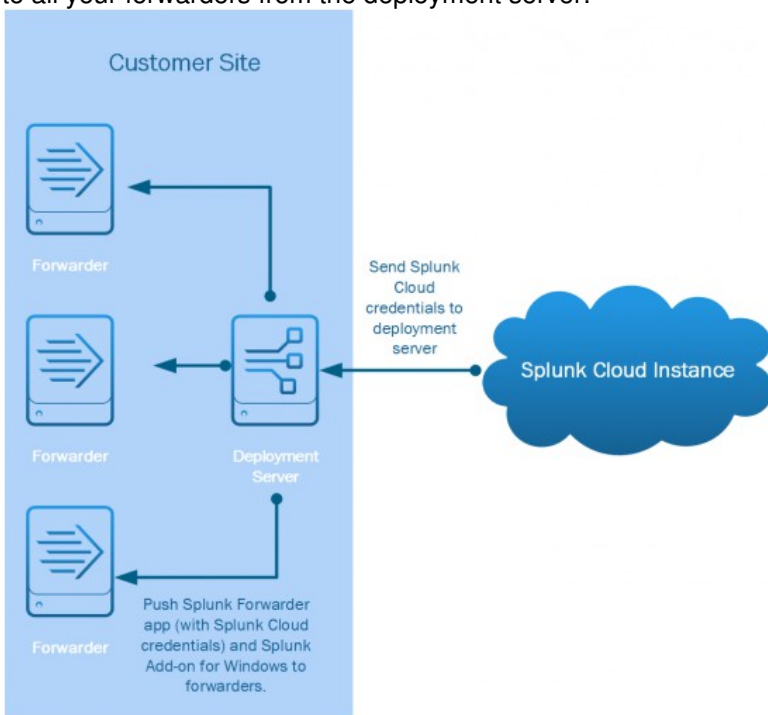
## Overview

In this procedure, you'll get your Windows applications, security, and system data into your Splunk Cloud Platform instance. There are many other types of Windows data you may want to include in your production environment, but you will likely want these logs at a minimum, and you can later add other types of logs.

To get Windows data into Splunk Cloud Platform, complete the following high-level steps:

1. [Set up your Windows environment.](#)
2. [Set up your Splunk Cloud Platform environment.](#)
3. [Configure indexes on your Splunk Cloud Platform instance.](#)  
You create an index for each of the types of data you want to bring into you Splunk Cloud deployment.
4. [Configure your deployment server.](#)  
The deployment server allows you to centrally manage the Splunk Forwarders in your environment. Using the deployment server you can configure what data gets collected and where to send it. In this case, you use the deployment server to send data to your Splunk Cloud instance.
5. [Configure apps and add-ons on your deployment server.](#)  
You configure the Splunk Universal Forwarder app on the deployment server, and you configure the Splunk Add-on for Windows on your deployment server. Then you set up server classes so that you can push the configurations to the forwarders on your Windows machines.
6. [Configure Universal Forwarders on your Windows Machines.](#)  
Forwarders are used to collect data and forwarder data to your Splunk Cloud Platform instance.
7. [Verify that data is flowing to your Splunk Cloud Platform instance.](#)  
After configuring the deployment server, add-on and forwarders, check to see if data is flowing to your Splunk Cloud Platform instance.

The following graphic shows how add-on settings and forwarder settings are configured on the deployment server and pushed to groups of forwarders on the customer site. When you have configured all your settings, you can push updates to all your forwarders from the deployment server:



## Step 1: Set up your Windows environment

Complete the following steps to set up your Windows environment.

- On the server that will host your deployment server, open tcp/8089 inbound to allow communication with the deployment server from deployment clients. This can be a Linux or a Windows server.
- Open port tcp/9997 outbound on your network firewall to allow communication with the Splunk Cloud indexers.

If you have security concerns that prevent you from opening multiple ports on your firewall, you may want to create an intermediate forwarding tier to limit the number of open ports. For more information about this topic, see <https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf>.

## Step 2: Set up your Splunk Cloud Platform environment

Complete the following steps to set up your Splunk Cloud Platform environment.

- Confirm that you are assigned the `sc_admin` role on your Splunk Cloud Platform instance.
- Request that Splunk Support install the Splunk Add-on for Microsoft Windows on your Splunk Cloud Platform instance. Ensure you allow adequate time to complete this task before you attempt to get data in.
- Request a 0 MB deployment server license from Splunk Support. Ensure you allow adequate time to complete this task.

## Step 3: Configure indexes on your Splunk Cloud Platform instance

Create indexes to store the events you send from your Windows machines. It's a best practice to create separate indexes for different types of data. This can be useful if you want different storage settings for different types of data. For example, you may need to store wineventlogs for a specified time period for compliance purposes.

In this step, you create the following indexes:

- **wineventlog**: Store windows event logs
- **perfmon**: Store windows performance data
- **msad**: Store Microsoft Active Directory data
- **dns**: If collecting, store dns data
- **dhcp**: If collecting, store dhcp data

Complete the following steps to create an index.

1. From your Splunk Cloud Platform instance, go to **Settings > Indexes**.
2. Click **New Index**.
3. For the index name, enter **wineventlog**.
4. For index data type, select **Events**.
5. For **searchable time (days)**, enter a value that indicates the number of days the data is searchable. The image shows an example of 90 days of searchable storage.
  - ◆ Storage is based on your subscription type. For more information on an appropriate storage value per your subscription type, see Storage in the *Splunk Cloud Platform Service Description*. Be sure to refer to the correct service description version for your deployment.
  - ◆ Optionally, you can extend your storage for longer if you have different requirements. Discuss your storage requirements with your Splunk account representative.
6. Click **No Additional Storage**, and click **Save**:

You can also set up different types of storage for expired Splunk Cloud Platform data (such as self-storage or archiving).

The screenshot shows the 'New Index' configuration interface. The 'Index name' field is set to 'wineventlog'. The 'Index Data Type' is set to 'Events'. The 'Searchable time (days)' is set to '90'. Under 'Dynamic Data Storage', the 'No Additional Storage' option is selected. The 'Save' button is highlighted in green.

7. Repeat these steps for each of the following indexes:

- ◆ perfmon
- ◆ msad
- ◆ dns
- ◆ dhcp

## Step 4: Configure your Splunk Deployment Server

Complete the following steps to configure the deployment server (Windows OS) with the deployment server license and the Universal Forwarder App.

1. Download a Splunk Enterprise instance as your deployment server.

From Splunk.com download an instance of Splunk Enterprise and install it on its own Windows machine.

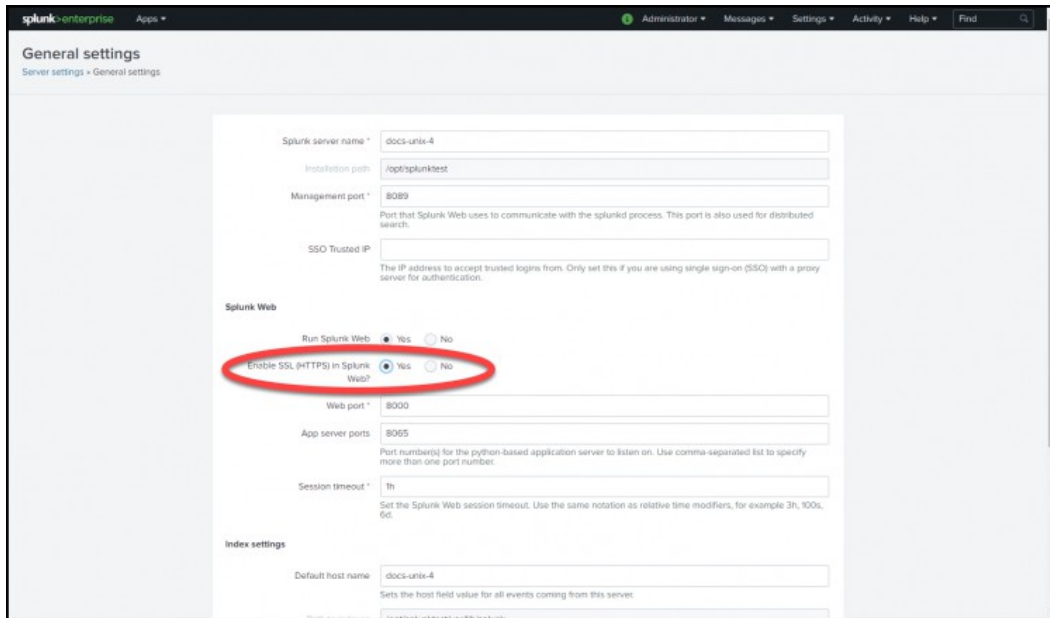
Do not install on the same machine as a Universal Forwarder.

2. Configure HTTPS for Splunk Web.

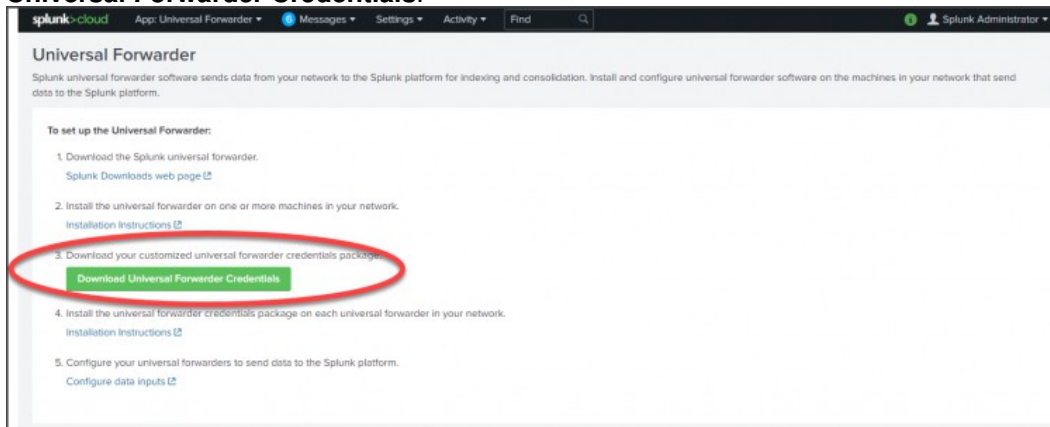
From the Splunk Enterprise instance you installed, go to **Settings > Server settings > General Settings**.

In the field **Enable SSL (HTTPS) in Splunk Web**, click **Yes**, and click **Save**.

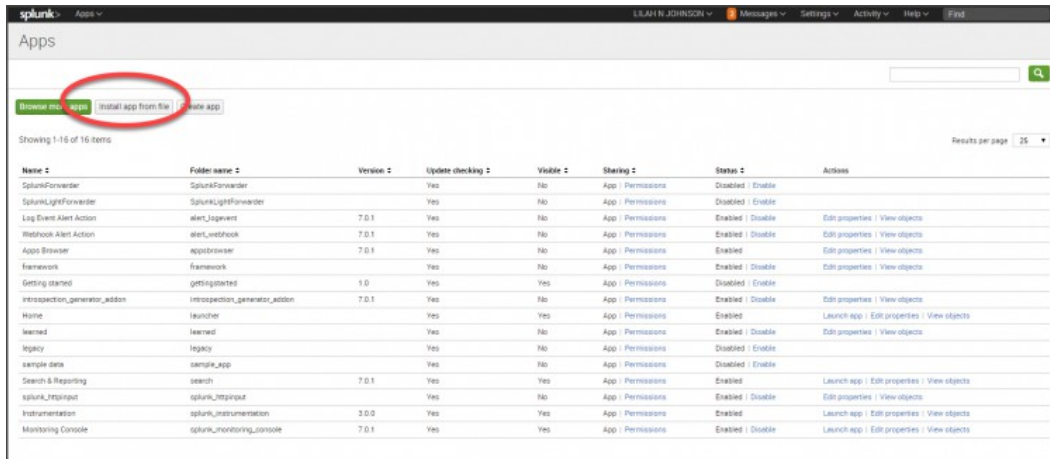
This is a best practice for security. For additional security you can add your own certificate instead of using the default certificates.



3. Download the Universal Forwarder credentials to install the Universal Forwarder App. Log into your Splunk Cloud Platform instance. Under **Apps**, click **Universal Forwarder**, then click **Download Universal Forwarder Credentials**.



4. Upload the Universal Forwarder credentials on your deployment server. On your deployment server (the Splunk Enterprise instance you will use as a deployment server), go to **Apps > Manage Apps > Install Apps from file**. Click **Upload** to upload the Universal Forwarder app.



5. Configure the licensing for the deployment server.

From **Settings > Licensing** use the license to configure the Splunk instance as a deployment server.

This is the license you requested from Splunk Support in Step 2.

6. Click **Restart later**.

## Step 5: Configure Apps and Add-Ons on your Deployment Server

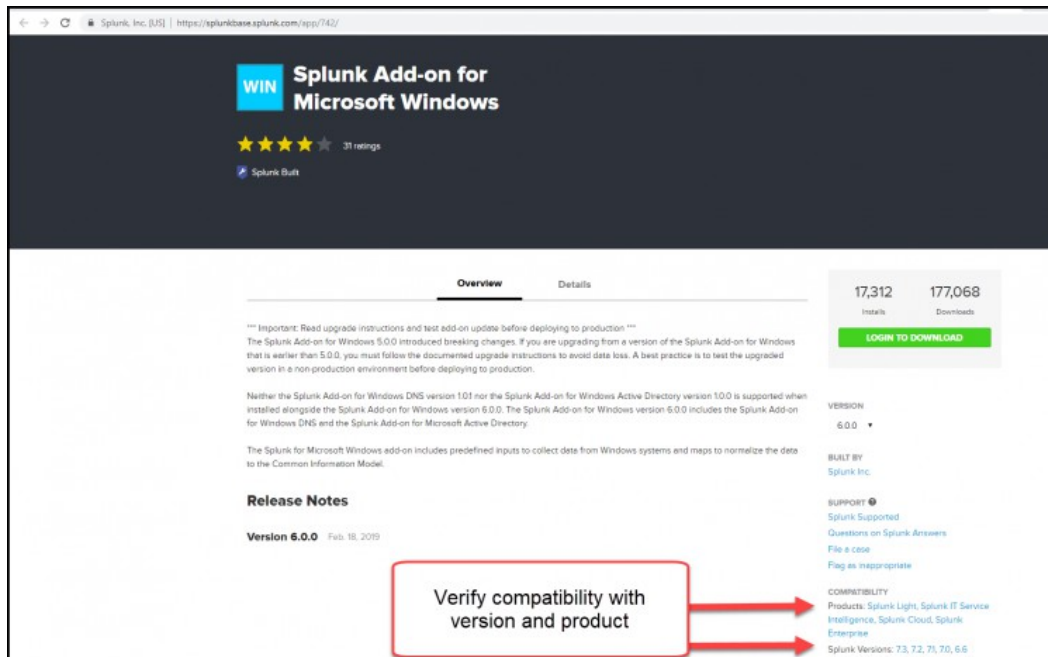
Add the Universal Forwarder app and the Splunk Add-on for Windows to your deployment server so that it can push forwarder and add-on configurations to all of the forwarders you install.

Complete the steps in the following sections to configure apps and add-ons on your deployment server.

### **Download and install the add-on**

1. Go to Splunkbase and download the Splunk Add-on for Microsoft Windows.

As a best practice, use the **COMPATIBILITY** field on the left side to verify that the add-on is valid for Splunk Cloud Platform and the version you have installed.



2. On your deployment server, click **Apps > Manage Apps > Install Apps from file**, then click **Upload** to upload the Splunk Add-on for Microsoft Windows you downloaded from splunkbase.
3. Verify these folders are in the right directory by going to **Windows > Program Files > Splunk > etc > apps** and checking for the following folders:
  - ◆ 100\_<splunk cloud stack name>\_splunkcloud
  - ◆ Splunk\_TA\_windows
4. Copy these folders to the following directory: **Windows > Program Files > Splunk > etc > deployment-apps**. After copying the folders, make sure that no local folder exists under the Splunk Forwarder app **Windows > Program Files > Splunk > etc > deployment-apps > 100\_<splunk cloud stack name>\_splunkcloud**. If a local folder exists, delete it. This folder gets created when the app is installed but you need a unique outputs.conf for each forwarder. This gets recreated when the Universal Forwarder restarts.
5. Perform a verification step:  
Return to the Forwarder Management console by going to **Settings > Forwarder management**. The Universal Forwarder app and the Splunk Add-On for Microsoft Windows should be listed under the Apps tab.

### **Configure and customize the Windows data collection add-ons**

1. Navigate to **Windows > Program Files > Splunk > etc > deployment-apps**.
2. Make copies of the Splunk\_TA\_windows folder for each of the types of Windows instances that you want to get data from.
3. Rename each of the folders so that they represent your different Windows servers. For this example, create the following folders:
  - ◆ Splunk\_TA\_windows\_DomainController
  - ◆ Splunk\_TA\_windows\_server
  - ◆ Splunk\_TA\_windows\_client
  - ◆ Splunk\_TA\_windows\_GlobalCatalogServer
4. Navigate to **Windows > Program Files > Splunk > etc > deployment-apps > Splunk\_TA\_windows\_server**.
5. In the folder, create a new folder called **local**.  
This is a Splunk best practice and ensures that your configuration changes are saved during an upgrade. Also, this provides a way to revert back to the original configurations if some settings are misconfigured.

6. From **Windows > Program Files > Splunk > etc > deployment -apps > Splunk\_TA\_windows\_server > default**, copy the file `inputs.conf` into your **local** folder.
7. Using a file editor, open the `inputs.conf` file for editing.  
Tip: Use Wordpad or Notepad++ rather than Notepad, which does not handle word wrapping correctly by default.
8. Review the Source Types for Windows Add-Ons in the documentation to ensure that your sources are represented by this add-on.

In this instance, you configure the add-on to get data in for the following Windows Event Logs:

- ◆ Application
- ◆ Security
- ◆ System

9. To get the Application log data in, modify the `inputs.conf` file as follows:
  - ◆ For `WinEventLog://Application`, set `disabled = 0`. This enables the input.
  - ◆ Add an entry for the location of the index by adding the following line to the stanza: `index=wineventlog`. This is the index you previously configured).

The example shows resulting stanza. The bold font shows which lines are changed or added.

```
[WinEventLog://Application]
disabled = 0
start_from = oldest
current_only = 0
checkpointInterval = 5
renderXml=true
index=wineventlog
```

10. To get the Security Log data in, modify the `inputs.conf` file as follows:
  - ◆ For `WinEventLog://Security`, set `disabled = 0`.
  - ◆ Add an entry for the location of the Security log files by adding the following stanza: `index=wineventlog`.

The example shows resulting stanza. The bold font shows which lines are changed or added.

```
[WinEventLog://Security]
disabled = 0
start_from = oldest
current_only = 0
evt_resolve_ad_obj = 1
checkpointInterval = 5
blacklist1 = EventCode="4662" Message="Object Type:(?!\\s*groupPolicyContainer)"
blacklist2 = EventCode="566" Message="Object Type:(?!\\s*groupPolicyContainer)"
renderXml=true
index=wineventlog
```

11. Save the `inputs.conf` file.
12. Perform a verification step:  
Return to the Forwarder Management console by going to **Settings > Forwarder management**. The Universal Forwarder app and the Splunk Add-On for Microsoft Windows and the modified ones you created should be listed under the **Apps** tab.

### **Create and configure a server class inside the Forwarder Management Console**

1. Under the **Server Classes** tab, click **New Server Class**.
2. Enter **outputs** as the server class name and click **Save**.  
In this case, name the server class **outputs** because it sets the `outputs.conf` file for the forwarders.
3. When you save these changes, you are taken to a screen to add apps or a client. Click **Add Apps** and select the Universal Forwarder app.
4. The name of the Universal Forwarder app is unique to your Splunk Cloud instance (example: `100_<stack_name>_splunkcloud`).  
Click the name to add it to the right side and then click **Save**.
5. Under **Actions for the Universal Forwarder** app, click **Edit**.



6. Select the **Enable App and Restart Splunkd** the checkbox and click **Save**.  
Setting **Restart Splunkd** lets you to restart the forwarder after you push changes to the apps via the deployment server.
7. Navigate to the **Server Classes** tab.
8. For the outputs server class, click **Edit > Edit Clients**.
9. In the **Include (whitelist)** box, enter a wildcard (\*) so that the Universal Forwarder app is deployed to all of your Universal Forwarders as they get installed and phone home to the deployment server.
10. Repeat steps 1-9 for another server class called Windows servers.
  - ◆ For step 2 customize for Windows servers.
  - ◆ For step 3 customize to the Splunk\_TA\_windows\_Server app.
11. Perform a verification step:  
When you view the apps from the deployment server, you should see that the app and **Restart Splunkd** are enabled.

## Step 6: Install the Splunk Universal Forwarder on your Windows Servers

Install a Universal Forwarder on each of the Windows servers from which you want data. The easiest way to do this is to run the installer on your server.

Complete the following steps.

1. From splunk.com, download the Universal Forwarder to your Windows server.
2. Once the download is complete, click on the file to start the install.
3. Clear the checkbox **Uncheck if you want to use Splunk Cloud**.
4. Set a username and password.
5. In the **Deployment Server** field, enter the name of the deployment server.  
For example, **win2016-splk-ds**. As a best practice, include the full DNS name. For the port, enter port **8089** to allow the Universal Forwarder to communicate with the deployment server.
6. Click **Next**, and click **Install**.  
The forwarder is installed on your server, and you have instructed it to check the deployment server for configuration settings. Once the forwarder is running, it checks with the deployment server and downloads any apps you have configured. In this case, it downloads the Universal Forwarder app and the Splunk Add-on for Windows.
7. Repeat these steps for each of the Windows machines where you want to send data to Splunk Cloud Platform.
8. Perform a verification step:  
To verify that your forwarders are configured correctly, you can return to your deployment server, and from the Forwarder Management page, check to see if your clients have checked in. If the clients (forwarders) have checked in, you can see them listed in the Clients tab on the Forwarder Management page.

## Step 7: Verify that Data is Flowing to Splunk Cloud Platform

Complete the following steps to verify that data is flowing to Splunk Cloud Platform.

1. From your Splunk Cloud instance, go to **Apps > Search and Reporting**.
2. In the search field, enter **index=\_internal host!= "\*.splunkcloud.com"**.  
This search allows you view events from any host that is not a splunkcloud.com instance. So, you should be able to see any other hosts that are sending data to your Splunk Cloud Platform instance.
3. For the time range, select **presets > last 30 days**.  
This allows you to start seeing data more quickly because the oldest events populate first.
4. Click the search icon. Events from your Windows machines should display.
5. In the left pane, a list of fields displays. Under the **Host** field you can see which forwarders are sending data to Splunk Cloud Platform.

## What's next?

The table lists additional topics of interest in the Splunk Cloud Platform *Getting Data In* manual.

### See also

| For more information about                     | See                                        |
|------------------------------------------------|--------------------------------------------|
| Getting, managing, and monitoring Windows data | The The Get Windows data section           |
| Testing and troubleshooting data input         | The Improve the data input process section |

## Forward data from files and directories to Splunk Cloud Platform

This topic tells you how to configure and run the universal forwarder to forward the data from local files and directories. It also provides command examples for common scenarios.

### See also

| For more information about                                          | See                                                                                                                      |
|---------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Getting data from files and directories using Splunk Cloud Platform | The topics in the Get data from files and directories chapter in the Splunk Cloud Platform <i>Getting Data In</i> manual |
| Details about other options for forwarding data                     | Splunk Universal Forwarder Manual                                                                                        |

## Start and restart the universal forwarder

To start the universal forwarder, go to the `$SPLUNK_HOME/bin/` directory and run the `splunk start` command. After changing settings for a forwarder, you must restart the forwarder by issuing the `splunk restart` command. To verify that the desired data is being forwarded to Splunk Cloud Platform, use the Splunk Web Search app.

## Configure the universal forwarder to forward data

To configure forwarding, use the commands and parameters listed in the following tables.

### Commands

To configure forwarding of data in files, use the commands in this table.

| Command               | Command syntax                                                | Description                                                                                                                                                                                                                                                                                                               |
|-----------------------|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>add monitor</b>    | <pre>add monitor &lt;source&gt; [-parameter value] ...</pre>  | Start monitoring the specified input. The forwarder watches for changes to the specified source and forwards data to your Splunk Cloud Platform deployment until you remove the source. For example, to continuously monitor the files in the <code>/var/log/</code> directory: <code>splunk add monitor /var/log/</code> |
| <b>edit monitor</b>   | <pre>edit monitor &lt;source&gt; [-parameter value] ...</pre> | Edit a data input that Splunk Cloud Platform is monitoring.<br><br>For example, to move a log file from the default location to <code>C:\windows\system32\LogFiles\W3SVC</code> , run the following command:<br><br><code>splunk edit monitor C:\windows\system32\LogFiles\W3SVC</code>                                   |
| <b>remove monitor</b> | <pre>remove monitor &lt;source&gt;</pre>                      | Stop monitoring the specified input                                                                                                                                                                                                                                                                                       |

| Command                             | Command syntax                                                                                                                                     | Description                                                                                                                                                                                                                                                                                     |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                     |                                                                                                                                                    | For example, to stop monitoring of the Windows log file that contains all automatic update activity, run the following command:<br><br>splunk remove monitor C:\Windows\windowsupdate.log                                                                                                       |
| <b>list monitor</b>                 | list monitor                                                                                                                                       | Displays a list of all configured data inputs.                                                                                                                                                                                                                                                  |
| <b>add oneshot<br/>or<br/>spool</b> | add oneshot<br><source><br>[-parameter value]<br>...<br><br>Or:<br>spool <source><br>[-parameter value]<br>...<br><br>splunk spool /var/log/applog | Use this command to forward the contents of the specified data source once.<br><br>For example, the following commands perform a one-time forwarding of the contents of the /var/log/applog directory.<br><br>splunk add oneshot /var/log/applog<br><br>or:<br><br>splunk spool /var/log/applog |

### Parameters

You can use the parameters in the following table with data input commands.

| Parameter              | Required | Description                                                                                                                                                                                                                                                                                                       |
|------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <source>               | Yes      | Specify the path to the file or directory that contains the data you want to monitor or upload.<br><br>The syntax for this parameter is the value. It is not preceded with the <code>-source</code> parameter flag. For example, enter <code>&lt;source&gt;</code> ", not <code>"-source &lt;source&gt;</code> ". |
| sourcetype             | No       | Specify a single source type for the data <code>&lt;source&gt;</code> . The source type determines how events are formatted and is a default field that is included in all events.                                                                                                                                |
| hostname<br>or<br>host | No       | Specify a single host or host name for the data <code>"&lt;source&gt;</code> ". This default field is included in all events.                                                                                                                                                                                     |

### Common command examples

This section provides command examples for monitoring files and logs and uploading a file.

| Description                                                                                      | Command                                                                                 |
|--------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Monitor the files in the /var/log/ directory (Unix)                                              | splunk add monitor /var/log/                                                            |
| Monitor C:\Windows\windowsupdate.log                                                             | splunk add monitor C:\Windows\windowsupdate.log                                         |
| Monitor the default location for Windows IIS logging                                             | splunk add monitor C:\windows\system32\LogFiles\W3SVC                                   |
| Monitor a set of log files in a directory, specifying metadata to be used by the Splunk indexers | splunk add monitor /tmp/foo/*.log -index se_test -sourcetype insurgency -host vm_host01 |
| One-time upload of a file                                                                        | splunk add oneshot /var/log/applog                                                      |

### Upgrade your Forwarders

If you are using either heavy or universal forwarders, maintaining version compatibility between your forwarders and Splunk Cloud Platform environment ensures there is no interruption to your service. In addition, when forwarders are

version compatible with your Splunk Cloud Platform environment, you can immediately take advantage of new capabilities.

As a best practice, run the most recent forwarder version, even if the forwarder is a higher version number than your Splunk Cloud Platform environment.

To upgrade a heavy or universal forwarder for your Splunk Cloud Platform environment, see the appropriate section in this topic.

**See also**

| For more information about                                                                                                   | See                                                                                       |
|------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Supported forwarder versions, their compatible Splunk Cloud Platform versions, and respective end-of-support milestone dates | Supported forwarder versions in the <i>Splunk Cloud Platform Service Description</i>      |
| The Splunk universal forwarder                                                                                               | Splunk Universal Forwarder <i>Forwarder Manual</i>                                        |
| Upgrading a universal forwarder to a heavy forwarder                                                                         | Upgrade the universal forwarder in the Splunk Universal Forwarder <i>Forwarder Manual</i> |

### Upgrade the \*nix universal forwarder

To upgrade a \*nix universal forwarder for a Splunk Cloud Platform deployment, see Upgrade the universal forwarder in the Splunk Universal Forwarder *Forwarder Manual*.

### Upgrade the Windows universal forwarder

To upgrade a Windows universal forwarder for a Splunk Cloud Platform deployment, see Upgrade the universal forwarder in the Splunk Universal Forwarder *Forwarder Manual*.

### Upgrade a heavy forwarder on \*nix

This section describes how Splunk Cloud Platform administrators can upgrade a heavy forwarder on a \*nix machine for their Splunk Cloud Platform deployment.

#### **Before you upgrade**

Before you upgrade, see About upgrading: READ THIS FIRST for information on changes in the new version that can impact you if you upgrade from an existing version.

Your Splunk Heavy Forwarder does not provide a means of downgrading to previous versions. If you need to revert to an older Splunk Forwarder, uninstall the upgraded version and reinstall the version you want.

#### **Back your files up**

Before you perform the upgrade, **back up all of your files**.

For information on backing up configurations, see Back up configuration information in the *Splunk Enterprise Admin Manual*.

## How upgrading works

To upgrade a heavy forwarder installation, you must install the new version directly on top of the old version (into the same installation directory.) When the Splunk Heavy Forwarder starts after an upgrade, it detects that the files have changed and asks whether or not you want to preview the migration changes before it performs the upgrade.

If you choose to view the changes before proceeding, the upgrade script writes the proposed changes to the `$$SPLUNK_HOME/var/log/splunk/migration.log.<timestamp>` file.

Splunk Heavy Forwarder does not change your configuration until after you restart it.

As a best practice, run the most recent forwarder version, even if the forwarder is a higher version number than your Splunk Cloud Platform environment.

## Upgrade a Splunk Heavy Forwarder

1. Download the full version of Splunk Enterprise that you want to upgrade to from the [splunk.com](http://splunk.com) website.
2. Open a shell prompt on the machine that has the instance that you want to upgrade.
3. Change to the `$$SPLUNK_HOME/bin` directory.
4. Run the `$$SPLUNK_HOME/bin/splunk stop` command to stop the instance.
5. Confirm that no other processes can automatically start the Splunk Heavy Forwarder.
6. To upgrade and migrate, install the Splunk Heavy Forwarder package directly over your existing deployment.
  - ◆ If you use a `.tar` file, expand it into the same directory with the same ownership as your existing Splunk Heavy Forwarder instance. This overwrites and replaces matching files but does not remove unique files.  

```
tar xzf splunk-7.x.x-<version-info>.tgz -C $$SPLUNK_HOME
```
  - ◆ If you use a package manager, such as RPM, type `rpm -U splunk_package_name.rpm`
  - ◆ If you use a `.dmg` file on Mac OS X, double-click it and follow the instructions. Specify the same installation directory as your existing installation.

7. Run the `$$SPLUNK_HOME/bin/splunk start` command.

The Splunk Heavy Forwarder displays the following output.

```
This appears to be an upgrade of Splunk.
```

```
-----  
Splunk has detected an older version of Splunk installed on this machine. To  
finish upgrading to the new version, Splunk's installer will automatically  
update and alter your current configuration files. Deprecated configuration  
files will be renamed with a .deprecated extension.  
You can choose to preview the changes that will be made to your configuration  
files before proceeding with the migration and upgrade:  
If you want to migrate and upgrade without previewing the changes that will be  
made to your existing configuration files, choose 'y'.  
If you want to see what changes will be made before you proceed with the  
upgrade, choose 'n'.  
Perform migration and upgrade without previewing configuration changes? [y/n]
```

8. Choose whether or not you want to run the migration preview script to see proposed changes to your existing configuration files, or proceed with the migration and upgrade right away. If you choose to view the expected changes, the script provides a list.
9. After you review these changes and are ready to proceed with migration and upgrade, run `$$SPLUNK_HOME/bin/splunk start` again.

## Upgrade and accept the license agreement simultaneously

After you place the new files in the Splunk Heavy Forwarder installation directory, you can accept the license and perform the upgrade in one command.

- To accept the license and view the expected changes (answer 'n') before continuing the upgrade, use the following command.

```
$SPLUNK_HOME/bin/splunk start --accept-license --answer-no
```

- To accept the license and begin the upgrade without viewing the changes (answer 'y').

```
$SPLUNK_HOME/bin/splunk start --accept-license --answer-yes
```

## Upgrade a heavy forwarder on Windows

You can upgrade with either the GUI installer or the `msiexec` utility on the command line as described in "Install on Windows via the command line".

Splunk does not provide a means of downgrading to previous versions.

After you upgrade Splunk Heavy Forwarder, if you need to downgrade, you must uninstall the upgraded version and then reinstall the previous version of Splunk Heavy Forwarder that you were using. Do not attempt to install over an upgraded installation with an installer from a previous version, as this can result in a corrupt instance and data loss.

As best practice, run the most recent forwarder version, even if the forwarder is a higher version number than your Splunk Cloud Platform environment.

### **Before you upgrade**

Before you upgrade, see [About upgrading: READ THIS FIRST](#) for information on changes in the new version that can impact you if you upgrade from an existing version.

Splunk Heavy Forwarder does not provide a means of downgrading to previous versions. If you need to revert to an older Splunk Heavy Forwarder release, uninstall the upgraded version and reinstall the version you want.

### **The Windows domain user must match what you specified at installation**

If you installed Splunk Heavy Forwarder with a domain user, you must specify the same domain user explicitly during an upgrade. If you do not, Splunk Heavy Forwarder installs the upgrade as the Local System user. If you do not do this, or you specify the wrong user accidentally during the upgrade, then see [Correct the user selected during installation](#) to switch to the correct user before you start Splunk Heavy Forwarder.

### **Changing heavy forwarder ports during an upgrade is not supported**

Splunk Heavy Forwarder does not support changing the management or Splunk Web ports when you upgrade. If you need to change these ports, do so either before or after you upgrade.

### **Back your files up**

Before you upgrade, back up all of your files, including Splunk Heavy Forwarder configurations, indexed data, and binaries.

- For information on backing up configurations, see [Back up configuration information](#) in the *Admin Manual*.

### **Keep copies of custom certificate authority certificates**

When you upgrade on Windows, the installer overwrites any custom certificate authority (CA) certificates that you have created in %SPLUNK\_HOME%\etc\auth. If you have custom CA files, back them up before you upgrade. After the upgrade, you can restore them into %SPLUNK\_HOME%\etc\auth. After you have restored the certificates, restart Splunk Heavy Forwarder.

### ***Upgrade a heavy forwarder using the GUI installer***

1. Download the new MSI file from the Splunk download page.
2. Double-click the MSI file. The installer runs and attempts to detect the existing version of Splunk Heavy Forwarder installed on the machine. When it locates the older version, it displays a pane that asks you to accept the licensing agreement.
3. Accept the license agreement. The installer then installs the updated Splunk Heavy Forwarder. This method of upgrade retains all parameters from the existing installation. By default, the installer restarts Splunk Heavy Forwarder when the upgrade completes and places a log of the changes made to configuration files during the upgrade in %TEMP%.

### ***Upgrade using the command line***

1. Download the new MSI file from the Splunk download page.
2. Install the software, as described in Install on Windows via the command line.
  - ◆ If Splunk runs as a user other than the Local System user, specify the credentials for the user in your command-line instruction with the LOGON\_USERNAME and LOGON\_PASSWORD flags.
  - ◆ You can use the LAUNCHSPLUNK flag to specify whether Splunk Heavy Forwarder should start up automatically or not when the upgrade finishes, but you cannot change any other settings.
  - ◆ Do not change the network ports (SPLUNKD\_PORT and WEB\_PORT) at this time.
3. Depending on your specification, Splunk Heavy Forwarder might start automatically when you complete the installation.

# Configure your Splunk Cloud Platform Deployment

## Configure IP allow lists using Splunk Web

IP allow lists control which IP addresses on your network have access to specified features in your Splunk Cloud Platform deployment. You can use the IP allow list management page in Splunk Web to add IP subnets to allow lists and manage access to Splunk Cloud Platform features in a self-service manner without assistance from Splunk Support.

Alternatively, you can configure IP allow lists programmatically using the Admin Config Service (ACS) API. For more information, see [Configure IP allow lists for Splunk Cloud Platform](#) in the *Admin Config Service Manual*.

### Requirements

To configure IP allow lists using Splunk Web, you must:

- Have Splunk Cloud Platform version 8.2.2201 or higher.
- Hold a role that has the `edit_ip_allow_list` capability, including inherited roles. The `sc_admin` role has this capability by default.
- Enable token authentication. See [Enable or disable token authentication](#).

SAML users must enable a scripted authentication extension for proper authentication of IP allow list operations in Splunk Web. For more information, see [Configure authentication tokens to interface with your SAML IdP](#).

The Splunk Web UI supports configuring IP allow lists on the primary search head (sh1) or search head cluster (shc1) only. It does not support configuring IP allow lists on additional search heads, including premium search heads. You can configure IP allow lists on additional search heads and premium search heads using the ACS API directly. For more information, see [Configure IP allow lists for Splunk Cloud Platform](#) in the *Admin Config Service Manual*.

### Determine IP allow list use case

Splunk Cloud Platform supports several common IP allow list use cases. In each case, the IP allow list controls access to a particular Splunk Cloud Platform feature, for example Search head API access, HEC access for ingestion, and so on.

IP allow list management supports the following IP allow list use cases:

| Use Case                 | Description                                                                                           |
|--------------------------|-------------------------------------------------------------------------------------------------------|
| Search head API access   | Grants access for customer subnets to Splunk search head api (applies to automated interfaces)        |
| HEC access for ingestion | Allows customer's environment to send HTTP data to Splunk indexers.                                   |
| Indexer ingestion        | Allows subnets that include UF or HF to send data to Splunk indexers.                                 |
| Search head UI access    | Grant explicit access to search head UI in regulated customer environments.                           |
| IDM UI access            | Grant explicit access to IDM UI in regulated customer environments.                                   |
| IDM API                  | Grant access for add-ons that require an API. (Allows add-ons to send data to Splunk Cloud Platform.) |



IP allow list rules apply to the entire Splunk Cloud Platform deployment stack, not just to individual components. For example, any subnet that you add to "Search head API access" allow list will have access to the entire search head tier, including all individual search heads and search head clusters. Likewise, any forwarder whose subnet you add to the "Indexer ingestion" allow list will have access to all indexers.

## Add or remove subnets from IP allow lists

The IP allow list management page lets you add or remove subnets from IP allow lists for specified Splunk Cloud Platform features. You can add or remove one or more IP subnets for multiple different features in a single page update. You must click save for any changes that you make to the page to propagate through the system.

### Add subnets to IP allow lists

To add a subnet to an IP allow list:

1. In Splunk Web, click **Settings > Server settings > IP allow list**.
2. If token authentication is not enabled, click **Go to tokens page** and enable token authentication. Once token authentication is enabled, return to the IP allow list management page and refresh the page.
3. Select the tab of the feature to which you wish to grant access. For example, click the "Search head UI access" tab to grant access to the search head UI.
4. Click **Add IP subnet**.
5. Enter the subnet using CIDR notation. For example 192.0.0.0/24
6. Optionally, click **Add IP subnet** to add more subnets.
7. Click **Save**.

This saves all changes to the IP allow list management page since the last page update, including any subnets that you have added or removed, across all feature tabs.

### Remove subnets from IP allow lists

To delete a subnet from an IP allow list:

1. Select the tab for the feature from which you wish to revoke access.
2. Click **X** to delete the existing subnet.
3. Click **Save**.

This saves all changes to the IP allow list management page since the last page update, including any subnets that you have added or removed, across all feature tabs.

You cannot delete the final IP subnet on the allow list for a feature. This is a safety measure that prevents inadvertently revoking all access to a feature. To delete the final subnet on an IP allow list, you must contact Splunk Support.

Changes can take up to 15 mins or more to propagate through the system.

## Troubleshoot authentication related errors

If your deployment uses SAML authentication, and you receive a "Could not find ACS endpoint" error message when configuring IP allow lists in the UI, make sure you have enabled the appropriate scripted authentication extension for your SAML IdP. A scripted authentication extension is required for SAML users to properly authenticate IP allow list operations in Splunk Web.

For more information, see [Configure authentication extensions to interface with your SAML IdP](#).

For help troubleshooting scripted authentication extensions, see [Troubleshoot problems with authentication extensions](#).

## Configure Dashboards Trusted Domains List

The Dashboards Trusted Domains List is a list of authorized domains and URLs that aid the management of external content. For example, external images without a domain or URL specified in the list will not render in the dashboard. To permit external content, you can add the content's domain or URL to the list. You can turn off the enforcement of the domain list by configuring your `web-features.conf` file.

### Add domains

Add a domain or URL to the Dashboards Trusted Domains List using Splunk Web.

1. In Splunk Web, navigate to **Settings > Server settings > Dashboards Trusted Domains List**.
2. Enter a name. The name is a label for the corresponding domain or URL.
3. Select **Save**. This saves all changes to the Dashboards Trusted Domains List page since the last page update, including any domains or URLs that you have added or removed.

### Remove domains

Remove a domain or URL from the Dashboards Trusted Domains List using Splunk Web.

1. In Splunk Web, navigate to **Settings > Server settings > Dashboards Trusted Domains List**.
2. Select **X** to delete the domain or URL.
3. Select **Save**. This saves all changes to the Dashboards Trusted Domains List page since the last page update, including any domains or URLs that you have added or removed.

## Example of configured `dashboards_trusted_domains` settings

Add authorized domains or URLs to the `web-features.conf` file, instead of the previously used `web.conf` file.

If you want to troubleshoot the Dashboards Trusted Domains List or add to the list directly, you can add authorized domains or URLs to the `[feature:dashboards_csp]` stanza in the `web-features.conf` file. Each setting starts with the syntax `dashboards_trusted_domain.` followed by the domain name.

Domain and URL names can be specific or use an asterisk wildcard. The asterisk wildcard must be the leftmost domain in the domain name system. Asterisk wildcards in the middle or end of a domain name system do not work. For example, the domain name `*.buttercup-games.com` loads content from any subdomain under `buttercup-games.com`. The domain name `www.*.buttercup-games.com` is invalid.

The following is an example of configured `dashboards_trusted_domains` settings.

```
[feature:dashboards_csp]
dashboards_trusted_domain.everything=*.buttercup-games.com,
dashboards_trusted_domain.example=example.buttercup-games.com
```

## Subdomains allowed by default

The Dashboards Trusted Domains List (DTDL) allows select subdomains by default without adding the domains to the DTDL. Additionally, the subdomains do not trigger the content warning modals. The subdomains are part of an internal Splunk software list that is not visible to users.

The following lists the subdomains allowed by default:

- apps.splunk.com
- dev.splunk.com
- docs.flowmill.com
- docs.splunk.com
- help.rigor.com
- help.victorops.com
- lantern.splunk.com
- splunkbase.com
- splunkbase.splunk.com
- splunkui.splunk.com
- splunk.com/download
- splunk.com/products

## External content and redirection feature settings

Do not set the feature settings to false. Turning the feature settings to false removes safeguards for external content and external redirection modals.

Dashboard Studio and Classic SimpleXML dashboards use feature settings in web-features.conf to turn the enforcement of the Dashboards Trusted Domains List on and off.

`Enable_dashboards_external_content_restriction` is true by default and shows the external content warning if a domain or URL is not in the Dashboards Trusted Domains List.

`Enable_dashboards_redirection_restriction` is true by default and shows the redirection warning modal if a domain or URL is not in the Dashboards Trusted Domains List.

## Dashboard Studio dashboards

The warning modals for Dashboard Studio dashboards differ in how they handle external or redirection content. Both modals have configurable feature settings that default to true for enablement.

### ***External content warning modal***

Dashboard Studio dashboards that attempt to load external content not listed in the Trusted Domains List receive an error message and the content doesn't load.

To avoid the error, you can do one of the following:

- Add the domain or URL to the Dashboards Trusted Domains List.
- Upload external content to your app directory and reference the content locally.
- Upload images directly with the Dashboard Studio UI. For more details, see Add an image.

### ***Redirection content warning modal***

Dashboard Studio dashboards that attempt to redirect to external content not listed in the Trusted Domains List receives a warning message confirming that you want to leave the Splunk Platform.

To avoid the warning modal, you can add the domain or URL to the Dashboards Trusted Domains List.

### **Classic SimpleXML dashboards**

The warning modals for Classic SimpleXML dashboards differ in how they handle external or redirection content. Both modals have configurable feature settings that default to true for enablement.

#### ***External content warning modal***

When viewing SimpleXML dashboards that attempt to load external content, a warning modal prompts the following:

- Load content by acknowledging the external domain or URL is trusted.
- Not load content by selecting **Cancel** because the external domain or URL is not trusted.

To avoid the warning modal, you can do one of the following:

- Add the domain or URL to the Dashboards Trusted Domains List.
- Upload external content to your app directory and reference the content locally.

#### **Tags that load external content**

The warning modal checks HTML tags that load external content. The following is a list of HTML tags in SimpleXML that load external content:

- applet
- audio
- base
- embed
- form
- frame
- iframe
- img
- object
- script
- style
- track
- video

### ***Redirection content warning modal***

The redirection content warning modal applies to any links in HTML tags or custom URLs. When viewing Classic SimpleXML dashboards that attempt to redirect to external content, a warning modal prompts the following:

- Redirect to the content by acknowledging the external domain or URL is trusted.
- Not redirect to the content by selecting Cancel because the external domain or URL is not trusted.

## Tags that load external content

The warning modal checks HTML tags that redirect to external content. The following is a list of HTML tags in SimpleXML that redirect to external content:

- a
- link

## Manage custom bookmarks

On the Home page of the Splunk platform, you can create, save, and share bookmarks to in-product locations or external links for quick access to resources such as apps, knowledge objects, reference guides, or documents.

Each user regardless of role can create a maximum of 30 personal bookmarks, which are visible to that user only. Users with the admin role can create a separate pool of a maximum of 30 shared bookmarks, which they can display to all users or just to a subset of roles. The Home page also includes Splunk-recommended bookmarks, which administrators can choose to hide or display to all users.

Bookmarks must be within one of the following domains:

- The Splunk platform instance
- External domains that an administrator has explicitly allowed

## Configure the list of allowed domains

The list of allowed domains for bookmarks includes the Splunk platform instance by default. Administrators can complete the following steps to specify additional allowed domains.

1. From the Home page, select **Home page settings**.
2. Enter a domain name and value for each allowed domain.
3. Select **Save**.

Domain and URL names can be specific or use an asterisk wildcard. The asterisk wildcard must be the leftmost domain in the domain name system. Asterisk wildcards in the middle or end of a domain name system do not work.

In the following chart, each row provides an example domain value, and each column shows if that domain value would permit users to bookmark that column's link:

| Domain value       | splunk.com | docs.splunk.com | lantern.splunk.com |
|--------------------|------------|-----------------|--------------------|
| *.splunk.com       | Yes        | Yes             | Yes                |
| splunk.com*        | No         | No              | No                 |
| splunk.com         | Yes        | Yes             | Yes                |
| docs.splunk.com    | No         | Yes             | No                 |
| lantern.splunk.com | No         | No              | Yes                |

If an administrator later removes an allowed domain that has associated bookmarks, those bookmarks are marked invalid and no longer function, but they are not deleted.

## Configure webhook allow list using Splunk Web

The webhook allow list is a list of URL endpoints to which webhook alert actions in Splunk Cloud Platform are permitted to send HTTP POST requests. Before a triggered alert can send a request to a specified webhook URL, Splunk Cloud Platform checks to ensure that the URL is on the allow list. You can add URLs to the allow list using the webhook allow list page in Splunk Web.

For more information on webhook alert actions, see *Use a webhook alert action* in the *Alerting Manual*.

### Requirements

To configure the webhook allow list using Splunk Web, you must have:

- Splunk Cloud Platform version 8.2.2203 or higher.
- The `sc_admin` role.
- The `edit_webhook_allow_list` capability. `sc_admin` has this capability by default.

### Add or remove URL endpoints from the webhook allow list

The webhook allow list page lets you add or remove target URL endpoints for webhook alert actions. You can add or remove multiple URL endpoints in a single page update. You must click save for any changes that you make to the page to propagate through the system.

#### *Specify URLs using restrictive regular expressions*

Splunk Cloud Platform does a regular expression match against URLs in the allow list. If there is a string match, then an alert (HTTP POST request) is sent to the specified webhook URL. When adding a URL to the webhook allow list, make sure to define the URL as completely as possible to achieve the most restrictive match. For example, the following URLs appear in order from most restrictive to least restrictive:

1. `https://splunk.m.pipedream.net`
2. `pipedream.net`
3. `pipe`

If you send an alert to `http://orange.pipedream.net`, it will be restricted (not match) in the first case. But it will not be restricted in the second case, since the regular expression `pipedream.net` matches.

Similarly if you send an alert to `http://mywebsite.pipeline.com`, it will be restricted in the first and second case. But it will not be restricted in the third case, since the regular expression `pipe` matches. Hence, it is best to use the first URL for a more restrictive policy.

In most cases, it is best to use `https://` as the starting string of the URL.

#### *Add URL endpoints to the webhook allow list*

To add a URL endpoint to the webhook allow list using Splunk Web:

1. In Splunk Web, click **Settings > Server settings > Webhook allow list**.
2. Enter a name for the endpoint. The name is just a label for the corresponding URL. You cannot use the name field in the search and reporting app to send an alert .
3. Specify the endpoint URL value. See [Specify URLs using restrictive regular expressions](#).
4. Click **Save**  
This saves all changes to the webhook allow list page since the last page update, including any URLs that you have added or removed.

### **Remove URL endpoints from the webhook allow list**

1. In Splunk Web, click **Settings > Server settings > Webhook allow list**.
2. Click **X** to delete the URL endpoint.
3. Click **Save**.  
This saves all changes to the webhook allow list page since the last page update, including any URLs that you have added or removed.

### **Check alert failures due to URL not in allow list**

Upon upgrade to version 8.2.2203, Splunk Cloud Platform automatically adds all URLs currently associated with a webhook alert action to the webhook allow list. However, after upgrade to 8.2.2203 or higher, you must manually add any URL associated with a webhook alert action to the webhook allow list, or that alert will fail.

To see which webhook alerts will fail because the webhook URL is missing from the allow list, run the following search:

```
index="_internal" source=*splunkd.log "did not match an entry" URL=* | stats values(URL) by sid
```

## **Configure limits using Splunk Web**

Splunk Cloud Platform supports self-service configuration of select `limits.conf` settings, which can be useful for optimizing search performance. You can use the Configure limits page in Splunk Web to view and edit `limits.conf` settings, without assistance from Splunk Support.

Alternatively, you can configure `limits.conf` settings programmatically using the Admin Config Service (ACS) API. For more information, see [Manage limits.conf configurations in Splunk Cloud Platform in the Admin Config Service Manual](#).

### **Requirements**

To configure `limits.conf` using Splunk Web:

- You must have the `sc_admin` role.
- You must have the `edit_limits_conf` capability. The `sc_admin` role includes this capability by default.
- You must have Splunk Cloud Platform version 9.0.2209.
- Your Splunk Cloud Platform deployment must be on Victoria Experience. See [Determine your Splunk Cloud Platform Experience](#).
- Automatic UI updates and token authentication must be enabled for your deployment.
- Your deployment must have one or more separate search heads or a search head cluster.

The Configure limits UI does not currently support AWS GovCloud or FedRAMP environments.

Changing limits.conf settings can affect the performance of your Splunk Cloud Platform deployment.

## View and edit limits.conf settings

This section shows you how to view and edit select `limits.conf` settings using Splunk Web.

The table shows editable `limits.conf` settings by stanza, with minimum, maximum, and default values:

| Stanza          | Setting                    | Description                                                                                                                                                                           | Values (min/max/default)                                     |
|-----------------|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------|
| [join]          | subsearch_maxout           | The maximum number of result rows to output from subsearch to join against.                                                                                                           | "minValue": 0<br>"maxValue": 100000<br>"defaultValue": 50000 |
|                 | subsearch_maxtime          | Maximum search time, in seconds, before auto-finalization of subsearch.                                                                                                               | "minValue": 0<br>"maxValue": 120<br>"defaultValue": 60       |
| [kv]            | maxchars                   | Truncate <code>_raw</code> to this size and then do auto KV. A value of 0 means that no truncation occurs.                                                                            | "minValue": 1<br>"maxValue": 20480<br>"defaultValue": 10240  |
|                 | limit                      | The maximum number of fields that an automatic key-value field extraction (auto kv) can generate at search time.                                                                      | "minValue": 1<br>"maxValue": 200<br>"defaultValue": 100      |
|                 | maxcols                    | When non-zero, the point at which kv stops creating new fields.                                                                                                                       | "minValue": 256<br>"maxValue": 2048<br>"defaultValue": 512   |
| [pdf]           | max_rows_per_table         | The maximum number of rows that will be rendered for a table within integrated PDF rendering.                                                                                         | "minValue": 500<br>"maxValue": 5000<br>"defaultValue": 1000  |
| [scheduler]     | max_per_result_alerts      | Maximum number of alerts to trigger for each saved search instance (or real-time results preview for RT alerts). Only applies in non-digest mode alerting.                            | "minValue": 250<br>"maxValue": 5000<br>"defaultValue": 500   |
|                 | max_per_result_alerts_time | Maximum amount of time, in seconds, to spend triggering alerts for each saved search instance (or real-time results preview for RT alerts). Only applies in non-digest mode alerting. | "minValue": 150<br>"maxValue": 1800<br>"defaultValue": 300   |
| [searchresults] | maxresultrows              | Maximum number of events generated by search commands                                                                                                                                 | "minValue": 0                                                |



| Stanza      | Setting           | Description                                                                                                                                                                                                                      | Values (min/max/default)                                        |
|-------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
|             |                   |                                                                                                                                                                                                                                  | "maxValue": 1000000<br>"defaultValue": 50000                    |
| [spath]     | extraction_cutoff | For 'extract-all' spath extraction mode, this setting applies extraction only to the first <integer> number of bytes. This setting applies both the auto kv extraction and the spath command, when explicitly extracting fields. | "minValue": 2500<br>"maxValue": 2000000<br>"defaultValue": 5000 |
| [subsearch] | maxout            | Maximum number of results to return from a subsearch.                                                                                                                                                                            | "minValue": 0<br>"maxValue": 10400<br>"defaultValue": 10000     |
|             | maxtime           | Maximum number of seconds to run a subsearch before finalizing                                                                                                                                                                   | "minValue": 0<br>"maxValue": 120<br>"defaultValue": 60          |

All editable limits.conf settings are reloadable.

For more detailed information on each of the supported limits.conf settings, see Limits.conf in the Splunk Enterprise *Admin Manual*.

### **Enable automatic UI updates and token authentication**

Before you can access and use the Configure limits page in Splunk Web, you must enable automatic UI updates and token authentication for your deployment.

To enable automatic UI updates:

1. In Splunk Web, select **Settings > Automatic UI updates**.
2. Set the switch to enable automatic UI updates.
3. Select **Save**.

After you enable automatic UI updates the Configure Limits menu option appears under **Settings > Server settings**.

To enable token authentication:

1. In Splunk Web, select **Settings > Tokens > Token Settings**.
2. Set the **Token Authentication** switch to **Enabled**.

You can also find the **Token Settings** page through the interactive search bar.

## **Configure *limits.conf* settings**

To view, edit, or reset *limits.conf* settings using Splunk Web:

1. Select **Settings > Server settings**.
2. Edit one or more of the available *limits.conf* settings values.
3. Select **Save**. A successful request message means that your edits have been submitted successfully, but setting changes can still take time to propagate.

## **Manage HTTP Event Collector (HEC) tokens in Splunk Cloud Platform**

### **Topic has moved**

This topic has moved to the *Admin Config Service Manual*. See [Manage HTTP Event Collector tokens in Splunk Cloud Platform](#).

If you have bookmarks to this page, update them to use the new page.

# Manage Splunk Cloud Platform Users and Roles

## Manage Splunk Cloud Platform users and roles

### Topic has moved

This topic has moved to the *Securing the Splunk Platform* Manual. See Manage Splunk Cloud users and roles.

If you have bookmarks to this page, update them to use the new page.

## Configure Splunk Cloud to use SAML for authentication tokens

### Topic has moved

This topic has moved to the *Securing the Splunk Platform* Manual. See Configure Splunk Cloud to use SAML for authentication tokens.

If you have bookmarks to this page, update them to use the new page.

# Monitor your Splunk Cloud Platform Deployment

## Introduction to the Cloud Monitoring Console

The Cloud Monitoring Console (CMC) lets Splunk Cloud Platform administrators view information about the status of your Splunk Cloud Platform deployment. CMC dashboards provide insight into how the following areas of your Splunk Cloud Platform deployment are performing:

- Data ingestion and data quality
- Forwarder connections
- HTTP Event Collection tokens
- Indexing
- Indexer clustering and search head clustering, if applicable
- License usage
- Search
- User behavior
- Workload management

The Cloud Monitoring Console does not store or retain any customer data displayed in the dashboards. Customer data remains local to the customer stack.

You must have the `sc_admin` (Splunk Cloud Platform Administrator) role to use the Cloud Monitoring Console.

### Preview disclaimer

Some features in the CMC may contain a **(preview)** label.

Preview features are provided by Splunk to you "as is" without any warranties, maintenance and support, or service level commitments. Splunk makes this preview feature available in its sole discretion and may discontinue it at any time. Use of preview features is subject to the Splunk General Terms.

### Locate the Cloud Monitoring Console

To locate the CMC app in your Splunk Cloud Platform deployment, follow these steps:

1. From anywhere in Splunk Web, select **Apps**.
2. Select **Cloud Monitoring Console**.

On the Apps page that you access through **Apps > Managed Apps**, the CMC is named `splunk_instance_monitoring`.

### Select the correct documentation version for your deployment

Ensure that you are viewing the correct CMC documentation version for your Splunk Cloud Platform deployment.

To determine your Splunk Cloud Platform deployment version, follow these steps:

1. In the CMC app, select **Support & Services > About**. The CURRENT APPLICATION area at the bottom of the About page shows the app's version and build numbers.
2. In this documentation, select the correct version from the **Version** dropdown menu in the upper right corner.

## Set your default time zone

The CMC app displays time-based data in panels, charts, and tables based on the default time zone set for your user profile. To review or reset your current time zone setting, perform the following steps:

1. In the CMC app, select your user profile adjacent to **Support & Services**, then select **Preferences**.
2. In the **Preferences** page, select **Global**.
3. Specify an option for the **Time zone** field and select **Apply**.

## Enable platform alerts

CMC provides alerting functionality with preconfigured platform alerts for missing forwarders and skipped searches that you can enable. If either alert is triggered, CMC displays a notification on the Triggered Alerts page. You can also set up custom alerts on the global Searches, Reports, and Alerts page, which is accessible from the Triggered Alerts page. For more information, see [Use the Alerts panel](#).

## Troubleshoot the CMC dashboards

If you have any issues with a CMC dashboard, try these methods of troubleshooting:

- Check that any necessary setup is properly configured and enabled. For example, you must first configure the **Forwarders Monitoring Setup** page to use the Forwarders dashboards.
- If you have a support contract, log in and file a new case using the Splunk Support Portal. Otherwise, contact Splunk Customer Support.

Do not modify any part of a CMC dashboard. Any local changes that you make might break the CMC application and override its automatic update process.

## Variables in panel titles

A number of panels in CMC have variable titles. You set the variable when you select a specific filter option in the panel. Panels with variable titles are noted in this manual.

## Splunk Cloud Platform documentation site

For more information about features and functionality of Splunk Cloud Platform, see the Splunk Cloud Platform documentation. The manuals hosted on this site provide comprehensive information on how to configure your deployment, search the ingested data, and create dashboards, visualizations, and reports for your data analysis. This site also hosts the CMC release notes, along with release notes for other Splunk Cloud Platform products.

## Use the Overview dashboard

The Cloud Monitoring Console (CMC) Overview dashboard enables Splunk Cloud Platform administrators to quickly understand the general state and health of their deployment.

A blue progress bar may appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify this dashboard. Changing any of the search criteria, formatting, or layouts might cause inaccurate results and also override the automatic update process.

## Review the Overview dashboard

The Overview dashboard displays 12 summary panels of information about the health of your deployment, with each panel linked to its respective source CMC dashboard. Select a panel to view more detailed information about that particular metric.

The Release Notes link near the top of the dashboard accesses the latest version of the CMC release notes in the Splunk Cloud Platform documentation.

The File with Local Overwrites panel displays if your deployment contains modifications to the original delivered app files. The table lists all modified files and the date and time that they were changed. Modifications to any custom deployment-specific file are not considered a local overwrite.

Local overwrites prevent the CMC app from automatically updating. If your deployment contains modifications to the original delivered app files, you must contact Splunk Customer Support to remove the local overwrites and re-enable the automatic update functionality.

To investigate your panels, go to **Cloud Monitoring Console > Overview**. Use the following table to understand the dashboard interface.

| Panel                             | Description                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Current Active Users (Last Hour)  | Shows the number of active users in the deployment as of the last 60 minutes from when you accessed the dashboard. For example, if you access the dashboard at 4:30 PM, this panel shows data from 3:30-4:30 PM.<br><br>This panel accesses the <a href="#">User Activity dashboard</a> .                                                                                                    |
| Average Daily Users (Last 7 Days) | Shows the number of daily users in the deployment averaged over the last seven days from the previous day. For example, if you access the dashboard on June 8, this panel shows data from June 1, 12:00 AM to June 7, 11:59 PM.<br><br>This panel accesses the <a href="#">User Activity dashboard</a> .                                                                                     |
| Search Count (Yesterday)          | The large number shows the number of searches performed during the previous day. For example, if you access the dashboard on June 8, this panel shows data from June 7, 12:00 AM to 11:59 PM. The smaller number and arrow indicates the increase or decrease in searches from the previous search count.<br><br>This panel accesses the <a href="#">Search Usage Statistics dashboard</a> . |
| Indexes with Events               | Shows the number of indexes that have processed events.<br><br>This panel accesses the <a href="#">Indexing Performance dashboard</a> .                                                                                                                                                                                                                                                      |

| Panel                                   | Description                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                         | <b>Note:</b> You must have the <code>indexes_edit</code> capability to view accurate data in this panel.                                                                                                                                                                                                                                                                         |
| Total Indexes                           | <p>Shows a snapshot of the currently active indexes that contain events.</p> <p>This panel accesses the <a href="#">Indexing Performance dashboard</a>.</p> <p><b>Note:</b> You must have the <code>indexes_edit</code> capability to view accurate data in this panel.</p>                                                                                                      |
| Ingest Volume                           | <p>The large number shows the amount of data ingested in gigabytes in the previous day. See <b>Search Count (Yesterday)</b> for an explanation of the time range for the previous day. The smaller number and arrow indicates the increase or decrease in data ingestion from the previous ingestion total.</p> <p>This panel accesses the <a href="#">Ingest dashboard</a>.</p> |
| Searches by Type (Last 24 Hours)        | <p>Shows a color-coded bar graph of searches performed over the last 24 hours. For example, if you access the dashboard on June 2 at 9:00 AM, this panel shows data from June 1, 9:00 AM to June 2, 9:00 AM.</p> <p>This panel accesses the <a href="#">Search Usage Statistics dashboard</a>.</p>                                                                               |
| Throughput by Index (Last 24 Hours)     | <p>Shows a color-coded bar graph of data throughput performance per index over the last 24 hours. See <b>Searches by Type (Last 24 Hours)</b> for an explanation of the 24-hour time range.</p> <p>This panel accesses the <a href="#">Indexing Performance dashboard</a>.</p>                                                                                                   |
| Splunk TCP Port Closures (Last 4 Hours) | <p>Shows the percentage of your active indexers in the last 4 hours that have Splunk TCP port closures. For example, if you access the dashboard at 4:00 PM, this panel shows data from 12:00-4:00 PM.</p> <p>This panel accesses the <a href="#">Indexing Performance dashboard</a>.</p>                                                                                        |
| Long Running Searches (Last 4 Hours)    | <p>Shows the number of ad hoc searches in the last 4 hours that have taken more than 30 minutes to complete. See <b>Splunk TCP Port Closures (Last 4 Hours)</b> for an explanation of the 4-hour time range.</p> <p>This panel accesses the <a href="#">Search Usage Statistics dashboard</a>.</p>                                                                               |
| Scheduled Search Skip Ratio (Last Hour) | <p>Shows the percentage of your scheduled searches that encountered an issue and had to be skipped in the last hour.</p> <p>See <b>Current Active Users (Last Hour)</b> for an explanation of the 1-hour time range.</p> <p>This panel accesses the <a href="#">Skipped Scheduled Searches dashboard</a>, enabling you to resolve the issue and run a skipped search again.</p>  |
| Data Parsing Issues (Last Hour)         | <p>Shows a bar chart of the line breaking, timestamp parsing, and aggregation issues the Splunk platform encountered when parsing your data for indexing. See <b>Current Active Users (Last Hour)</b> for an explanation of the 1-hour time range.</p> <p>This panel accesses the <a href="#">Data Quality dashboard</a>.</p>                                                    |

### ***Interpret these results***

Because the Overview dashboard provides a high-level view of the overall health of your deployment, investigate any anomalous spikes or dips and take the necessary mitigation action. For example, if you see a sudden increase in skipped scheduled searches, audit these searches to determine the cause and correct any issues.

## **Use the Health dashboard**

Review the status of your Splunk Cloud Platform deployment using the **Health** dashboard. This dashboard provides information about the overall health of the deployment and its data collection, indexing, and search performance.

The **Health** dashboard data updates every 3 hours. This allows the health dashboard to load data quickly. See the **Last updated** column to learn when the data was last updated. The health indicator **View details** pages update at page load time. Because of this, there might be a discrepancy between the **Health dashboard** main page and the health indicator details pages.

### **Navigate the Health dashboard**

The **Health** dashboard offers status information and suggested actions for health indicators so you can identify metrics that need updating or optimizing.

#### ***Review the health indicator panels***

Select a health indicator panel to show indicators that affect a particular health category of your Splunk Cloud Platform deployment.

Each list item shows the corresponding indicators, the health check validation criteria, the results of the health check, and the option to configure an alert for it. The individual results data for a specific indicator correlate to the **Conform**, **Warning**, and **Critical** totals that display in the corresponding panel.

Selecting each panel provides details on the following health categories:

- **Overall health:** Provides a combined summary view of your deployment's data collection, data indexing, and data search performance in context of indicators provided in the indicator table.
- **Data collection:** Shows deployment's universal forwarders and heavy forwarders days remaining before expiration.
- **Data indexing:** Shows the current state of bucket size and range per index for your deployment.
- **Data search:** Shows the current state of skipped search percentage, high memory searches, and cache transfer activity in your deployment.

#### ***Review health indicator toggled summary view***

Select the toggle next to an indicator to view a description of what the indicator evaluates, when the indicator is marked as warning or critical, and suggested actions to maintain the health of the indicator.

#### ***Configure alerts for health indicators***

CMC includes preconfigured alerts so you can get informed when health indicators reach certain thresholds. Select **Configure** to enable an alert. The following table describes the threshold for each health indicator preconfigured alert:



| Health indicator                     | Alert trigger                                                                            |
|--------------------------------------|------------------------------------------------------------------------------------------|
| Universal forwarder software version | When your universal forwarder is going to expire within 15 days.                         |
| Heavy forwarder software version     | When your universal forwarder is going to expire within 15 days.                         |
| Bucket size and range                | When your bucket size is not well distributed.                                           |
| Skipped search percentage            | When your skipped search percentage is greater than 25%.                                 |
| Cache transfer activity              | When your SmartStore download size exceeds 10% of total disk space.                      |
| High memory searches                 | When your searches are consuming more than 10% of Splunk Cloud Platform instance memory. |

**Review health indicator details**

In the toggled expanded view, select **View details** for any of the health indicators to view a drilldown of the indicator. Select a status card to filter the list by **Conforming**, **Warning**, or **Critical** status. The status column correlates to the **Conforming**, **Warning**, and **Critical** totals that display in top status cards.

The detailed view for each health indicator displays the following information:

| Health indicator                     | Chart description                                                                                                                                                                                                    | Health indicator importance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Universal forwarder software version | The <b>Universal forwarder software version</b> detailed views show the forwarder names, versions, days to expiration, expiration timestamp, and status.                                                             | <p>This indicator informs you of upcoming expiry dates of your universal forwarder software version so you can maintain version compatibility with Splunk Cloud Platform. Maintaining version compatibility allows you to immediately take advantage of new capabilities and ensures uninterrupted service.</p> <p>The <b>Health</b> dashboard follows the end of full support date listed in the Splunk Support Policy.</p>                                                                                                                                                              |
| Heavy forwarder software version     | The <b>Heavy forwarder software version</b> detailed view shows the heavy forwarder names, versions, days to expiration, expiration timestamp, and status.                                                           | <p>This indicator informs you of upcoming expiry dates of your heavy forwarder so you can maintain version compatibility with Splunk Cloud Platform. Maintaining version compatibility allows you to immediately take advantage of new capabilities and ensures uninterrupted service.</p> <p>The <b>Health</b> dashboard follows the end of full support date listed in the Splunk Support Policy.</p>                                                                                                                                                                                   |
| Bucket size and range                | The <b>Bucket size and range</b> detailed view shows the index, bucket type, caller, quarantined percentage, full percentage, exceeded count, small percentage, small count, total count, and status for each bucket | <p>This indicator evaluates buckets and their size in an index to help you manage optimal bucket sizes. If bucket sizes fall below or exceed the range 375MB to 750MB, your stack might experience degraded performance from excessive cache calls. If buckets frequently exceed their maximum configured sign, it might be due to insufficient indexing capacity.</p> <p>Select a critical status index to view more information about source types and other bucket details. A longer period of time between event time and processing time might indicate a need for optimization.</p> |

| Health indicator                 | Chart description                                                                                                                                                                                                                                                                                                          | Health indicator importance                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Skipped search percentage</b> | The <b>Skipped search percentage</b> detailed view shows the app, saved search, user, skip ratio, percentage skipped, reason, and status.                                                                                                                                                                                  | This indicator evaluates the skipped search ratio of all scheduled searches. A high ratio of skipped scheduled searches might indicate one of the following causes: <ul style="list-style-type: none"> <li>• The number of searches being run exceeds your deployment's capacity.</li> <li>• The searches being run are taking too long or using too large amount of memory or CPU.</li> </ul> |
| <b>Cache transfer activity</b>   | The <b>Cache transfer activity</b> detailed view shows the index, download amount, cache churn percentage, and status.                                                                                                                                                                                                     | This indicator evaluates cache download size per index and informs you when data downloaded from SmartStore exceeds 10% of total disk space. Keeping cache download size below 5% ensures proper use of infrastructure resources and prevents unnecessary cache churn that slows down searches.                                                                                                |
| <b>High memory searches</b>      | The <b>High memory searches</b> detailed view shows the search IDs, memory used, percentage memory used, and status. <b>Note:</b> The High memory searches detailed view returns the first 50,000 searches sorted by critical, warning, then conforming status. This prevents the results from timing out on large stacks. | This indicator evaluates search size and informs you when searches take up a large amount of memory. High memory searches might cause your Splunk Cloud Platform instance to not function if it runs out of memory.                                                                                                                                                                            |

**Get recommended actions based on status**

The **Health** dashboard includes help panels that provide recommended actions and information that can help you proactively update expiring forwarders, improve search queries and search time, and maintain well-distributed buckets.

The following health indicators include help panels:

- **Universal forwarder software version**
- **Heavy forwarder software version**
- **High memory searches**
- **Skipped search percentage**
- **Bucket size and range**
- **Cache transfer activity**

To view a help panel, navigate to the **indicator dropdown > View details > any list item**. The panel provides status information, recommended action, and additional information about maintaining the health of the selected indicator.

**Health indicator information and additional resources**

The following table provides information on each health indicator, what the health indicator informs you of, and additional resources for further knowledge and troubleshooting:

| Health indicator                            | Description                                                                                                                                                                                                                                                                      | Additional resources                                                                                                                                                                                                                                   |
|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Universal forwarder software version</b> | The universal forwarder streams data from your machine to a data receiver, formats data before sending it to your Splunk platform, and lets you monitor data in real time. Maintaining version forwarder version compatibility ensures there is no interruption to your service. | To learn more about the universal forwarder, see About the universal forwarder in the Splunk Universal Forwarder <i>Forwarder Manual</i> .<br><br>For details on upgrading your universal forwarder, see <a href="#">Upgrade your Forwarder</a> in the |

| Health indicator                               | Description                                                                                                                                                                                                                                                                                    | Additional resources                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                | <p>This indicator evaluates the software version for all universal forwarders and informs you of upcoming expiry dates so you can maintain version compatibility.</p>                                                                                                                          | <p><i>Splunk Cloud Platform Admin Manual.</i></p> <p>For more information on forwarder version compatibility, see <a href="#">Monitor forwarder deployments</a> in the <i>Splunk Cloud Platform Admin Manual</i> and Supported forwarder versions in the <i>Splunk Cloud Platform Service Description</i>.</p>                                                                                                                                                                                                                                         |
| <p><b>Heavy forwarder software version</b></p> | <p>A heavy forwarder parses data before forwarding or indexes data locally while forwarding the data to another indexer.</p> <p>This indicator evaluates the software version for all heavy forwarders and informs you of upcoming expiry dates so you can maintain version compatibility.</p> | <p>To learn more about the heavy forwarder, see Heavy and light forwarders in the Splunk Enterprise <i>Forwarding Data</i> manual.</p> <p>For details on upgrading your heavy forwarder, see <a href="#">Upgrade your Forwarder</a> in the <i>Splunk Cloud Platform Admin Manual</i>.</p> <p>For more information on forwarder version compatibility, see <a href="#">Monitor forwarder deployments</a> in the <i>Splunk Cloud Platform Admin Manual</i> and Supported forwarder versions in the <i>Splunk Cloud Platform Service Description</i>.</p> |
| <p><b>High memory searches</b></p>             | <p>High memory searches use a significant amount of your Splunk platform instance memory.</p> <p>This indicator evaluates search size and informs you of when searches take up a high amount of memory.</p>                                                                                    | <p>See the <b>Expensive searches</b> dashboard in the CMC for more details about searches that are using a lot of memory.</p> <p>To learn more about how to troubleshoot high memory, see TLimit search process memory usage in the Splunk Cloud Platform <i>Search Manual</i>.</p> <p>For information on how to write more efficient searches, see Write better searches in the <i>Splunk Cloud Platform Search Manual</i>.</p>                                                                                                                       |
| <p><b>Bucket size and range</b></p>            | <p>An index typically consists of buckets, directories that contain processed external data and index files.</p> <p>This indicator evaluates buckets and their size in an index to help you manage optimal bucket sizes.</p>                                                                   | <p>To learn more about buckets and how they work in an index, see How the indexer stores indexes in the Splunk Enterprise <i>Managing Indexers and Clusters of Indexers</i> manual.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <p><b>Skipped search percentage</b></p>        | <p>Skipped searches occur when the load on your system is higher than the available resources.</p> <p>This indicator evaluates the skipped search ratio of all scheduled searches.</p>                                                                                                         | <p>To learn more about why searches are being skipped and how to avoid skipped searches, see Are you Skipping? Please read!.</p> <p>To investigate skipped scheduled searches, use the <b>Search &gt; Skipped scheduled searches</b> dashboard. See the documentation at</p>                                                                                                                                                                                                                                                                           |

| Health indicator               | Description                                                                                                                                                                                                                                             | Additional resources                                                                                          |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
|                                |                                                                                                                                                                                                                                                         | <a href="#">Investigate skipped scheduled searches</a> to learn how to review this dashboard.                 |
| <b>Cache transfer activity</b> | Cache transfer activity refers to your deployment's local storage, aggregated by bucket size.<br><br>This indicator evaluates bucket download size and informs you of high download size so you can optimize your deployment's cache transfer activity. | To investigate cache transfer activity further, use the <b>Search &gt; Search usage statistics</b> dashboard. |

## Provide feedback

You can provide feedback and ask questions directly from the Health dashboard. Select the **Feedback** button to ask the Splunk community a question or submit an idea for the Splunk team.

## Use the Maintenance dashboard

The **Maintenance** dashboard shows upcoming and past maintenance windows for Splunk initiated changes. Currently, the **Maintenance** dashboard only displays Splunk platform upgrade information. Splunk is working to display other maintenance information in this dashboard.

The Maintenance dashboard does not display any maintenance information for AWS GovCloud Splunk Cloud Platform deployments.

## Prerequisite

The **Maintenance** dashboard requires token authentication to be enabled. See [Enable or disable token authentication](#) to learn how.

Enabling token authentication doesn't impact CMC availability and requires no maintenance or downtime.

If you're using a Splunk Cloud Platform version lower than 9.0.2208 and SAML authentication, the **Maintenance** dashboard requires authentication extensions to be enabled or an identity provider that supports Attribute Query Request. See [Configure Splunk Cloud Platform to use SAML for authentication tokens](#) to learn more.

## Review next maintenance window

If you have a scheduled maintenance window, you'll see a panel with information about your next maintenance window. This panel provides a timestamp, maintenance ID number, and progress timeline.

Select **Configure global banner** to get a sample message you can use to display a global banner that informs users about the next maintenance window.

## Review upcoming and past maintenance windows

You can review upcoming and past maintenance windows in the **Maintenance** dashboard. Filter by **Upcoming**, **Past**, or **All** using the dropdown selection. The default display shows **Upcoming** maintenance windows.

The **Upcoming** filter shows upcoming maintenance windows for the next 30 days. The **Past** filter shows past maintenance windows for the past 180 days.

If there are no maintenance windows in this dashboard, there are no scheduled maintenance windows that are supported by this dashboard. See the following table for more maintenance type information.

The **Maintenance windows** table shows maintenance windows and their following details:

| Field                   | Description                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ID</b>               | Unique ID for the maintenance window.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Maintenance type</b> | Splunk Cloud Service has three classes of maintenance types: <ul style="list-style-type: none"><li>• Service Updates</li><li>• Routine Maintenance</li><li>• Emergency Maintenance</li></ul> See <a href="https://www.splunk.com/en_us/legal/splunk-cloud-service-maintenance-policy.html?locale=en_us">https://www.splunk.com/en_us/legal/splunk-cloud-service-maintenance-policy.html?locale=en_us</a> to learn more. |
| <b>Operation type</b>   | Type of maintenance being performed. Currently, only the <b>Splunk Upgrade</b> operation type is shown. The <b>Splunk Upgrade</b> operation type refers to a Splunk Cloud Platform version upgrade. This operation type belongs to the service updates category of maintenance types.                                                                                                                                   |
| <b>Start time</b>       | Refers to when the maintenance occurred for <b>Past</b> maintenance windows and refers to when the maintenance will occur for <b>Upcoming</b> maintenance windows. The start time is the local time that your Splunk Cloud Platform stack is set to.                                                                                                                                                                    |
| <b>Status</b>           | The current state of the maintenance window. The maintenance window status is updated at around 7:00pm PST/PDT daily, and does not reflect real time status.                                                                                                                                                                                                                                                            |
| <b>Duration</b>         | The length of time the maintenance window will last. This value is an estimate for upcoming maintenance windows and an actual time for past maintenance windows.                                                                                                                                                                                                                                                        |
| <b>Last updated</b>     | When information for this maintenance window was last updated.                                                                                                                                                                                                                                                                                                                                                          |
| <b>Operations</b>       | Tasks required for maintenance.                                                                                                                                                                                                                                                                                                                                                                                         |

Select a maintenance window to view a brief description of the maintenance being performed, operation type, target version, and timeline. The dashboard informs you if the maintenance does not impact any Splunk Cloud Platform services.

The **Operation type** and **Status** values come from the Admin Config Service (ACS) API. To learn more about viewing maintenance windows with the ACS API, see *View maintenance windows in the Splunk Cloud Platform Admin Config Service Manual*.

## Review the upcoming maintenance window operations and timeline

The **Upcoming maintenance** dashboard displays operation details and a progress timeline panel. Select **Operations** to view a description, the operation type, and target version for the operation. Select **Timeline** to view a timeline for the operation.

In CMC version 3.6.0, the timestamps represent the last time the status was updated, which might not be when the actual event occurred.

The timeline provides the following statuses:

| Status            | Description                                                           |
|-------------------|-----------------------------------------------------------------------|
| <b>Tentative</b>  | Splunk has tentatively scheduled a maintenance window.                |
| <b>Scheduled</b>  | Splunk has scheduled the maintenance window.                          |
| <b>InProgress</b> | Splunk is currently performing the maintenance.                       |
| <b>Completed</b>  | Splunk has successfully completed at least one maintenance operation. |
| <b>Canceled</b>   | Splunk or the customer has canceled the maintenance window.           |

### Add an upcoming maintenance window to your calendar

Select the calendar icon next to any scheduled upcoming maintenance window to download a .ics file. To add the event to your calendar, import this file to your calendar application.

The information in the .ics file does not update if any changes are made to the upcoming maintenance window.

### Request a change freeze

In CMC version 3.24.0 and higher, you can request a change freeze up to one year in advance using the **Maintenance** dashboard. A change freeze is a suspension of maintenance for service updates during specific dates or periods in the calendar month.

To learn more, see the Splunk Cloud Platform maintenance policy and the Maintenance section in the Splunk Cloud Platform Service Details.

Change freeze requests don't impact maintenance windows that have already been scheduled. To cancel an existing maintenance window, submit a cancellation request to Splunk Support at least 72 hours before its start time.

#### **Requirements**

Users with the sc\_admin role can submit change freeze requests.

#### **Request a change freeze**

1. In the **Maintenance dashboard**, select the **Change freeze** tab.
2. Select **Request change freeze**.
3. Acknowledge that you've reviewed the Splunk Cloud Platform maintenance policy.
4. Enter the start date, end date, and reason for your change freeze request. Based on the dates you enter, the modal notifies you if a maintenance window is affected by the request.
5. Select whether the change freeze applies to Splunk initiated changes only or both customer and Splunk initiated changes.
6. Select **Request**.

### ***Review change freeze requests***

In the **Maintenance dashboard**, select the **Change freeze** tab to view your change freeze requests. The table displays requests made in the past 1 year.

Select a request to view its details such as scope, category, and impacted maintenance windows.

### ***Edit change freeze request***

1. In the **Maintenance dashboard**, select the **Change freeze** tab.
2. Select the request you want to edit.
3. Select **Edit change freeze**.
4. Edit the request.
5. Select **Update**.

### ***Delete change freeze request***

1. In the **Maintenance dashboard**, select the **Change freeze** tab.
2. Select the request you want to delete.
3. Select **Delete change freeze**.
4. Select **Delete**.

## **Use the Alerts dashboard**

CMC provides preconfigured platform alerts for missing forwarders and skipped searches that you can enable. You can also create custom platform alerts using the global Searches, Reports, and Alerts page accessible through the CMC Alerts functionality.

When a CMC platform alert is triggered, a message alert displays on registered mobile devices that are equipped with a Splunk Mobile app for Splunk Cloud Platform administrators. The alert does not display in **Messages** in the top **Splunk Cloud** bar in Splunk Web.

Splunk Cloud Platform administrators can also review alerts on the Triggered Alerts page of the CMC app and the **Alerts** count column on the Searches, Reports, and Alerts page.

You must be on at least app version 2.1.1 to use the CMC platform alerts functionality. To check the app version, select Support & Services > About. The CURRENT APPLICATION area at the bottom of the About page shows the app's version and build numbers.

### **Review triggered alerts**

To view triggered alerts:

1. In the CMC navigation bar, select **Alerts > Triggered Alerts**.
2. The page displays the name of any triggered alert and a timestamp of when it was triggered.

When a preconfigured alert is triggered, CMC displays an alert with a **3** severity level on the **Triggered Alerts** page, which indicates medium severity.

Starting with CMC 2.6.0, preconfigured alerts use the prefix CMC. Alerts with the prefix SIM are retained for backwards compatibility.

The table describes the situations that trigger a preconfigured alert and the CMC dashboards to review to take further action.

| Preconfigured alert                                                | Description                                                                                                                                                                                                                                                                                                                            | Dashboards                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CMC Alert - 503 errors                                             | Triggers when the server returns a 503 error from attempting to process an HTTP Event Collector (HEC) request.                                                                                                                                                                                                                         | See <a href="#">Review the HTTP Event Collector dashboard</a> to view your HEC functionality status.                                                                                                                                                                                 |
| CMC Alert - Bucket size and range                                  | Triggers when an index meets any of the following critical thresholds: <ul style="list-style-type: none"> <li>• 10% of that index's buckets are quarantined</li> <li>• More than 50% of buckets on an index are less than half the max size of that bucket</li> <li>• Less than 50% of buckets have reached their full size</li> </ul> | See <a href="#">Use the Health dashboard</a> to learn more about bucket size and range health.                                                                                                                                                                                       |
| CMC Alert - Cache activity transfer                                | Triggers when bucket download size is higher than 10% of total disk space on all indexers. This is the critical threshold for cache activity transfer.                                                                                                                                                                                 | See <a href="#">Use the Health dashboard</a> to learn more about maintaining bucket size download rate.                                                                                                                                                                              |
| CMC Alert - Heavy forwarder software version                       | Triggers when less than 15 days are remaining before end of support for the Heavy forwarder.                                                                                                                                                                                                                                           | See <a href="#">Use the Health dashboard</a> to learn more about maintaining your Heavy forwarder software version.                                                                                                                                                                  |
| CMC Alert - High memory searches                                   | Triggers when a search size uses more than 10% of your Splunk platform instance memory. This is the critical threshold for search memory usage.                                                                                                                                                                                        | See <a href="#">Use the Health dashboard</a> to learn more about optimizing searches.                                                                                                                                                                                                |
| CMC Alert - Indexers blocked queues                                | Triggers when 50% or more of stack indexers are blocked from processing.                                                                                                                                                                                                                                                               | See <a href="#">Review the Indexing Performance dashboard</a> to investigate blocked indexer queues.                                                                                                                                                                                 |
| CMC-Alert - Ingest volume exceeds 80% of entitlement value         | Triggers when your ingest volume exceeds 80%.                                                                                                                                                                                                                                                                                          | See <a href="#">Monitor current usage of your ingestion-based subscription</a> to learn more about monitoring your ingest volume.                                                                                                                                                    |
| CMC Alert - New Data in Index Specified as "lastchanceindex"       | Runs at 12 minutes past midnight every day and is triggered if there is new data in the index specified as the <code>lastchanceindex</code> in the last 24 hours.                                                                                                                                                                      | See the following: <ul style="list-style-type: none"> <li>• <a href="#">Manage indexes in the Splunk Cloud Platform Admin Config Service (ACS) API endpoint reference</a></li> <li>• <code>lastChanceIndex</code> definition in the <i>Splunk Enterprise Admin Manual</i></li> </ul> |
| CMC Alert - S3 scanned volume exceeds 80% of the entitlement value | Triggers when your Federated Search for Amazon S3 data scan entitlement usage exceeds 80%                                                                                                                                                                                                                                              | See <a href="#">Documentation:SplunkCloud:Admin:MonitoringLicenseUsage</a> to learn more about monitoring your federated search for Amazon S3 data scan entitlement.                                                                                                                 |
| CMC Alert - Skipped search percentage                              | Triggers when a search head has a skip search ratio higher than 25%.                                                                                                                                                                                                                                                                   | See <a href="#">Use the Health dashboard</a> to learn more about lowering your skip search ratio.                                                                                                                                                                                    |
| CMC Alert - Storage Capacity Exceeds 80%                           | Runs at 4:16 AM every day and is triggered if the searchable storage                                                                                                                                                                                                                                                                   | See the table in <a href="#">Review the Searchable Storage (DDAS) dashboard</a> , especially the Searchable Storage Usage Percent panel description.                                                                                                                                 |



| Preconfigured alert                                 | Description                                                                                                                                       | Dashboards                                                                                                                                                          |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                     | usage percent value for your deployment exceeds 80%.                                                                                              |                                                                                                                                                                     |
| CMC Alert - SVC Utilization Exceeds 80% for 3 Hours | Runs every hour at 12 minutes past the hour and is triggered if the SVC utilization value for your deployment exceeds 80% over a 3-hour timespan. | See the table in <a href="#">Review the Workload dashboard</a> , especially the SVC Usage panel description.                                                        |
| CMC Alert - Universal forwarder software version    | Triggers when less than 15 days are remaining before end of support for the Universal forwarder forwarder.                                        | See <a href="#">Use the Health dashboard</a> to learn more about maintaining your Universal forwarder software version.                                             |
| SIM Alerts - Missing Forwarders                     | Runs every 15 minutes and is triggered if there are any forwarders with a status of <b>Missing</b> .                                              | See the <a href="#">Forwarders: Deployment dashboard</a> , especially the Missing Forwarder Alerts and Status and Configuration - As of <current_timestamp> panels. |
| SIM Alerts - Skipped Searches                       | Runs every 60 minutes and is triggered if the number of skipped searches exceeds 20%.                                                             | See the <a href="#">Skipped Scheduled Searches dashboard</a> .                                                                                                      |

## Review preconfigured alerts

In the CMC navigation bar, select **Alerts > Configured Alerts**. The table displays the preconfigured CMC alerts and any custom alerts that you or another Splunk Cloud Platform administrator configured for your organization's deployment. **Last Updated** shows when an alert was edited.

Select the **Enabled** toggle to enable or disable an alert.

Select the **Mobile Alert** toggle to enable or disable an alert on mobile devices. Enabling an alert automatically enables it for display for Splunk Cloud Platform administrators on Splunk Web and registered mobile devices equipped with a Splunk Mobile app. For more information on downloading and registering a Splunk Mobile app, see the following:

- Download Splunk Mobile for iOS
- Download Splunk Mobile for Android
- Log in to a Splunk platform instance in a Connected Experiences app

Select **Edit** to access the Searches, Reports, and Alerts page. You can view detailed information about an alert and perform specific actions, such as reviewing the alert definition and running the alert.

Do not edit the search field for preconfigured alerts.

## Manage CMC Alerts on the Searches, Reports, and Alerts page

To manage CMC platform alerts on the Searches, Reports, and Alerts page, follow these steps:

1. Access this page through one of the following methods:

- Select the **Edit** link adjacent to an alert on the **Alerts > Configured Alerts** page in the CMC app.
- In the **Splunk Cloud** bar at the top of the page, select **Settings**. In the **KNOWLEDGE** section, select **Searches, reports, and alerts**.

- Set **Type** to Alerts.
- Set **App** to Cloud Monitoring Console (splunk\_instance\_monitoring).

- Set **Owner** to All or Nobody. The CMC and SIM alerts for CMC appear.
- In the **Actions** column, select **Edit > Enable**.

## Create custom alerts

You can also create custom platform alerts using the Searches, Reports, and Alerts page. You can access this page through one of the two methods noted in step one of [Manage CMC Alerts on the Searches, reports, and alerts page](#). Select the New Alert button to define an alert and the corresponding action to be performed when the alert is triggered. For example, you can send an email to the email account in a Splunk Cloud Platform administrator's profile, or an alert to their registered mobile device equipped with a Splunk Mobile app.

For more information, see the following:

- Set up alert actions in the *Alerting Manual*
- The global Alert Actions page. To access this page, in the **Splunk Cloud** bar at the top of the page, select **Settings**. In the **KNOWLEDGE** section, select **Alert actions**.

## Use the Indexing dashboards

Data that you send to Splunk Cloud Platform is stored in indexes. Managing your indexes and their data is important to ensuring the speed and quality of your search results and the accuracy of your data insights.

The dashboards accessed from the **Cloud Monitoring Console > Indexing** tab let you to administer the following indexing and related functionality in your deployment:

- Thoroughly review your indexes, including their performance, current data consumption, and remaining storage capacity, and events and indexing rate of an individual index.
- Manage the quality of your data and correct parsing errors encountered during the conversion process.
- Monitor the progress of HTTP Event Collection tokens within your deployment, if you enabled this functionality.

You can self-manage your Splunk Cloud Platform index settings. See [The Indexes page](#) in the Splunk Cloud Platform *Admin Manual*.

A blue progress bar might appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify any Cloud Monitoring Console (CMC) dashboard. Changing any of the search criteria, formatting, or layouts may cause inaccurate results and also override the automatic update process.

## Check indexing performance

The CMC Indexing Performance dashboard provides information to Splunk Cloud Platform administrators on incoming data consumption. Use this dashboard to analyze the thrupt rate of your indexers and determine if the rate needs to be optimized.

## Review the Indexing Performance dashboard

This dashboard contains four panels. The **Time Range** in the **Historical Charts** area controls the date range of the data displayed in the bottom three panels.

To investigate your panels, go to **Cloud Monitoring Console > Indexing > Indexing Performance**. Use the following table to understand the dashboard interface.

| Panel or Filter                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Indexing thrupt                         | Shows the speed of the indexing rate in KB per second for all of your indexers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Historical Data                         | <p>This area includes the three panels shown under this section.</p> <p>Set a <b>Time Range</b> value to refresh the data in these panels.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Estimated Indexing Rate                 | <p>Provides a bar chart of the estimated indexing rate over time, based on KB ingested per second.</p> <p>You can split by index, source, or source type, or view the total of all these inputs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <variable> Queue Fill Ratio             | <p>The title of this panel is dynamic and depends on the specified <b>Aggregation</b> value, which can be one of the following:</p> <ul style="list-style-type: none"> <li>• Median</li> <li>• Maximum</li> <li>• Minimum</li> <li>• 90th Percentile</li> <li>• Sampled</li> </ul> <p>After you select an <b>Aggregation</b> value, select a <b>Queue</b> value to view the latency performance of each queue in the graph. Queue options are the following:</p> <ul style="list-style-type: none"> <li>• Splunk Tcpin Queue</li> <li>• Parsing Queue</li> <li>• Aggregation Queue</li> <li>• Typing Queue</li> <li>• Indexing Queue</li> </ul> <p>Comparing the queues against one another shows you which queue has the lowest latency and is hindering indexing performance. Note that latency performance is also known as fill percentage over time.</p> |
| Splunk TCP Port Closures                | Shows the percentage of indexers that have closed their forwarder connection port at least once in the specified time range. <b>Note:</b> A high percentage value could indicate that the ingest pipeline is overwhelmed or misconfigured, and data is not being ingested. Contact Splunk Support to resolve this issue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Indexers - Blocked Queues by Queue Type | <p>Shows indexer queues that are blocked from processing, categorized by queue type. Indexers with many blocked queues and no restarts may indicate the following:</p> <ul style="list-style-type: none"> <li>• High CPU</li> <li>• Bad data distribution</li> <li>• Bad data quality</li> </ul> <p>The <b>Time Span</b> field in this panel works in conjunction with the <b>Time Range</b> selector in the <b>Historical Data</b> panel. Select a <b>Time Range</b> value for the chart's x-axis, then select a <b>Time Span</b> value to group data within time increments. For example, you could set a time range of 60 minutes with a time span of 5 minutes.</p>                                                                                                                                                                                       |

## Interpret indexing performance results

When interpreting your indexing performance results, note the following:

- Regularly review your indexing performance and ensure that on average it is adequately handling the load. Though occasional spikes are normal, a consistently high load degrades performance.
- Check for these issues:
  - ◆ An indexing rate lower than expected. An example of this is an indexing rate of 0 with a forwarder outgoing rate of 100.
  - ◆ A TCP port closure percentage value that is high. This percentage indicates an ingestion pipeline issue and indicates that data is potentially being lost.
  - ◆ Source types that are sending a larger volume than expected.
  - ◆ Spikes in blocked queues at regular intervals, specific times, or both. Investigate why the queues become blocked so you can remediate the underlying issue.

## Check index detail

The CMC Index Detail dashboard provides Splunk Cloud Platform administrators with a more granular view about the events in and performance of a specific index. Use this dashboard to more thoroughly investigate individual indexes.

### Review the Index Detail dashboard

This dashboard shows six panels of information for a specified index.

To investigate your panels, go to **Cloud Monitoring Console > Indexing > Index Detail**. Use the following table to understand the dashboard interface.

To view this dashboard, you must have the `indexes_edit` capability.

| Panel or Filter                | Description                                                                                                                                                                                                                                                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index                          | The selected index value affects all panels in this dashboard.<br>The indexes available to you are based on user access levels.                                                                                                                                                                               |
| Overview                       | Shows the uncompressed raw data size and total bucket count of the specified index.                                                                                                                                                                                                                           |
| Events                         | Shows the total number of events in the specified index, and the timestamps of the earliest and latest events.                                                                                                                                                                                                |
| Throughput: Last 24 hours (GB) | Shows the speed of the indexing rate in KB per second for the specified index over the last 24 hours.<br>You can split by host, source, or source type. This value is the y-axis in the graph.<br>If the <b>Undefined Host</b> value appears, see the <a href="#">Interpret index detail results</a> section. |

### Interpret index detail results

Use the Index Detail dashboard to monitor the flow of data into the system by index. If there is an issue that affects one or more indexes, analyzing the metadata for each affected index can help you diagnose the underlying issue.

The value **Undefined Host** appears in the **Throughput: Last 24 Hours (GB)** chart when the CMC app encounters an index configuration issue and can't correctly parse the data. This issue generally indicates that the index host name is either not configured or incorrectly configured for a forwarder. For information about configuring the host for a forwarder, see the entry for `hostname` or `host` in [Forward data from files and directories to Splunk Cloud Platform](#).

## Check the status of HTTP event collection

The CMC HTTP Event Collector dashboard provides the status of your Splunk HTTP Event Collection (HEC) functionality to Splunk Cloud Platform administrators, if you use HEC tokens to securely transmit event and application data. Use this dashboard to view summarized and detailed information about your HEC token usage and performance.

See also Set up and use HTTP Event Collector in the Splunk Cloud Platform *Getting Data In* manual.

### Review the HTTP Event Collector dashboard

This dashboard contains a number of panels about your HEC token data.

Panels are grouped into one of three views, with a fourth view that combines the other three views so you can see all the data concurrently. You can also opt to see all your HEC token data in the results, or specify a particular token for analysis.

The **Historical Data** view contains two graphs with a variable in the panel title that you set with a filter option: **<variable> Count** and **Data <variable>**.

For a HEC token to display in this dashboard, it must meet either of the following conditions:

- Be enabled and have received data within the last 7 days.
- Be recently disabled but have received messages within the last 7 days, prior to being disabled.

To investigate your views, go to **Cloud Monitoring Console > Indexing > HTTP Event Collector**. Use the following table to understand the dashboard interface.

| View or Filter  | Description                                                                                                                                                                                                                                                                                                    |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HEC Token       | Specify an option to see data for all HEC tokens or one specific token.<br>See the information in the previous section as to valid tokens that display in this dashboard.                                                                                                                                      |
| Select View     | Select <b>Usage</b> , <b>Current Thruput</b> , or <b>Historical Data</b> to see a specific view of the data, or select <b>All</b> to see a combined view.                                                                                                                                                      |
| Usage           | The <b>HTTP Event Token Usage (Last 7 Days)</b> panel shows a table that lists the token name, all hosts associated with the token, trend line, and count.                                                                                                                                                     |
| Current Thruput | The <b>Current Thruput</b> panel shows information on the Thruput of your requests and data, per second.<br>The <b>Activity (Last 30 Minutes)</b> graph shows the count of requests and data received (MB) over time.                                                                                          |
| Historical Data | Set the time range for the historical data display.<br>The <b>Request Overview</b> panel shows the event count, valid request count, and invalid request count. This panel is associated with the <b>&lt;variable&gt; Count</b> graph. The title variable depends on the selected <b>Activity Type</b> option. |

| View or Filter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                | <p>The <b>Split by Token</b> checkbox displays only for Events and Valid Requests options.</p> <p>The <b>Data Overview</b> panel shows the total MB received and indexed. This panel is associated with the <b>Data &lt;variable&gt;</b> graph. The title variable depends on the selected <b>Data Type</b> option. The <b>Split by Token</b> checkbox displays only for the Indexed and Valid Received options.</p> <p>The <b>Errors</b> graph shows the count of all or only specific token errors over time. Select an error type from the <b>Reason</b> filter. The <b>Split by Token</b> checkbox displays when you select one of the following error type options:</p> <ul style="list-style-type: none"> <li>• Authentication errors</li> <li>• Requests to disable token</li> <li>• Requests to incorrect URL</li> <li>• Parser errors</li> <li>• 503 errors</li> </ul> <p>The <b>Data received indexed</b> panel shows the amount of data received and indexed by the HTTP event collector. The title variable depends on the selected <b>Data Type</b> option.</p> <p>The <b>Data delay</b> panel shows the seconds between the time that the event was seen by the thrupt processor in the indexing queue, and the time when the event occurred. Select a statistic to show the max or average time difference between the current time and the perceived time of the events coming through the thrupt processor.</p> |

### Interpret HTTP event collection results

When interpreting your HTTP event collection results, note the following:

- Use the **Errors** panel in the **Historical Data** view to identify HEC token processing issues that you must resolve, such as authentication failures, parser errors, and invalid requests.
- A **Data Received** value that is greater than the **Data Indexed** value indicates that Splunk couldn't process the received messages. This generally occurs because of parsing issues, such as missing timestamps. You can check these values in the **Current Throughput** and **Historical Data** views.

See also Detecting scaling problems in the Splunk Cloud Platform *Getting Data In* manual.

### Verify data quality

The CMC Data Quality dashboard provides information to Splunk Cloud Platform administrators on issues that prevented the Splunk platform from correctly parsing your incoming data. Use this dashboard to analyze and resolve common issues that happen during the ingestion process.

Your data quality can have a great impact on both your system performance and your ability to achieve accurate results from your queries. If your data quality is degraded enough, it can slow down search performance and cause inaccurate search results. Be sure to regularly check and repair any data quality issues before they become a problem.

Generally, data quality issues fall under three main categories:

- Line breaks: When there are problems with line breaks, the ability to parse your data into the correct separate events that it uses for searching is affected.

- Timestamp parsing: When there are timestamp parsing issues, the ability to determine the correct time stamp to use for the event is affected.
- Aggregation: When there are problems with aggregation, the ability to break out fields correctly is affected.

**Review the Data Quality dashboard**

The tables in this dashboard list the issues Splunk Cloud Platform encountered when processing your events at both the source type and source levels. To help you better identify which of your data sources have quality issues, you can opt to exclude Splunk source types in the results.

This dashboard contains one panel with a variable in the title: **Issues by source type <variable> by source**.

To investigate your panels, go to **Cloud Monitoring Console > Indexing > Data Quality**. Use the following table to understand the dashboard interface.

| Panel or Filter                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Range                                 | Set the time range for the data display.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Include Splunk Source Types                | Specify whether to include or exclude Splunk source types from the results. Choose <b>No</b> to exclude Splunk source types and filter the results to only your source types.                                                                                                                                                                                                                                                                                                                                                                                    |
| Event Processing Issues by Source Types    | <p>The results table lists the following information:</p> <ul style="list-style-type: none"> <li>• Sourcetype: Select to open the Issues by source type &lt;variable&gt; by source panel.</li> <li>• Total issues</li> <li>• Source count: Total number of individual sources contained in the source type.</li> <li>• Line breaking, timestamp parsing, and aggregation issues</li> </ul> <p>When any cell shows a number greater than 0, select the cell to view the underlying search and related information. This data will help you resolve the issue.</p> |
| Issues by source type <variable> by source | <p>The &lt;variable&gt; value depends on the selected sourcetype. The results table lists the following information:</p> <ul style="list-style-type: none"> <li>• Source: Select any source to open its related <b>Event Line Count</b>, <b>Event Size</b>, and <b>Event Time Disparity</b> panels.</li> <li>• Total issues</li> <li>• Line breaking, timestamp parsing, and aggregation issues</li> </ul>                                                                                                                                                       |

**Interpret data quality results**

This section discusses how to check the quality of your data and how to repair issues you may encounter. However, the concept of data quality depends on what factors you use to judge quality. For the purposes of this section, data quality means that the data is correctly parsed.

**Guidelines**

Finding and repairing data quality issues is unique to each environment. However, using the following guidelines can help you address your data quality:

- It's a good idea to check your most important data sources first. Often, you can have the most impact by making a few changes to a critical data source.
- Data quality issues may generate hundreds or thousands of errors due to one root cause. Sort by volume and work on repairing the source that generates the largest volume of errors first.
- Repairing data quality issues is an iterative process. Repair your most critical data sources first, and then run queries against the source again to see what problems remain.

- For your most critical source, resolve all data quality issues. This helps to ensure that your searches are effective and your performance is optimal.
- Run these checks on a regular cadence to keep your system healthy.

For more information, see [Resolve data quality issues in the Splunk Cloud Platform \*Getting Data In\* manual](#).

### Example

The following example shows the process of resolving a common data quality issue using information from the CMC Data Quality dashboard, specifically, resolving timestamp parsing issues in a source. The steps to resolve your particular data quality issues may differ, but you can use this example as a general template for resolving data quality issues.

1. In the Data Quality dashboard, view the **Event Processing Issues by Source Type** panel. For this example, you are most concerned with timestamp errors in the syslog source, so you need to drill down into that source.



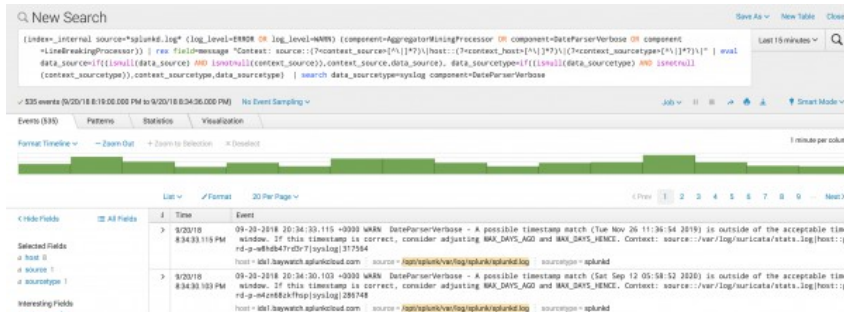
| Source type           | Total Issues | Source Count | Line Breaking Issues | Timestamp Parsing Issues | Aggregation Issues |
|-----------------------|--------------|--------------|----------------------|--------------------------|--------------------|
| splunk_python         | 3385         | 1            | 0                    | 3385                     | 0                  |
| splunkd               | 1949         | 11           | 1482                 | 97                       | 0                  |
| restapi               | 1471         | 1            | 0                    | 4                        | 1466               |
| ps                    | 1325         | 1            | 0                    | 1081                     | 234                |
| osint                 | 545          | 2            | 0                    | 545                      | 0                  |
| tcp                   | 454          | 1            | 0                    | 70                       | 384                |
| suricata              | 410          | 4            | 0                    | 410                      | 0                  |
| package               | 192          | 1            | 0                    | 19                       | 173                |
| splunk_bundl          | 135          | 2            | 0                    | 135                      | 0                  |
| linux_messages_syslog | 65           | 2            | 0                    | 65                       | 0                  |

2. Drilling down, you can see that the majority of issues are with the following source: `/var/log/suricata/stats.log`.



| Source                      | Total Issues | Line Breaking Issues | Timestamp Parsing Issues | Aggregation Issues |
|-----------------------------|--------------|----------------------|--------------------------|--------------------|
| /var/log/suricata/stats.log | 511          | 0                    | 511                      | 0                  |
| /var/log/auth/auth.log      | 3            | 0                    | 3                        | 0                  |
| /var/log/auth.log           | 2            | 0                    | 2                        | 0                  |

3. Select the source to drill down further and see the searches against this source.



```

index=_internal source="*syslog*.log" (log_level=ERROR OR log_level=WARN) (component=aggregator@indexingProcessor OR component=DataParser@Verbose OR component=LineBreakingProcessor) | rex field=message "Context: source:(?<context_source%{1})*%{1}host:(?<context_host%{1})*%{1}((?<context_sourcetype%{1})*%{1})" | eval data_source=if(!isnull(data_source) AND isnull(context_source),context_source,data_source), data_sourcetype=if(!isnull(data_sourcetype) AND isnull(context_sourcetype),context_sourcetype,data_sourcetype) | search data_sourcetype=syslog component=DataParser@Verbose
  
```

| Time                       | Event                                                                                                                                                                                                                                                                                         |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 9/23/18 11:34:53.115 +0000 | WARN DataParser@Verbose - A possible timestamp match (Tue Nov 26 11:34:54 2018) is outside of the acceptable time window. If this timestamp is correct, consider adjusting MAX_DAYS_AGO and MAX_DAYS_HENCE. Context: source::/var/log/suricata/stats.log host:p rd-p-wkhd47r2d7fjyslog 177584 |
| 9/23/18 12:05:54.523 +0000 | WARN DataParser@Verbose - A possible timestamp match (Sat Sep 12 05:54:52 2020) is outside of the acceptable time window. If this timestamp is correct, consider adjusting MAX_DAYS_AGO and MAX_DAYS_HENCE. Context: source::/var/log/suricata/stats.log host:p rd-p-wkhd47r2d7fjyslog 181748 |



4. From here, you can look at a specific event. You can see that the issue is that the Splunk platform was unable to parse the timestamp in the **MAX\_TIMESTAMP\_LOOKAHEAD** field.

```
> 9/20/18 09-20-2018 20:34:25.005 +0000 WARN DateParserVerbose - Failed to parse timestamp in first MAX_TIMESTAMP_LOOKAHEAD (100) characters of event. D
8:34:28.005 PM efaulting to timestamp of previous event (Thu Sep 20 20:34:25 2018). Context: source:///var/log/suricata/stats.log|host::dev-rainmakdev-q-mas9
%k&ku2|syslog|106544
host = id3.baywatch.splunkcloud.com | source = /opt/splunk/var/log/splunk/splunkd.log | sourcetype = splunkd
```

5. To fix this, go to **Settings** in the search bar and select **Source types** in the **DATA** section.
6. In the filter, enter **syslog** for the source type.
7. Select **Actions > Edit**. The **Edit Source Type** page opens.
8. Select **Timestamp > Advanced...** to open the **Timestamp** page for editing. Ensure you are satisfied with the timestamp format and the **Lookahead** settings. In this case, you need to edit the **Lookahead** settings so that the Splunk platform can parse the timestamp correctly.

Timestamp

Extraction: Auto Current time Advanced...

Time zone: Auto

Timestamp format:   
A string in strftime() format that helps Splunk recognize timestamps. [Learn More](#)

Timestamp prefix:   
Timestamp is always prefaced by a regex pattern eg: 'd+abc123\d{2,4}'

Lookahead:   
Timestamp never extends more than this number of characters into the event, or past the Regex if specified above.

9. Return to the main **Edit Source Type** page and go to the **Advanced** menu. From here you can make other changes if needed.

Advanced

| Name                    | Value                             |
|-------------------------|-----------------------------------|
| CHARSET                 | UTF-8                             |
| ADD_EXTRA_TIME_FIELDS   | True                              |
| ANNOTATE_PUNCT          | true                              |
| AUTO_KV_JSON            | true                              |
| BREAK_ONLY_BEFORE_DATE  | true                              |
| DEPTH_LIMIT             | 1000                              |
| LEARN_MODEL             | true                              |
| LEARN_SOURCETYPE        | true                              |
| LINE_BREAKER_LOOKBEHIND | 100                               |
| LOOKUP-inventory_host   | aws_inventory FQDN AS host OUTPUT |
| LOOKUP-zaws_accounts    | aws_accounts aws_account_id OUTPL |
| MATCH_LIMIT             | 100000                            |
| MAX_DAYS_AGO            | 10951                             |
| MAX_DAYS_HENCE          | 2                                 |

## Use the Search dashboards

Healthy search loads are critical to the performance of your entire Splunk Cloud Platform environment. Understanding search patterns can help you understand if your search workload is aligned with best practices and optimized for the best performance. The dashboards accessed from the **Cloud Monitoring Console > Search** tab help you determine if a specific user, search, dashboard, or app is inhibiting your performance. If you encounter an issue, you can then work with users to improve performance.

When searches run for a long time, they may use too much computation and memory, causing an overall slowness of the Splunk instance. This commonly occurs when a few poorly formed searches are taking a large amount of resources. It can also occur if you have a dashboard that is being frequently used by multiple users concurrently. In each of these cases, investigating further can help you to pinpoint the searches that are long-running and determine if you can optimize them.

Because each company's environment is different, it's not easy to set benchmarks for search performance. Generally, the best way to understand your search performance is to compare your historical search times with your current search times to see if there is a change. If search runtimes have slowed, review changes to your environment and new searches to determine if you need to optimize your searches or environment. For example, you may have added a poorly formed search, or you may have added a dashboard that has attracted a lot of traffic.

For more information about optimizing searches, see About search optimization in the Splunk Cloud Platform *Search Manual*.

A blue progress bar might appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify any Cloud Monitoring Console (CMC) dashboard. Changing any of the search criteria, formatting, or layouts may cause inaccurate results and also override the automatic update process.

## Analyze search usage statistics

The CMC Search Usage Statistics dashboard provides information about your searches on a per-instance basis. You can split the data by user, host, and source type to better review the results. Analyzing this data helps you answer questions like these:

- Which users run a large number of searches, or long-running searches, or both?
- Which hosts handle the greatest number of searches, and what is the median runtime?
- What are the most heavily used search types in the deployment?

You can then examine searches in more detail to determine if they can be optimized.

See also:

- Write better searches in the Splunk Cloud Platform *Search Manual*
- Configure the priority of scheduled reports in the Splunk Cloud Platform *Reporting Manual*

## Review the Search Usage Statistics dashboard

This dashboard includes panels of summary, graphical, and tabular data about your search usage statistics. The filter lets you specify a time range and instance values and opt to include or exclude ad hoc searches.

This dashboard contains one panel with a variable in the title: **Search Activity by <variable>**.

To investigate your panels, go to **Cloud Monitoring Console > Search > Search Usage Statistics**. Use the following table to understand the dashboard interface.

| Panel or Filter               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instances                     | Choose to view all instances or specify a particular instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Time Range                    | Set the time range for the data display.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Only Ad Hoc Searches          | Choose <b>Yes</b> to limit the displayed data to only ad hoc searches. Choose <b>No</b> to view results for both ad hoc and scheduled searches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Overview                      | Shows the total number of searches finished, successfully and unsuccessfully completed, and the total search runtime of all searches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Searches                      | <p>Provides a graphical representation of the information in the <b>Overview</b> panel. Specify a <b>Split by</b> option to view the related Searches graph by user, host, or search type. Search types displayed grouped in the following general and granular categories:</p> <ul style="list-style-type: none"> <li>• <b>acceleration and summarization</b>: See Overview of summary-based search acceleration.</li> <li>• <b>ad hoc</b>: See Ad hoc search.</li> <li>• <b>dashboard</b>: Searches that populate a dashboard. Search types in this granular category are <b>ad hoc</b>, <b>acceleration</b>, <b>other</b>, or <b>scheduled</b>.</li> <li>• <b>other</b>: Searches that don't fall into the other categories. These searches are generally performed by a Splunk Cloud Platform administrator.</li> <li>• <b>realtime</b>: This is a granular category for <b>scheduled</b> searches that are continually running in the background.</li> <li>• <b>scheduled</b>: See Scheduled search. This category excludes <b>realtime</b> searches.</li> <li>• <b>subsearch</b>: See Subsearch. Search types in this granular category are <b>ad hoc</b> and <b>other</b>.</li> </ul>                                                                                                                                                                                               |
| Search Activity by <variable> | <p>The &lt;variable&gt; is determined by <b>Split by</b> choice. The table lists the following information:</p> <ul style="list-style-type: none"> <li>• Search type and count</li> <li>• Median and cumulative runtimes</li> <li>• Timestamp of the last search</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Cache Activity                | <p>Shows your deployment's cache (local storage) activity, aggregated by bucket count or GB.</p> <p>The chart shows the following bucket rates tracked over time:</p> <ul style="list-style-type: none"> <li>• <b>bucket_download</b>: Buckets that have automatically transferred from AWS S3 storage to the cache. Spikes occur when searches need to localize data into the cache and are a part of your deployment's normal execution. Spikes over a longer duration may indicate indexer overload, which affects optimal search and ingest performance.</li> <li>• <b>bucket_eviction</b>: Buckets that have been automatically cleared from the cache.</li> <li>• <b>bucket_upload</b>: Spikes occur when more data is being transferred to AWS S3, though this does not impact system performance. This metric tracks the following: <ul style="list-style-type: none"> <li>◆ Buckets that have automatically been added to AWS S3 storage. These are generally hot buckets that are converting to warm buckets.</li> <li>◆ Buckets that already exist in AWS S3 storage but their internal details have changed; this results in an automatic resubmit and re-upload. For example, when a Splunk Cloud Platform administrator uses the delete command, this action affects the bucket's internal details and results in a re-upload of the modified bucket.</li> </ul> </li> </ul> |

| Panel or Filter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Search Details  | <p>The table lists the following information:</p> <ul style="list-style-type: none"> <li>• Report name/search string: Shows the report name for saved searches or the search string for ad hoc searches.</li> <li>• Search runtime: Enter a value in the <b>Search runtime &gt;= (seconds)</b> field to restrict the results to searches that meet or exceed this runtime.</li> <li>• Search start: The initial start time of the search job.</li> <li>• Earliest and latest time: The earliest and latest times of the search's time range.</li> <li>• Search type</li> <li>• User and host names</li> <li>• SID: Search identifier</li> </ul> |

### **Interpret search usage statistics results**

To review searches by user, follow these steps:

1. Change the time frame to widen the scope. For example, set it to the week prior to the current date.
2. Split the search by users so that you can see if there are a few users who are typically running longer searches.
3. Sort by **Cumulative Runtime** to see which users have the most cumulative search time.
4. Sort by **Median Runtime** to see which users are running the median longest searches.
5. If the user running the most or longest searches is the system user, you may want to review your applications to make sure that you have optimized them, and that they are providing the expected value. You may discover that some applications are not needed or are not used.

To review long-running searches, follow these steps:

1. Expand the time range to at least 24 hours. Searches are automatically sorted by long-running searches.
2. Set **Only Ad Hoc Searches** to **No** if it isn't already. This setting ensures that you see only scheduled searches, which are more likely to be long-running searches than ad hoc searches.
3. Scroll to the **Search Details** panel where the searches are sorted by search runtime.
4. Select the search name to view more details, and scroll to the bottom of the screen. Two events are displayed. In the second event, you can see the search query.

If you discover a long-running query that runs frequently, you may want to expand the time range to a week or longer to see how commonly this search is run. If it is running frequently, consider optimizing the search.

Be sure to monitor the **Cache Activity** panel for any spikes in bucket download size and count that have a long duration. Constant and excessive downloading of data into the cache churns the local file system on your indexers, and sustained bucket download spikes may cause performance impacts in searching and ingestion. To remediate these issues, you may need to modify search patterns, alter retention, or upgrade your workload-based or ingestion-based subscription entitlements so more of your searchable data is able to reside in the cache.

### **Check scheduler activity**

The CMC Scheduler Activity dashboard provides information to Splunk Cloud Platform administrators about how search jobs, also known as reports, are scheduled. Use this dashboard to understand the performance of your scheduled activity and determine if it needs any optimization.

See also Configure the priority of scheduled reports in the Splunk Cloud Platform *Reporting Manual*.

## Review the Scheduler Activity dashboard

This dashboard contains a number of panels about your scheduler activity.

Panels are grouped into one of three views, with a fourth view that combines the other three views so you can see all the data concurrently. Filter the results by specifying a time range and opting to include or exclude acceleration searches.

To investigate your panels, go to **Cloud Monitoring Console > Search > Scheduler Activity**. Use the following table to understand the dashboard interface.

| View or Filter                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Range                    | Set the time range for the data display.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Select View                   | Select a specific view, <b>Status</b> , <b>Activity</b> , or <b>Performance</b> , or select <b>All</b> to see a combined view.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Include Acceleration Searches | Acceleration searches are summaries of large datasets, used to help efficiently report on large volumes of data.<br><br>Select <b>Yes</b> to include summary-based acceleration searches in the displayed results. Select <b>No</b> to include only searches run on the complete dataset in the results.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Status                        | Contains the following two panels:<br><br><ul style="list-style-type: none"> <li>• <b>Skip Ratio (Last Hour)</b> shows the ratio of schedulers that skipped in the last hour.</li> <li>• <b>Average Execution Latency (Last Hour)</b> shows the latency in seconds.</li> </ul><br>The smaller number and the arrow indicate if this is an increase or decrease from the previous reading.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Activity                      | Contains the following four panels:<br><br><ul style="list-style-type: none"> <li>• <b>Scheduler Executions Detail</b> table with a <b>Group by</b> filter. The <b>Total</b> value is affected by the specified <b>Time Range</b> and <b>Include Accelerated Searches</b> filter values.</li> <li>• <b>Scheduler Executions</b> bar graph with a <b>Group by</b> drop-down list.</li> <li>• <b>Scheduled Report Completions</b> bar graph with a <b>Runtime Aggregation</b> filter.</li> <li>• <b>&lt;variable&gt; Concurrency of Scheduled Searches</b> bar graph with <b>Aggregation</b>, <b>Search Mode</b>, and <b>Group by</b> filters. The specified <b>Aggregation</b> value determines the variable. If you select <b>maximum</b> aggregation and <b>historical</b> search mode, The results show the maximum concurrency of completed scheduled searches. The results do not include searches that are currently running or searches that haven't completed within the displayed 5 minute intervals.</li> </ul> |
| Performance                   | Contains the following two panels:<br><br><ul style="list-style-type: none"> <li>• <b>Scheduler Errors and Warnings</b> table. The <b>Total</b> value is affected by the specified <b>Time Range</b> and <b>Include Accelerated Searches</b> filter values.</li> <li>• <b>Execution Latency</b> bar graph with <b>Group by</b> and <b>Aggregation</b> filters.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

### Interpret scheduler activity results

When interpreting your scheduler activity results, note the following:

- Be sure to include and exclude acceleration searches and see how that affects your results.
- Check for any significant spikes or dips in scheduler status, activity, and performance, particularly as compared against historical data. Large spikes and dips generally indicate an issue that you must resolve as soon as possible.
- Choose various time range values when reviewing the **Scheduler Errors and Warnings** panel. Doing this ensures that you see the errors and warnings reported for that specific time range.

## Investigate skipped scheduled searches

The CMC Skipped Scheduled Searches dashboard provides information to Splunk Cloud Platform administrators on skipped searches and search errors. Use this dashboard to identify why the Splunk platform can't process your scheduled searches and take steps to correct the issues.

See also Prioritize concurrently scheduled reports in Splunk Web and Offset scheduled search start times in the Splunk Cloud Platform *Reporting Manual*.

### **Review the Skipped Scheduled Searches dashboard**

This dashboard includes six panels of summary, graphical, and tabular data. Filter the results by specifying a time range and opting to include or exclude acceleration searches.

To investigate your panels, go to **Cloud Monitoring Console > Search > Skipped Scheduled Searches**. Use the following table to understand the dashboard interface.

| Panel or Filter                               | Description                                                                                                                                                                                                                                                                                |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Range                                    | Set the time range for the data display.                                                                                                                                                                                                                                                   |
| Include Acceleration Searches                 | Acceleration searches are summaries of large datasets, used to help efficiently report on large volumes of data.<br><br>Select Yes to include summary-based acceleration searches in the displayed results. Select No to include only searches run on the complete dataset in the results. |
| Total Skipped Searches                        | Shows the total number of skipped searches.                                                                                                                                                                                                                                                |
| Scheduled Search Skip Ratio                   | Shows the percentage of your scheduled searches that had to be skipped.                                                                                                                                                                                                                    |
| Skipped Scheduled Searches Detail             | Shows a table with <b>Group by</b> filter.                                                                                                                                                                                                                                                 |
| Skipped Searches                              | Shows a bar graph of skipped scheduled searches with <b>Group by</b> filter.                                                                                                                                                                                                               |
| Skipped Scheduled Searches by Name and Reason | Shows a table that lists the following: <ul style="list-style-type: none"> <li>• Report name</li> <li>• Skip reason and count</li> <li>• Alert actions</li> <li>• Total skips</li> </ul>                                                                                                   |
| Scheduler Errors and Warnings                 | Shows a table that lists the following: <ul style="list-style-type: none"> <li>• Error message</li> <li>• Count, or number of times the message was issued</li> <li>• Percent of Total, indicating what percent of the total this error was issued</li> </ul>                              |

### **Interpret skipped scheduled searches results**

If you are skipping searches, it can be indicative of the following possible problems with your search scheduling or query formation:

- You have scheduled too many searches to run at the same time. Alleviate this problem by staggering the scheduled searches.
- You have a search that attempts to run before the previously scheduled search has completed. For example, you schedule Search\_A to run every 5 minutes, but the first instance of the search takes 10 minutes to complete. The

next time the search is scheduled to run, it is skipped because the first search has not yet completed. Correct this issue by either adjusting the time range (set it to 10 minutes instead of 5) or optimizing your search to improve its performance. See *About search optimization* in the Splunk Cloud Platform *Search Manual*.

- You have skipped searches because your users have met the threshold for concurrency limits that you set in your Splunk System Limits. This is expected behavior, but it may also indicate that your users need help in optimizing their searches.

To check for skipped searches, perform these steps:

1. In the **Time Range** field of Skipped Scheduled Searches, select **24 hours** to get a better picture of your searches historically.
2. In the **Skipped Scheduled Searches Detail** panel, sort by **Reason**. Frequently, there are a number of skipped searches for the same reason. Note the primary reason or reasons that searches are skipped.
3. Scroll down to see which report is generating the primary issues. Note the report name and determine the following:

- If this is an expected behavior, you don't need to research any further.
- If the skipped searches are unexpected, continue to the next step.

- Go to **Settings > Searches Reports and Alerts**.

- If you know the App associated with the search or report, you can sort by the App. Otherwise, search by the report or search name.

- Locate the needed search or report and select it to open the **Edit Search** dialog box.

- Determine if the problem is with the search formation or the scheduling:

- If you need to troubleshoot the formation of the search, look for wild cards and check whether an index is specified.
- If scheduling is the problem, continue to the next step.

- Go to **Edit > Edit Schedule** to review the schedule for the search.

- Verify that the schedule for the report or search is in line with how long the search takes to complete. For example, if the report runs every hour but it takes 1.5 hours to run the search, the searches are skipped.

### ***Enable the Skipped Scheduled Searches alert***

Skipped searches can be indicators of non-optimal performance in your deployment. A high ratio of skipped scheduled searches can indicate the following:

- The number of searches being run exceeds your deployment's capacity.
- The searches being run are taking too long, or using large amounts of memory or CPU.

CMC provides an alert that lets you know when the ratio of skipped scheduled searches exceeds 20% in a 60-minute period. For more information about using this alert, see [Use the Alerts panel](#). For general information about managing alerts, see the Splunk Cloud Platform *Alerting Manual*.

## **Analyze expensive searches**

The CMC Expensive Searches dashboard provides information to Splunk Cloud Platform administrators on searches that are high consumers of your Splunk Cloud Platform resources. Use this dashboard to determine if these expensive and possibly inefficient searches are worth their cost.

## Review the Expensive Searches dashboard

This dashboard provides four panels of data regarding expensive and inefficient searches. Set a time range to filter the results.

To investigate your panels, go to **Cloud Monitoring Console > Search > Expensive Searches**. Use the following table to understand the dashboard interface.

| Panel or Filter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time Range                               | Set the time range for the data display.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Maximum Runtime Searches                 | Shows a line graph of search duration in seconds over time, comparing maximum ad hoc searches against scheduled searches.                                                                                                                                                                                                                                                                                                                                                                                            |
| Top 20 Most Memory Consuming Searches    | Shows a table that lists the following: <ul style="list-style-type: none"> <li>• Splunk platform instance label</li> <li>• Provenance</li> <li>• Percentage memory used (KB)</li> <li>• Search duration and start time and date</li> <li>• Search type, mode, and app</li> <li>• User name and role</li> </ul>                                                                                                                                                                                                       |
| Top 20 Most Expensive Ad Hoc Searches    | Shows a table that lists the following: <ul style="list-style-type: none"> <li>• Search time</li> <li>• User</li> <li>• Time range start and end</li> <li>• Search duration and result count</li> <li>• Memory usage (KB)</li> <li>• Total number of events scanned</li> <li>• Search query</li> </ul>                                                                                                                                                                                                               |
| Top 20 Most Expensive Scheduled Searches | Shows a table that lists the following: <ul style="list-style-type: none"> <li>• Search time</li> <li>• User</li> <li>• Scheduled time</li> <li>• Status</li> <li>• Search duration and result count</li> <li>• Memory usage (KB)</li> <li>• Saved search name</li> </ul>                                                                                                                                                                                                                                            |
| Potentially Inefficient Searches         | Shows a table that lists the following: <ul style="list-style-type: none"> <li>• User</li> <li>• Search Processing Language (SPL)</li> <li>• Events scanned</li> <li>• Search time range in days</li> <li>• Search duration</li> <li>• Splunk query score: A calculated number based on weighted indicators within an SPL string. A high score indicates a very inefficient search.</li> <li>• Potentially inefficient behavior: Shows the indicators within an SPL string that reduce search efficiency.</li> </ul> |

## Interpret expensive searches results

When interpreting your expensive searches results, note the following:



- After you identify the expensive and inefficient searches in your deployment, collaborate with users to improve the queries, using the information in the Write better searches topic in the Splunk Cloud Platform *Search Manual*.
- Review the score range of your searches using the Splunk Query Score column in the **Potentially Inefficient Searches** panel, and optimize searches that received a high score as soon as possible.

## Use the Usage dashboards

The dashboards accessed from the **Cloud Monitoring Console > Usage** tab enable Splunk Cloud Platform administrators to monitor the users who have access to your Splunk Cloud Platform deployment. The User Activity dashboard lets you review user activity in the system over a specified time range. The User Detail dashboard lets you drill down into the activities of a specific user.

A blue progress bar might appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify any Cloud Monitoring Console (CMC) dashboard. Changing any of the search criteria, formatting, or layouts may cause inaccurate results and also override the automatic update process.

### View user activity

The CMC User Activity dashboard provides statistical information to Splunk Cloud Platform administrators about all users with access to your Splunk Cloud Platform deployment, including their pageviews and apps. Use this dashboard to understand your users' activity and take the appropriate action to monitor their behavior.

#### *Review the User Activity dashboard*

This dashboard shows summary and detailed information about user activity within a specified time range. The time range value affects all panels. The dashboard generally shows information for multiple users, but this is dependent on the number of users active during the specified time range.

The top three panels are summarized totals of a particular area: users, apps, or pageviews. You use the middle three panels to analyze more specific data about each area by setting a Sort by option. The bottom panel provides a tabular summary of the Total Pageviews panel.

The **Pageviews** area contains two panels with a variable in the panel title that you set with a filter option: **Top <variable>** and **<variable> Access over Time**.

To investigate your panels, go to **Cloud Monitoring Console > Usage > User Activity**. Use the following table to understand the dashboard interface.

| Panel or Filter | Description                                                                                                         |
|-----------------|---------------------------------------------------------------------------------------------------------------------|
| Time Range      | Set the time range for the data display.<br>This time range setting affects the display of all panels on this page. |
| Distinct Users  | Shows the total number of distinct users who were active in the specified time range.                               |

| Panel or Filter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | To view the individual users included in this total and their access activity, sort by <b>User</b> in <b>Pageviews</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Distinct Apps    | Shows the total number of distinct apps that were active in the specified time range.<br>To view the individual apps in this total and their access activity, sort by <b>App</b> in <b>Pageviews</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Total Pageviews  | Shows the number of total pageviews that were accessed in the specified time range.<br>To view the individual pages in this total and their access activity, sort by <b>Page</b> in <b>Pageviews</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Pageviews        | This area includes the <b>Top &lt;variable&gt;</b> and <b>&lt;variable&gt; Access over Time</b> panels. The <variable> changes depending on the selected <b>Sort by</b> option.<br><br>Select a <b>Sort by</b> option to view data by user, app, or page. The panels show the following pie chart and a bar chart combinations: <ul style="list-style-type: none"> <li>• <b>Top Users</b> and <b>User Access</b>: Graphical information for the <b>Distinct Users</b> panel.</li> <li>• <b>Top Apps</b> and <b>App Access</b>: Graphical information for the <b>Distinct Apps</b> panel.</li> <li>• <b>Top Pages</b> and <b>Page Access</b>: Graphical information for the <b>Total Pageviews</b> panel.</li> </ul> |
| Activity by Page | Shows the number of distinct users and pageviews for a specific app page.<br>Like the <b>Top Pages</b> and <b>Page Access over Time</b> panels, this panel provides additional information for the <b>Total Pageviews</b> panel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### **Interpret user activity results**

When interpreting your user activity results, note the following:

- When analyzing a particular area, be sure to select the corresponding **Sort by** option in **Pageviews** so that you view the summarized total and detail information together. For example, select **Users** when you are investigating distinct users to include the information at the **Top Users** and **User Access over Time** panels.
- If a particular user has a large number of pageviews, this might indicate that users are sharing credentials, or that the account is being used on a dashboard with multiple refreshes, creating additional search load.

### **View user detail**

The CMC User Detail dashboard provides comprehensive information to Splunk Cloud Platform administrators about the activities of a specific user in your Splunk Cloud Platform deployment. Use this dashboard to investigate a user, particularly their searches and pageviews.

#### **Review the User Detail dashboard**

This dashboard contains 10 panels of summary, graphical, and tabular data about a specific user of your Splunk Cloud Platform deployment. Filter the results by specifying a time range.

To investigate your panels, go to **Cloud Monitoring Console > Usage > User Detail**. Use the following table to understand the dashboard interface.

| Panel or Filter | Description |
|-----------------|-------------|
|                 |             |

| Panel or Filter        | Description                                                                                                                                                                                                                                     |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User and Time Range    | Select a user and time range to populate the dashboard. <b>Note:</b> When you view this dashboard, the User field is automatically populated with the first menu value. Be sure to change this default value to the user you are investigating. |
| Name                   | Shows the user's full name as defined in their user profile.                                                                                                                                                                                    |
| Username               | Shows the user's system name, which may be formatted as an email account.                                                                                                                                                                       |
| Roles                  | Shows the user's roles as assigned in <b>Settings &gt; Users and Authentication &gt; Roles</b> .                                                                                                                                                |
| User Logins            | Shows the number of times the user has logged into the deployment.                                                                                                                                                                              |
| Search Count           | Shows the number of searches the user has done.                                                                                                                                                                                                 |
| Average Search Runtime | Shows the average runtime of all searches performed.                                                                                                                                                                                            |
| Total Search Time      | Shows the total of the combined runtimes for all the user's searches.                                                                                                                                                                           |
| Searches by Type       | Provides a bar chart of the types of searches the user performed over time, such as ad hoc.                                                                                                                                                     |
| Pageviews by App       | Provides a bar chart of pageviews, color-coded by app.                                                                                                                                                                                          |
| User Search Detail     | Provides detailed information for each search done by the user, such as its type, start and run times, and host and search identifiers.                                                                                                         |

### **Interpret user detail results**

When interpreting your user detail results, note the following:

- Especially for new deployments, use this dashboard to monitor usage and determine if your users are adopting or rejecting the Splunk platform. A downward trend may indicate users need more training, or that there are other issues that require investigation and resolution.
- Check if the values in the second row of panels (**Search Count**, **Average Search Runtime**, and **Total Search Time**) increase significantly over time. Depending on the user's role and responsibilities, this may indicate behavior that needs further investigation.
- Analyze if the **Searches by Type** and **Pageviews by App** graphs reveal any patterns of behavior that could be optimized.

## **Use the License Usage dashboards**

The first three dashboards accessed from the **Cloud Monitoring Console > License Usage** tab enable Splunk Cloud Platform administrators to monitor their Splunk Cloud Platform subscription entitlement and ensure they don't exceed their license limits.

To review all of your organization's subscription entitlements, see the [Entitlements dashboard](#).

If your organization has an ingest-based subscription that measures by the amount of data ingested, see the [Ingest dashboard](#).

If your organization has a workload-based subscription that measures by Splunk Virtual Compute (SVC) units, see the [Workload dashboard](#).

For more detailed information about the different subscription types, see the Splunk Cloud Platform *Service Description*. Be sure to choose the correct service description version for your Splunk Cloud Platform deployment from the **Version** drop-down menu.

For more information about your organization's particular subscription entitlement, or to convert from an ingest-based subscription to a workload-based subscription, contact your Splunk account representative.

The last three dashboards accessed from the **Cloud Monitoring Console > License Usage** tab enable Splunk Cloud Platform administrators to monitor their Splunk Cloud Platform storage and usage entitlement. Splunk Cloud Platform retains data based on index settings that enable you to specify when data is to be deleted. Data retention capacity space in your Splunk Cloud Platform service is based on the volume of uncompressed data that you want to index on a daily basis.

Storage is based on your subscription type. You can also purchase additional data retention capacity. For more information, see the following information in the Splunk Cloud Platform Service Description:

- Storage
- Subscription types

For more information about creating and managing Splunk Cloud Platform indexes, see [Manage Splunk Cloud Indexes](#) in the Splunk Cloud Platform *Admin Manual*.

A blue progress bar might appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify any Cloud Monitoring Console (CMC) dashboard. Changing any of the search criteria, formatting, or layouts may cause inaccurate results and also override the automatic update process.

## Monitor your entitlements

Splunk Cloud Platform administrators use the Entitlements dashboard on the CMC to review the entitlement limits for their organization's subscription.

The panels show numerical values for the following entitlement limits:

- `<variable>` License Entitlement: The variable in this title displays either **Ingest** or **Workload**, based on your subscription type.
- Searchable storage: Dynamic Data Active Searchable (DDAS)
- Archive storage: Dynamic Data Active Archive (DDAA)

Entitlement limits are specific to and based on your organization's unique requirements for ingesting and storing data with Splunk Cloud Platform. In particular, searchable and archive storage limits are specific to your Splunk Cloud Platform subscription because your organization may opt to purchase additional storage. For more information, see the following:

- The Storage section in the *Splunk Cloud Platform Service Description*.
- The topics about managing indexes and archived data in the [Manage your Indexes and Data in Splunk Cloud Platform](#) chapter of the *Splunk Cloud Platform Admin Manual*.

### Review the Entitlement dashboard

To investigate your panels, go to **Cloud Monitoring Console > License Usage > Entitlement**.

| Panel | Description |
|-------|-------------|
|-------|-------------|

|                                                          |                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <variable> License Entitlement                           | Shows title of <b>Workload License Entitlement</b> and total number of SVCs if your organization has a workload-based subscription.<br><br>Shows title of <b>Ingest License Entitlement</b> and ingest limit in GB if your organization has an ingest-based subscription. |
| Searchable Storage (DDAS) Entitlement                    | Shows your Dynamic Data Active Searchable (DDAS) storage entitlement in GB.                                                                                                                                                                                               |
| Archive Storage (DDAA) Entitlement                       | Shows your Dynamic Data Active Archive (DDAA) entitlement in GB. Shows <b>N/A</b> if this isn't applicable for your organization's subscription.                                                                                                                          |
| Data scan entitlement for Federated Search for Amazon S3 | Shows your amount of data scan entitlement available. If your organization doesn't have a license for Federated Search for Amazon S3, this panel is not visible.                                                                                                          |

### ***Interpret the entitlement results***

Because entitlement limits are determined by your organization's Splunk Cloud Platform subscription, contact your Splunk account representative with any questions about the displayed values.

## **Monitor current usage of your ingestion-based subscription**

If your Splunk Cloud subscription plan measures the search workload consumption by the amount of data ingested, Splunk Cloud Platform administrators use the Ingest dashboard on the CMC to monitor usage and stay within their subscription entitlement.

Splunk Cloud Platform administrators can also use the **SVC Usage** panel in the Workload dashboard to view basic information about their organization's projected SVC utilization. Workload-based subscriptions use Splunk Virtual Compute (SVC) as a unit of measure. To understand the potential SVC equivalent for your ingest-based subscription, see Performance considerations in the Splunk Cloud Platform *Service Description*. Be sure to view the correct service description version for your Splunk Cloud Platform deployment version.

For any questions about your organization's ingest-based subscription, or to convert from an ingest-based subscription to a workload-based subscription, contact your Splunk account representative.

### ***About the Ingest dashboard***

The Ingest dashboard contains four panels visible to Splunk Cloud Platform administrators:

- **License Entitlement** shows the licensed limit in GB for your organization's ingest-based subscription. This entitlement also displays as a red horizontal line in the **Daily License Usage** panel.
- **Daily License Usage summary**, **Daily License Usage details**, and **Average and Peak Daily Volume** show data ingestion in GB over a 30-day time range. These panels derive information from your organization's license manager and present data in a bar chart.
  - ◆ To view split-by details from the **Daily License Usage summary** or **Daily License Usage details** panels, click and drag an area of the panel to focus on a time range. Then use the **Split by** drop-down list to split the displayed results by **host**, **index**, **source**, or **source type**.

The Daily License Usage summary panel uses the UTC timezone. The Daily License Usage details panel uses the timezone that your Splunk Cloud Platform instance is set to. You might see a discrepancy between the panels if your Splunk Cloud Platform instance timezone is not set to UTC.

The Daily License Usage summary, Daily License Usage details, and Average and Peak Daily Volume panels use daily totals event data collected from the license\_usage\_summary.log file when you choose No Split. When you choose a Split by option, the panels use event data collected from the license\_usage.log file. If the license manager is down at its local midnight, it won't generate the events for that day, and you won't see that day's data in the panels.

### Review the Ingest dashboard

To investigate your panels, go to **Cloud Monitoring Console > License Usage > Ingest**.

Chart series values are color-coded. See the key on the side of a panel for the specific values included in a chart.

| Filter option       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| License Entitlement | Shows the licensed limit in GB for your organization's ingest-based subscription. See the red <b>license limit</b> horizontal line in the <b>Daily License Usage</b> panel to determine if your organization's ingestion rate stays under the limit.<br><br>Shows <b>N/A</b> if your organization has a workload-based subscription to Splunk Cloud Platform.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| No Split            | The panels show license volume and usage data for all data pools.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Split by value      | Select a <b>Split by</b> option of Source Type, Host, Source, or Index. The panels may show the following behavior: <ul style="list-style-type: none"> <li>• <b>Daily License Usage:</b> Shows up to 11 color-coded series of the selected option. This includes the top 10 series and OTHER, a summary category that includes series not in the top 10.</li> <li>• <b>Average and Peak Daily Volume:</b> Shows the average and peak daily values for the top five series of the selected option.</li> </ul> <p>Data may display as SQUASHED when you split by host or source. This is because every license peer periodically reports to the license manager its stats for the data indexed, broken down by source, source type, host, and index. If the number of distinct tuples (host, source, source type, index) grows beyond a configurable threshold, Splunk software squashes the host and source values and only reports a breakdown by source type and index. This is done to conserve internal resources.</p> <p>Because of squashing on the other fields, only the split-by source type and index guarantee full reporting. Split by source and host do not guarantee full reporting if those two fields represent many distinct values. The panels show the entire quantity indexed, but not the names. This means that you don't know who consumed a particular amount, but you know what the amount consumed is.</p> |

### Interpret ingestion-based results

The series in a bar chart are individually color coded so you can analyze usage patterns and take any appropriate action. For example:

- You set **Split by** to Index and see that a certain index shows an unusually high spike in usage. Investigate the cause of the spike and determine if it requires remediation.
- You see that your daily usage and average and peak volumes are consistently close to or exceeding your license limit. Contact your Splunk account representative to upgrade your subscription.

Select any bar in the chart to view the underlying data for the bar. Be sure to not modify the underlying data in any way.

You can also set up an alert action (for example, send an email) to be performed when a platform alert is triggered. Go to **Settings > Searches, Reports, and Alerts** and select **New Alert** to define a new alert action. See also the *Determine retention usage and set an alert section* in [Interpret index and storage capacity results](#) in the Splunk Cloud Platform

## Monitor current SVC usage of your workload-based subscription

If your Splunk Cloud Platform subscription plan measures your deployment's ingestion and search workload consumption by Splunk Virtual Compute (SVC) units, Splunk Cloud Platform administrators use the Workload dashboard on the CMC to monitor usage. For more information about the SVC entitlement for your workload-based subscription, see Performance considerations in the Splunk Cloud Platform *Service Description*. Be sure to view the correct service description version for your Splunk Cloud Platform deployment version.

### Review the Workload dashboard

The Workload dashboard contains panels visible to Splunk Cloud Platform administrators that show SVC entitlement and usage for either ingest-based or workload-based subscriptions over a specific time range.

This dashboard shows your deployment's overall SVC usage and can help locate where you can optimize your organization's SVC consumption. Hover your mouse pointer over a vertical bar or a point on a line to view data for a specific hour.

The **SVC usage per hour by search type** and **SVC usage per hour by top <variable>** panels represent less accurate data due to sampling rates. These panels use the search\_launcher process, which represents searches that take less than 10 seconds to complete. This process might hide a lot of data. For more accurate data, view the **Search time by search type** and **Search time by top 10 apps, users, and searches** panels.

To investigate your panels, go to **Cloud Monitoring Console > License Usage > Workload**. Use the following table to understand the dashboard interface.

| Panel                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total number of licensed SVCs | <p>Shows the number of SVCs assigned to your organization's subscription per your license entitlement.</p> <p>This panel displays an <b>N/A</b> for the following scenarios:</p> <ul style="list-style-type: none"> <li>Subscription status: Your organization has a new workload-based subscription and Splunk is still processing your SVC entitlement. Once this process is complete, your entitlement will appear.</li> <li>Subscription type: Your organization uses ingest-based licensing. Contact your Splunk account representative to convert your subscription type from ingest-based to workload-based.</li> </ul>                                                                                                                                                                                                                                                                                                         |
| Peak SVC usage                | <p>Shows your organization's SVC usage against the license limit.</p> <p>This chart shows hourly usage calculated in standard 1 hour time blocks, meaning 9:00-9:59 AM or 11:00-11:59 PM. Use the time picker to adjust the granularity by 1 hour, 15 minutes, or 5 minutes. Finer time granularity selection offers increased visibility into when SVC usage peaks or dips within a given timeframe, so you can understand whether usage is consistently high or if there might be specific workloads causing spikes in usage.</p> <p>The displayed data excludes data gathered during both the current hour and one previous hour. This means that if you are viewing this chart at 2:58 PM, data from 1:00-1:59 PM (the previous hour) and 2:00-2:59 PM (the current hour) is excluded from calculation. At 3:00 PM, data from 1:00-1:59 PM will be included, and at 4:00 PM, the data from 2:00-2:59 PM will be included. This</p> |

| Panel                                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                  | <p>exclusion is to ensure the correct calculation of your organization's SVC utilization.</p> <p>For workload-based subscriptions:</p> <ul style="list-style-type: none"> <li>• Color-coded vertical bars show the following about SVC usage: <ul style="list-style-type: none"> <li>◆ Blue bars indicate usage that is below the optimal threshold.</li> <li>◆ Yellow bars indicate usage that is at or above the optimal threshold of 80% of the licensed amount. Splunk Cloud Platform administrators might see issues with their deployment when the usage remains elevated for extended periods of time.</li> <li>◆ Red bars indicate usage that is above 90% of the licensed amount. This indicates a degraded state. Splunk Cloud Platform administrators will likely see issues with their deployment when the usage remains degraded for extended periods of time.</li> </ul> </li> <li>• Color-coded horizontal reference lines show the following: <ul style="list-style-type: none"> <li>◆ Green: Your organization's average SVC utilization.</li> <li>◆ Yellow: The optimal utilization threshold, which is calculated as 80% of the license limit.</li> <li>◆ Red: Your organization's SVC entitlement or license limit.</li> </ul> </li> </ul> <p>Generally, SVC usage should be less than 80% to maintain performance. 80% to 90% is considered elevated usage. Greater than 90% usage might cause degraded performance. If utilization exceeds 80%, look at the detail panels and consider optimizing processes that are high SVC consumers. Or, you can contact your Splunk account representative to discuss increasing your license entitlement.</p> <p>For ingest-based subscriptions, the following elements don't appear:</p> <ul style="list-style-type: none"> <li>• Reference lines for SVC entitlement and 80% utilization threshold.</li> <li>• The yellow elevated and red degraded usage bars.</li> </ul> <p><b>Note:</b> The displayed SVC values for ingest-based subscriptions are only a projected estimate. The actual appropriate SVC entitlement for your organization might be affected by various usage factors. To determine the appropriate SVC entitlement for your deployment and to convert your ingest-based subscription to a workload-based subscription, contact your Splunk account representative.</p> |
| <p>Peak SVC usage as a percentage of allocated SVCs per tier</p> | <p>Shows SVC peak usage as a percentage of SVCs provisioned by the search head and indexer tier. Use the time picker to adjust the granularity by 1 hour, 15 minutes, or 5 minutes.</p> <p>Provisioned SVCs are allocated to the search head and indexer tiers after initial sizing conversations about intended workloads and requirements, with intention to minimize the footprint for both tiers. Viewing the usage as a percentage of provisioned SVCs provides insight on a tier level and helps you understand what utilization looks like if one tier is over extended. Review the percentage usage on each tier to identify which tier is close to exceeding the optimal range of greater than 80%.</p> <p>This panel has the following limitations:</p> <ul style="list-style-type: none"> <li>• This panel uses a new calculation as of CMC version 3.12.0 and does not display historical data. The data requires history before it's visible in the CMC. On day of release, this panel will contain approximately a week's worth of data.</li> <li>• This panel does not break down usage percentage by individual search heads.</li> </ul> <p>The displayed data excludes data gathered during both the current hour and one previous hour. This means that if you are viewing this chart at 2:58 PM, data from 1:00-1:59 PM (the previous hour) and 2:00-2:59 PM (the current hour) is excluded from calculation. At 3:00 PM, data from</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



| Panel                                                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                         | 1:00-1:59 PM will be included, and at 4:00 PM, the data from 2:00-2:59 PM will be included.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Peak SVC usage per hour split by process                                                | <p>Shows SVC consumption per hour by system processes and resources.</p> <ul style="list-style-type: none"> <li>• <b>Ingestion:</b> Encompasses both ingestion and indexing processes. This includes any <code>index</code> or <code>scripted_input</code> process and also processes on indexers that are not counted in the search or shared services categories. See the <b>SVC Usage by Ingestion</b> panel for a breakdown of the ingested data by either index or source type.</li> <li>• <b>Search:</b> Encompasses any running search process where the <code>process_type</code> starts with <code>search</code>.</li> <li>• <b>Shared services:</b> Encompasses internal system processes necessary to maintain service to your deployment. This includes any other non-search process on the search head, such as <code>kvstore</code> and <code>splunk_web</code> processes.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <variable> (search seconds, SVC usage) per hour by search type                          | <p><b>Search seconds per hour by search type</b> shows search seconds per hour by search type. This is the default view for this panel.</p> <ul style="list-style-type: none"> <li>• <b>REST_API:</b> Searches that use the Splunk REST API. See <b>Basic concepts about the Splunk platform REST API</b>.</li> <li>• <b>ad-hoc:</b> Searches that are unscheduled and manually run. See <b>ad hoc search</b>.</li> <li>• <b>dashboard:</b> Searches run by your dashboards</li> <li>• <b>scheduled:</b> Searches that are saved and scheduled so they automatically run. See <b>scheduled search</b>.</li> <li>• <b>summary director:</b> Maintenance tasks that run in the background involving caching and summarization to ensure searches are processed.</li> </ul> <p>Select <b>estimated SVC</b> to view <b>SVC usage per hour by search type</b>. This shows SVC consumption per hour as categorized by one of the following assigned search types. If the consumption can't be categorized in an assigned search type, it is grouped in the general <b>other</b> category.</p> <ul style="list-style-type: none"> <li>• <b>ad-hoc:</b> Searches that are unscheduled and manually run. See <b>ad hoc search</b>.</li> <li>• <b>report acceleration:</b> Searches that are related to accelerated data models or reports. See <b>data model acceleration, report acceleration</b>, and How data model acceleration differs from report acceleration and summary indexing.</li> <li>• <b>scheduled:</b> Searches that are saved and scheduled so they automatically run. See <b>scheduled search</b>.</li> <li>• <b>scheduled realtime:</b> Searches where the <code>search_mode</code> field value is <code>realtime indexes (RT Indexes)</code> and the <code>search_type</code> field value is <code>scheduled</code>.</li> <li>• <b>search launcher:</b> Ephemeral searches that are managed by the search launcher, which is a splunkd helper process that is responsible for forking new search processes and managing a high number of fast-running searches on deployments. Because the individual ephemeral searches are being quickly processed, your deployment's SVC usage for these searches is based on the search launcher process to ensure an accurate SVC calculation.</li> </ul> |
| <variable> (Search seconds, SVC usage) by top 10 <process type> (apps, searches, users) | <p><b>Search seconds by top 10 &lt;process type&gt;</b> shows search seconds per hour grouped by consumer type and search head. You can identify which apps, users, and searches per search head have relatively high search times. This is the default view for this panel.</p> <p>Select <b>estimated SVC</b> to view <b>SVC usage by top 10 &lt;process type&gt;</b>. This shows high consumers of SVC per hour grouped by consumer type and search head so you can take steps to optimize their consumption. For example, by analyzing the users and searches data, you can contact high consumers of SVC and discuss ways to optimize their consumption, such as improving their search queries.</p> <p>Select one of the following options from the <b>Process type</b> drop-down menu:</p> <ul style="list-style-type: none"> <li>• <b>Apps:</b> Lists a maximum of the top 10 apps and their respective SVC consumption.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

| Panel                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                              | <ul style="list-style-type: none"> <li>• <b>Users:</b> Lists a maximum of the top 10 users and their respective SVC consumption. These users may be human or virtual administrators.</li> <li>• <b>Searches:</b> Shows which searches utilize the greatest SVC as a percentage of the total consumption.</li> </ul> <p>Select one of the following options from the <b>Search head</b> drop-down menu:</p> <ul style="list-style-type: none"> <li>• <b>All:</b> Shows all search heads in your Splunk Cloud Platform deployment. This category includes all the data ingested and processed in the deployment.</li> <li>• <b>Historical:</b> Shows a different view of <b>All</b>. This category includes all the data ingested, processed, and summarized in the deployment prior to the CMC 2.9.0 release.</li> <li>• <b>Specific search head name:</b> Shows data for a specific search head that has been ingested, processed, and summarized in the deployment as of and after the CMC 2.9.0 release.</li> </ul> <p><b>Note:</b> One virtual administrator is the internal splunk-system-user, which runs jobs and processes like summary refreshes, report accelerations, and data model accelerations for a deployment on behalf of a Splunk Cloud Platform customer. Running these processes consumes SVCs. If the SVC usage of splunk-system-user seems abnormal, Splunk Cloud Platform administrators should contact the deployment's administrator to investigate the increased consumption.</p> |
| Dispatched and skipped search count per hour | <p>Shows the number of searches per hour that are dispatched or skipped.</p> <p>The yellow vertical lines indicating elevated SVC usage and the red vertical lines indicating degraded SVC usage correlate to the same lines in the <b>SVC Usage</b> panel.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Peak SVC usage per hour by indexing source   | Shows SVC consumption per hour by ingestion source. Select either <b>Index</b> or <b>Sourcetype</b> from the drop-down menu.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Hourly rate of ingestion                     | Shows the hourly rate of ingestion in GB. When data ingestion rates are high, the indexer consumes more resources to process and ingest data. High ingestion rates can increase SVC usage.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### **Interpret SVC usage results**

See the table in [Review the Workload dashboard](#) in this topic for information on keeping your SVC usage within license limits.

In the **Events** tab for a search, the search\_label field includes the `_ACCELERATE_{SID_NUMBER}` value so you can search for an event using its SID value.

You can also set up an alert action (for example, send an email) to be performed when a platform alert is triggered. Go to **Settings > Searches, Reports, and Alerts** and select **New Alert** to define a new alert action.

## **Monitor the Storage Summary dashboard**

This dashboard shows searchable and archive storage license usage data so Splunk Cloud Platform administrators can ensure their organization stays within its licensed subscription limits.

### **About the Storage Summary dashboard**

The Storage Summary dashboard highlights important information that also displays on the **Entitlements**, **Searchable Storage (DDAS)**, and **Archive Storage (DDAA)** dashboards. This dashboard provides insights into your data retention based on the uncompressed data you have indexed.

To view this dashboard, you must have the indexes\_edit capability.

## Review the Storage Summary dashboard

To investigate your panels, go to **Cloud Monitoring Console > License Usage > Storage Summary**.

| Panel                                                                                                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Searchable Storage (DDAS) Entitlement                                                                                | Shows the amount of your entitled searchable storage based on your DDAS license entitlement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Searchable Storage (DDAS) Usage                                                                                      | Shows the amount of searchable storage used by both customer-created and metered internal indexes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Searchable Storage (DDAS) Usage Percent                                                                              | Shows your percentage of usage compared to your DDAS license entitlement.<br><br>The value displays in the following colors to indicate status: <ul style="list-style-type: none"> <li>• Green: Usage is well under the entitlement limit.</li> <li>• Yellow: Usage is at or above 80% of the entitlement limit.</li> <li>• Red: Usage at or above 90% and close to exceeding the entitlement limit.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Archive Storage (DDAA) Entitlement                                                                                   | Shows the amount of your archive storage entitlement based on your DDAA license.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Archive Storage (DDAA) Usage                                                                                         | Shows the amount of archive storage used by both customer-created and metered internal indexes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Archive Storage (DDAA) Usage Percent                                                                                 | Shows your percentage of usage compared to your DDAA license entitlement.<br><br>The value displays in the following colors to indicate status: <ul style="list-style-type: none"> <li>• Green: Usage is well under the entitlement limit.</li> <li>• Yellow: Usage is at or above 80% of the entitlement limit.</li> <li>• Red: Usage at or above 90% and close to exceeding the entitlement limit.</li> </ul> <p>If your organization doesn't have a DDAA subscription, this panel displays <b>N/A</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                   |
| Restored Entitlement, Restored Searchable Storage (DDAS) Usage, and Restored Searchable Storage (DDAS) Usage Percent | For more information, see the panel descriptions in the <a href="#">Review the Searchable Storage (DDAS) dashboard</a> section.<br><br>If your organization doesn't have a DDAA subscription, these panels don't appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Index Details                                                                                                        | Provides a tabular overview of index retention and storage usage, per index. <ul style="list-style-type: none"> <li>• <b>Searchable Storage (DDAS) Retention Days</b></li> <li>• <b>Searchable Storage (DDAS) Index Size GB</b></li> <li>• <b>Archive Storage (DDAA) Retention Days</b></li> <li>• <b>Archive Storage (DDAA) Usage GB</b></li> <li>• <b>Archived GB Last 90 Days</b></li> <li>• <b>Expired GB Last 90 Days</b></li> </ul> <p>For <b>Archived GB Last 90 Days</b> and <b>Expired GB Last 90 Days</b>, the 90-day count is up to midnight of the previous day from when you accessed the dashboard. This means if you access the dashboard on January 1 at 9:00 AM, the 90th day of data is December 31 at 11:59 PM. <b>Searchable Storage (DDAS) Retention Days</b> and <b>Archive Storage (DDAA) Retention Days</b> also display values as of midnight of the previous day.</p> |

### Interpret storage summary results

- If the Searchable Storage (DDAS) Usage Percent panel value displays in red or yellow, this indicates that you need to reduce your DDAS usage. See the [Searchable Storage \(DDAS\) dashboard](#) for more detailed information.

- If the Archive Storage (DDAA) Usage Percent panel value displays in red or yellow, this indicates that you need to reduce your DDAA usage. See the [Archive Storage \(DDAA\)](#) for more detailed information.

## Monitor current usage of Searchable Storage (DDAS)

This dashboard shows comprehensive Dynamic Data Active Searchable (DDAS) license usage data so Splunk Cloud Platform administrators can ensure their organization stays within its licensed subscription limits.

### About the Searchable Storage (DDAS) dashboard

Dynamic Data Active Searchable (DDAS) is used for searching ingested data. DDAS is also commonly known as searchable storage. Review the information to ensure that you are staying within your subscribed limits for data ingestion and retention. The displayed data updates every time you access or refresh the dashboard in the CMC app. For more information, see [Restore archived data to Splunk Cloud Platform](#).

Your organization determines their DDAS entitlement amount when subscribing to the Splunk Cloud Platform. For questions about your organization's DDAS entitlement, contact your Splunk account representative. See also the "Data retention" and "Dynamic Data Active Searchable (DDAS)" sections in the Storage section of the Splunk Cloud Platform Service Description.

The Searchable Storage (DDAS) dashboard provides insights into your data retention based on the uncompressed data you have indexed.

### Review the Searchable Storage (DDAS) dashboard

To investigate your panels, go to **Cloud Monitoring Console > License Usage > Searchable Storage (DDAS)**.

| Panel                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Searchable Storage Entitlement   | Shows the amount of your searchable storage entitlement.<br><br>If you are an ingest-based customer, this value includes any additional storage you have purchased. If you are a workload-based customer, this value is the storage you have purchased. For questions about these entitlement values, contact your Splunk account representative.                                                                                                                                                                                                                                                                    |
| Searchable Storage Usage         | Shows the amount of searchable storage used by customer-created and metered internal indexes in GB. This value includes only actively searched storage and is calculated when you load this dashboard. Though this value will generally correspond to the total of the individual index values displayed in the Searchable Storage Index Details table, there may be differences due to the time the queries are performed, data aging out of indexes, and similar reasons.<br><br>Use this information to compare your current storage consumption against your subscription entitlement and data retention limits. |
| Searchable Storage Usage Percent | Shows your percentage of usage compared to your DDAS license entitlement.<br><br>The value displays in the following colors to indicate status: <ul style="list-style-type: none"> <li>• Green: Usage is well under the entitlement limit.</li> <li>• Yellow: Usage is at or above 80% of the entitlement limit.</li> <li>• Red: Usage at or above 90% and close to exceeding the entitlement limit.</li> </ul>                                                                                                                                                                                                      |
| Restored Entitlement             | Shows your entitlement limit for DDAA restores. For most Splunk Cloud Platform customers, this value is generally 10% of the amount that displays in the <b>Searchable Storage (DDAS) Entitlement</b> panel. If your organization has expanded their license to increase restoring capacity, the restored entitlement limit reflects this increase up to 20%. For more                                                                                                                                                                                                                                               |

| Panel                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                  | <p>information, see the following:</p> <ul style="list-style-type: none"> <li>• Restoration entitlement limit: Dynamic Data Active Archive (DDAA) in the Splunk Cloud Platform Service Description</li> <li>• DDAA restores: <a href="#">Restore indexed data from a self storage location</a></li> </ul> <p>If your organization doesn't have a DDAA subscription, this panel doesn't appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Restored Searchable Storage (DDAS) Usage         | <p>Shows the amount of restored storage used by both customer-created and metered internal indexes. This panel calculates searchable storage as the amount of restored data minus the expired and cleared data.</p> <p>If your organization doesn't have a DDAA subscription, this panel doesn't appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Restored Searchable Storage (DDAS) Usage Percent | <p>Shows the percentage of restored data usage compared to your restored storage entitlement.</p> <p>If your organization doesn't have a DDAA subscription, this panel doesn't appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Searchable Storage Usage Against Entitlement     | <p>Shows the amount of searchable storage used by all applicable indexes compared to your entitlement limit.</p> <p>This bar chart is the visualization for the <b>Searchable Storage Usage</b> panel.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Searchable Storage Usage by Top 10 Indexes       | <p>Shows the top 10 indexes that are high consumers of searchable storage.</p> <p>Select the <b>Include Internal Indexes</b> checkbox to include Splunk internal indexes in the chart and analyze if internal indexes are consuming high amounts of storage. See also the Splunk Internal Index Details table.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Searchable Storage Index Details                 | <p>Provides a tabular overview of searchable storage details per index that includes the following data:</p> <ul style="list-style-type: none"> <li>• <b>Oldest Event</b></li> <li>• <b>Newest Event</b></li> <li>• <b>Event Count</b></li> <li>• <b>Storage Retention Days</b></li> <li>• <b>Index Size GB</b></li> </ul> <p>Shows a table of the indexes in your deployment and the current searchable amount in GB for each actively searchable index. The searchable indexes of your deployment only include those in a hot or warm bucket. The GB value that displays for each index is calculated when you load this dashboard. Use this information to determine which indexes are high consumers of storage, and also understand general usage patterns and trends. For more information about index retention settings, see <a href="#">Manage data retention settings</a> in the Splunk Cloud Platform <i>Admin Manual</i>.</p> |
| Splunk Internal Index Details                    | <p>Provides a tabular overview of internal index details that includes the following data:</p> <ul style="list-style-type: none"> <li>• <b>Oldest Event</b></li> <li>• <b>Newest Event</b></li> <li>• <b>Event Count</b></li> <li>• <b>Storage Retention Days</b></li> <li>• <b>Default Retention Days</b></li> <li>• <b>Unmetered Index Size GB</b></li> <li>• <b>Metered Index Size GB</b></li> <li>• <b>Total Index Size GB</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

| Panel | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|       | <p>Splunk internal indexes can be identified by the underscore prefix ( _) in the index name and appear on other storage dashboards, such as the <b>Storage Summary</b> dashboard. You can opt to include internal indexes in the <b>Searchable Storage Usage by Top 10 Indexes</b> chart.</p> <p>An index with a storage value that exceeds the default value delivered by Splunk consumes additional license data. The <b>Default Retention Days</b> column shows Splunk default values. The <b>Storage Retention Days</b> column shows the actual storage retention value set for an index.</p> |

### ***Interpret your searchable storage results***

- A good method to determine if your data usage is running higher than expected is to check the dates of the earliest and latest events and compare this time period to the retention setting for the individual index. For example, if the earliest event is 2020/01/25, the latest event is 2020/01/31, and the retention setting for the index is 90 days, then the data ingestion for the index was met long before the time retention setting was met. So, the data ingestion was greater than anticipated.
- If an internal index displays a **Storage Retention Days** value that exceeds the **Default Retention Days** value, contact your Splunk account representative.

## **Monitor current usage of Archive Storage (DDAA)**

This dashboard shows comprehensive Dynamic Data Active Archive (DDAA) license usage data so Splunk Cloud Platform administrators can ensure their organization stays within its licensed subscription limits.

### ***About the Archive Storage (DDAA) dashboard***

Dynamic Data Active Archive (DDAA) is used as a long term storage and data in DDAA can be restored to DDAS to be searched. For Splunk Cloud Platform administrators, this dashboard shows information about your archived data for indexes that are enabled with DDAA. Review the information to ensure that you are staying within your subscribed limits for data ingestion and retention. The displayed data updates every time you access or refresh the dashboard in the CMC app. For more information, see [Store expired Splunk Cloud Platform data to a Splunk-managed archive](#).

Your organization must have enabled DDAA as part of its Splunk Cloud Platform subscription to see data in this dashboard. For more information, see the Dynamic Data Active Archive (DDAA) section in the Storage section of the Splunk Cloud Platform Service Description. If you exceed your storage requirements by ingesting more data than your initial estimate, Splunk Cloud Platform service elastically expands the amount of storage to retain your data per your retention settings. Periodically, Splunk will review and charge your account for any overages.

The Archive Storage (DDAA) dashboard provides insights into your data retention based on the uncompressed data you have indexed.

### ***Review the Archive Storage (DDAA) dashboard***

To investigate your panels, go to **Cloud Monitoring Console > License Usage > Archive Storage (DDAA)**.

| Panel                       | Description                                                                         |
|-----------------------------|-------------------------------------------------------------------------------------|
| Archive Storage Entitlement | Shows the amount of your archive storage entitlement.                               |
|                             | Shows the total amount of archive storage currently used by all applicable indexes. |

| Panel                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Archive Storage Usage                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Archive Storage Usage Percent             | <p>Shows the percentage of usage compared to your DDAA license entitlement.</p> <p>The value displays in the following colors to indicate status:</p> <ul style="list-style-type: none"> <li>• Green: Usage is well under the entitlement limit.</li> <li>• Yellow: Usage is at or above 80% of the entitlement limit.</li> <li>• Red: Usage at or above 90% and close to exceeding the entitlement limit.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Archive Storage Usage Against Entitlement | <p>Shows the amount of archive storage used by all applicable indexes compared to your entitlement limit.</p> <p>This bar chart is the visualization for the <b>Archive Storage Usage</b> panel.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Archive Storage Usage by Top 10 Indexes   | Shows your Top 10 indexes that are high consumers of archive storage.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Data Archive and Restoration Summary      | <p>Shows a summary of restoration activity for all of your deployment's indexes that are enabled with the DDAA feature from the last 90 days. The 90-day count is up to midnight of the previous day from when you accessed the dashboard. This means if you access the dashboard on January 1 at 9:00 AM, the 90th day of data is December 31 at 11:59 PM.</p> <p>These totals in GB show the amount of uncompressed (raw) data in the following categories:</p> <ul style="list-style-type: none"> <li>• <b>Total Size Restored GB:</b> Copied archive data that has been temporarily restored to an index. Restored data expires from searchable storage after 30 days.</li> <li>• <b>Total Size Cleared GB:</b> Restored data that has been manually removed from an index. This data has a Jobstatus of Cleared.</li> <li>• <b>Total Size Expired GB:</b> Data that has been automatically removed from searchable storage as it has passed the 30-day retention period. This data has a Jobstatus of Expired</li> </ul> <p>The displayed totals depend on the data you have selected to restore or clear and also the conditions and limitations of the restoration process, as follows:</p> <ul style="list-style-type: none"> <li>• The archival and restoration process is complete.</li> <li>• The data doesn't overlap with other data.</li> <li>• The data size doesn't cause performance issues.</li> </ul> <p>For more information, see the following in the the Splunk Cloud Platform <i>Admin Manual</i>:</p> <ul style="list-style-type: none"> <li>• <a href="#">Restore archived data to Splunk Cloud Platform</a></li> <li>• Step 7 of <a href="#">Steps to restore data to Splunk Cloud Platform</a> for the different types of Jobstatus</li> </ul> |
| Index Storage Usage Details               | <p>Provides a tabular overview of archive storage details per index that lists the following information:</p> <ul style="list-style-type: none"> <li>• Archived index name</li> <li>• Timestamps formatted in UTC for the earliest and latest archived events</li> <li>• 90-day data growth and expiration data in GB</li> <li>• Current usage amount in GB</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Interpret your archive storage results**

- Compare the archive usage against the entitlement and the growth against the expiration. If the usage and the growth consistently exceed the entitlement and the expiration, this indicates the following:
  - ◆ You must re-evaluate your index ingestion and retention settings. See the topics listed in the [See also](#) section on how to manage indexes and DDAA settings.

- ◆ You may need to upgrade your subscription to better handle your true data ingest and retention rates. Contact your Splunk account representative for help.
- Review the restoration totals and determine if the amount of data restored, cleared, and expired in your deployment meets or exceeds your organization's actual requirements. For example, a high total for restored data or low total for cleared or expired data may indicate the need to re-evaluate your index management policies and procedures. Ensure that you are restoring and retaining only the data that your organization truly needs.
- Be sure to convert event timestamps from UTC to your local time when analyzing the data in the **Index Storage Usage Details** table.

**See also**

| For more information about                                                        | See                                                                                                                         |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Splunk Cloud Platform data retention policies and available storage subscriptions | Storage                                                                                                                     |
| Managing your indexes, including searchable and archive storage                   | The <a href="#">Manage your Indexes and Data in Splunk Cloud Platform</a> section in the Splunk Cloud Platform Admin Manual |

**Monitor your Federated Search for Amazon S3 resources**

Federated Search for Amazon S3 lets you search data from your Amazon S3 buckets from your Splunk Cloud Platform deployment without needing to ingest or index it first. The **Federated Search for Amazon S3** dashboard in the CMC shows comprehensive data scan entitlement usage so your organization can stay within its limits.

**About the Federated Search for Amazon S3 dashboard**

This dashboard shows what your total data scan entitlement is and how much of that entitlement is used to date by your Federated Search for Amazon S3 searches in your current license term.

The dashboard tracks the volume of data on disk that is being scanned, not the amount of events that are being searched. Scans of data stored in compressed formats such as Parquet or GZIP will likely take up less of your entitlement than scans of data stored in uncompressed formats.

Review the information to ensure that you're staying within your Federated Search for Amazon S3 entitlement.

Your organization must have Federated Search for Amazon S3 set up as part of its Splunk Cloud Platform deployment to see data in this dashboard.

**Review the Federated Search for Amazon S3 dashboard**

To investigate your panels, go to **Cloud Monitoring Console** then **License Usage** then **Federated Search for Amazon S3**. The following panels display **N/A** if your organization does not have a Federated Search for Amazon S3 entitlement.

| Panel                                    | Description                                                                                              |
|------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Total data scan entitlement              | Total amount of data scanning capabilities available for use during your current license term.           |
| Data scan entitlement usage              | Total amount of data scanned by your searches during your current license term.                          |
| Percentage of data scan entitlement used | The percentage of data scanning capabilities utilized by your searches during your current license term. |



## ***Interpret federated search for Amazon S3 data scan entitlement usage***

The **Percentage of data scan entitlement used** panel is color-coded so you can quickly understand your usage. If your data scan entitlement usage is less than 80%, the panel data is green. If your usage is greater than 80%, the panel data is yellow. If your usage is greater than 90%, the panel data is red.

You can configure an alert action (for example, send an email) to be performed when your data scan entitlement usage exceeds 80%. Navigate to the CMC **Alerts** page to enable this alert: **Alerts** then **Configured Alerts** then **CMC Alert - S3 scanned volume exceeds 80% of the entitlement value**.

To learn more about CMC configured alerts, see [Use the Alerts panel](#).

If your data scan entitlement usage is consistently high, consider upgrading entitlements by contacting your Splunk Sales representative.

## **Monitor your Federated Analytics licenses**

This dashboard shows comprehensive Federated Analytics usage data so your organization can stay within its licensed limits. Federated Analytics uses the core platform entitlement for local data ingestion and Data Scan Units (DSUs) for federated searches.

This dashboard contains entitlement and usage metrics attributed to Data Scan Units (DSUs) for Federated Searches on external data sets. To learn more about Federated Analytics for Data Lakes, see [About Federated Analytics](#).

## **Use the Archive Management panel**

For Splunk Cloud Platform administrators, the Archive Management panel in the Cloud Monitoring Console (CMC) app shows information about your archived data for indexes that are enabled with Dynamic Data Active Archive (DDAA). Review the information to ensure that you are staying within your subscribed limits for data ingestion and retention. The displayed data updates every time you access or refresh the panel in the CMC app.

Your organization must have enabled DDAA as part of its Splunk Cloud Platform subscription to see data in this panel.

If you exceed your storage requirements by ingesting more data than your initial estimate, Splunk Cloud Platform service elastically expands the amount of storage to retain your data per your retention settings. Periodically, Splunk will review and charge your account for any overages. For more information and to understand storage requirements based on your subscription type, see the [Storage](#) section of the [Splunk Cloud Platform Service Description](#).

### ***Archive Summary***

In CMC, select the **Archive Management** link in the first panel of the **Storage Summary** or **Archive Storage (DDAA)** dashboard, then select the **Archive Summary** tab.

The summary information in this tab shows data on the usage, entitlement, and 90-day growth and expiration in GB for all of your deployment's indexes enabled with DDAA.

The archived data details table lists the following information:

- Archived index name
- Current size (GB)

- Timestamps for the earliest and latest archived events
- 90-day data growth and expiration data in GB

The amounts for the summarized and detailed growth and expiration data are for uncompressed (raw) data.

### ***Interpret these results***

Compare the usage against the entitlement and the growth against the expiration. If the usage and the growth consistently exceed the entitlement and the expiration, this indicates the following:

- You must re-evaluate your index ingestion and retention settings. See the topics listed in the [See also](#) section on how to manage indexes and DDAA settings.
- You may need to upgrade your subscription to better handle your true data ingest and retention rates. Contact your Splunk account representative for help.

### ***Restoration Summary***

In CMC, select the **Archive Management** link in the first panel of the **Storage Summary** or **Archive Storage (DDAA)** dashboard, then select the **Restoration Summary** tab.

The information in this tab shows the restoration activity for all of your deployment's indexes that are enabled with the DDAA feature. These totals in GB show the amount of uncompressed (raw) data in the following categories:

- **Restored:** Copied archive data that has been temporarily restored to an index. Restored data expires from searchable storage after 30 days.
- **Cleared:** Restored data that has been manually removed from an index. This data has a Jobstatus of Cleared.
- **Expired:** Data that has been automatically removed from searchable storage as it has passed the 30-day retention period. This data has a Jobstatus of Expired.

The displayed totals depend on the data you have selected to restore or clear and also the conditions and limitations of the restoration process, as follows:

- The archival and restoration process is complete.
- The data doesn't overlap with other data.
- The data size doesn't cause performance issues.

For more information, see the following in the the Splunk Cloud Platform *Admin Manual*:

- [Restore archived data to Splunk Cloud Platform](#)
- Step 7 of [Steps to restore data to Splunk Cloud Platform](#) for the different types of Jobstatus

### ***Interpret these results***

Review these totals and determine if the amount of data restored, cleared, and expired in your deployment meets or exceeds your organization's actual requirements. For example, a high total for restored data or low total for cleared or expired data may indicate the need to re-evaluate your index management policies and procedures. Ensure that you are restoring and retaining only the data that your organization truly needs.

## See also

| For more information about                 | See                                                                                                   |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Managing your aged ingested data with DDAA | <a href="#">Store expired Splunk Cloud Platform data to a Splunk-managed archive</a>                  |
| Managing indexes                           | <a href="#">Manage Splunk Cloud Platform indexes</a> in the Splunk Cloud Platform <i>Admin Manual</i> |

## Use the Forwarder dashboards

The dashboards accessed from the **Cloud Monitoring Console > Forwarder** tab provide information to Splunk Cloud Platform administrators about forwarder connections and status. This information helps you ensure your forwarders are correctly transmitting data to the indexers.

For data to appear on the forwarder dashboards, you must first configure and enable the [Forwarder Monitoring Setup page](#).

A blue progress bar might appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify any Cloud Monitoring Console (CMC) dashboard. Changing any of the search criteria, formatting, or layouts may cause inaccurate results and also override the automatic update process.

## Manage the forwarder monitoring setup

The CMC **Forwarder Monitoring Setup** page helps Splunk Cloud Platform administrators manage your forwarder monitoring configuration. This includes periodically removing decommissioned forwarders to improve system performance.

Because they are configuration pages, the **Forwarder Monitoring Setup** pages for Splunk Cloud Platform CMC and Splunk Enterprise Monitoring Console are similar. For more information on understanding and using this configuration page, see About time settings and Rebuild the forwarder asset table in the *Monitoring Splunk Enterprise* manual.

A difference between Splunk Cloud Platform CMC and the Splunk Enterprise Monitoring Console is the lookup file name. For CMC, enabling forwarder monitoring runs a scheduled search that populates the `sim_forwarder_assets.csv.gz` lookup file.

### Review the Forwarder Monitoring Setup page

To investigate this page, go to **Cloud Monitoring Console > Forwarders > Forwarder Monitoring Setup**.

After upgrading to CMC version 3.22.0, select "Rebuild forwarder assets" to ensure the data displays accurately.

The top section of the page is where you set whether forwarder monitoring is enabled and the data collection time interval, or disable it. Complete the following steps to enable forwarder monitoring.

1. Select **Enable**.
2. Choose a time option in **Data Collection Interval**.
3. Select **Save**.

4. Choose an option in the the **Build Forwarder Assets Now** dialog box that appears.
  1. Select **Continue** to start the forwarder assets rebuild process. This process lets you immediately rebuild the forwarders assets table, which removes decommissioned forwarders from the deployment and improves performance. Messages appear that indicate the state of this process and when it completes.
  2. Select **Later** if you want the forwarder assets table automatically rebuilt during the next daily update process.

Select **Reset** to reset the forwarder monitoring setup back to the previous configuration.

Be sure to select Save after making any configuration changes.

The bottom section of the page lets you immediately rebuild the forwarder assets table to remove any decommissioned forwarders. Complete the following steps.

1. Select **Rebuild forwarder assets...**
2. Choose an option in **Time Range**.
3. Select **Start Rebuild**.

Depending on the number of forwarders in your deployment, rebuilding the forwarder assets table can affect indexer performance and take a significant amount of time to complete.

## Monitor forwarder instances

The CMC Forwarders: Instance dashboard provides information to Splunk Cloud Platform administrators about the status and health of the forwarders in your deployment.

### *Review the Forwarders: Instance dashboard*

This dashboard contains two panels with tabular and graphical data for a specified forwarder instance. Set a time range to filter the results.

To investigate your panels, go to **Cloud Monitoring Console > Forwarders > Forwarders: Instance**. Use the following table to understand the dashboard interface.

| Panel or Filter          | Description                                                                                                                                                                                                                                                                                                                                   |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance and Time Range  | Specify a forwarder instance and a time range. These settings apply to both panels in the dashboard. <b>Note:</b> When you view this dashboard, the Instance field is automatically populated with the first menu value. Be sure to change this default value to the forwarder instance you are investigating.                                |
| Status and Configuration | Lists the following information for the specified forwarder: <ul style="list-style-type: none"> <li>• GUID</li> <li>• Forwarder type</li> <li>• IP address</li> <li>• Splunk version</li> <li>• OS and architecture</li> <li>• <b>Receiver</b> and connection counts</li> <li>• Average kilobytes per second and events per second</li> </ul> |
| Outgoing Data Rate       | Shows a graph that compares events per second and KB per second processed by the instance over the selected time range. Select an <b>Aggregation</b> value of either Maximum or Average.                                                                                                                                                      |

## Interpret forwarder instance results

When interpreting your forwarder instance results, note the following:

- Check that your forwarder's version is up-to-date.
- Use the IP address information to identify any faulty receivers in your local network.
- Compare the receiver count against the number of deployed indexers. A significant difference in these numbers indicates that there is likely a misconfiguration in the system.
- Review the graph in the **Outgoing Data Rate** panel and ensure that the forwarder is emitting data within its normal expected range. In particular, check the rates for average KB per second and events per second against their historical average rates. A rate that is significantly different from this historical rate, such as being very high or very low, could indicate an issue on the forwarding host.

## Monitor forwarder deployments

The CMC Forwarders: Deployment dashboard provides comprehensive information to Splunk Cloud Platform administrators about the status and health of the forwarders in your deployment. You can also set alerts that trigger if a forwarder is missing from the deployment.

### Review the Forwarders: Deployment dashboard

This dashboard shows both current status and historical information for your forwarder deployments, with various filters so you can further refine the results. Use the top panel to enable or disable missing forwarder alerts.

This dashboard contains one panel with a variable in the title: **Forwarders by <variable>**.

To investigate your panels, go to **Cloud Monitoring Console > Forwarders > Forwarders: Deployment**. Use the following table to understand the dashboard interface.

| Panel or Filter          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Missing Forwarder Alerts | <p>Select <b>enable</b> to open this panel.</p> <p>Specify a <b>Filter by Last:</b> option to view all missing forwarder alerts reported in that time range.</p> <p>Select the Scheduled Search: <b>SIM Alert - Missing Forwarders</b> link to access the <b>Searches, reports, and alerts</b> page. You can do the following for this alert:</p> <ul style="list-style-type: none"><li>• Confirm that the alert is successfully running every 15 minutes.</li><li>• Run the alert query on an ad hoc basis.</li><li>• View recently run jobs.</li></ul> <p>You can also manage this alert with the <a href="#">CMC Alerts panel</a>. For general information about managing alerts, see the Splunk Cloud Platform <i>Alerting Manual</i>.</p> |
| Forwarders by <variable> | <p>The &lt;variable&gt; in the panel title and the data in the pie chart graph dynamically change, based on the selected <b>Split by</b> option. The panel title is one of the following:</p> <ul style="list-style-type: none"><li>• <b>Forwarders by Status</b></li><li>• <b>Forwarders by Forwarder Type</b></li><li>• <b>Forwarders by Splunk Version</b></li><li>• <b>Forwarders by OS</b></li><li>• <b>Forwarders by Architecture</b></li></ul>                                                                                                                                                                                                                                                                                          |

| Panel or Filter                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                      | <b>Total: &lt;number&gt; forwarders</b> indicates the total number of forwarders in the deployment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Status and Configuration - As of <current_timestamp> | <p>Set criteria to filter the returned results:</p> <ul style="list-style-type: none"> <li>• The <b>Instance</b> filter accepts an asterisk (*) wildcard.</li> <li>• Specify a <b>Status</b> of All, Active, or Missing.</li> <li>• Select the <b>Show instances forwarding internal logs</b> checkbox to further refine the results.</li> </ul> <p><b>Total: &lt;number&gt;</b> on the left side of the table indicates the number of returned instances that meet the filter criteria. The table lists the following information:</p> <ul style="list-style-type: none"> <li>• Instance</li> <li>• Type</li> <li>• Version</li> <li>• OS</li> <li>• Architecture</li> <li>• Status</li> <li>• Last Connected to Indexers</li> <li>• Total KB</li> <li>• Average KB/s Over Time</li> <li>• Average KB/s</li> <li>• Average Events/s</li> </ul> |
| Historical Data                                      | This area includes the <b>Total Count of Forwarders</b> and <b>Forwarder Connection Count</b> panels. The specified <b>Time Range</b> option set here affects both panels. Specify an <b>Overlay</b> option to view a bar graph of the average KB per second or average events per second over time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

### ***Interpret forwarder deployment results***

Use this dashboard to identify misconfigurations or unhealthy behavior of the forwarders, such as outliers in the forwarder deployment. Misconfigurations means forwarders are sending too much or too little data. You also want to investigate any sudden spike of missing forwarders, as this could indicate a systemic failure.

### **Check forwarder versions**

The CMC Forwarder Versions dashboard shows the current installed version of Splunk Cloud Platform to Splunk Cloud Platform administrators and also indicates if your Splunk forwarders are outdated. Use this dashboard to determine which forwarders in your deployment are degrading its performance or have known compatibility issues with the deployed Splunk Cloud Platform version.

### ***Review the Forwarder Versions dashboard***

This dashboard provides four panels of information about your deployment and forwarders.

To investigate your panels, go to **Cloud Monitoring Console > Forwarders > Forwarder Versions**. Use the following table to understand the dashboard interface.

| Panel or Filter                       | Description                                                                                                                                                         |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Version Summary                       | Bar chart that shows forwarder version over forwarder count. The bars are color-coded to indicate if the forwarders are out-of-date (red) or up-to-date (green).    |
| Current Splunk Cloud Platform Version | Shows the version number of your current Splunk Cloud Platform deployment. This version number also appears in the <b>Support &amp; Services &gt; About</b> window. |
| Upgrade Recommendations               |                                                                                                                                                                     |

| Panel or Filter                       | Description                                                                                                                                                                |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                       | Shows upgrade recommendations based on comparing the forwarder version and the Splunk Cloud Platform version. Lists the forwarder name, version, type, and recommendation. |
| Flagged Forwarders (Based on Version) | Shows all forwarders that have been identified as broken or not operating as expected. Lists the forwarder name and version.                                               |

### **Interpret forwarder version results**

Use the CMC Forwarder Versions dashboard to identify which forwarders you must update as soon as possible. For more information, see Troubleshoot forwarder/receiver connection in the Splunk Cloud Platform *Forwarding Data* manual.

## **Use the Workload Management Monitoring dashboard**

The CMC Workload Management Monitoring dashboard shows the effect of triggered workload management rules, which is controlled by configuration settings. Use the dashboard to determine if you must adjust the underlying configuration to optimize workload performance. This dashboard appears only for deployments on Splunk Cloud Platform versions 8.x or higher.

This dashboard is driven by the workload management configuration that controls how Splunk Cloud Platform manages your workload performance. Go to **Settings > Workload management** to access the Workload Management page that contains the categories, pools, and rules defined by your configuration. You can use the workload management functionality to optimize search job processing and other tasks in your deployment.

An error message appears and the panels display 0 when the rules aren't configured or triggered.

For more information about Splunk Cloud Platform workload management, see Workload Management overview in the *Splunk Cloud Platform Admin Manual*. This topic describes how workload management works and how to use pools and rules to create a configuration that meets your organizational requirements.

A blue progress bar might appear above a panel, indicating that the Splunk platform is still generating data. Wait for the bar to disappear before reviewing the panel.

Do not modify any Cloud Monitoring Console (CMC) dashboard. Changing any of the search criteria, formatting, or layouts may cause inaccurate results and also override the automatic update process.

## **Review the Workload Management Monitoring dashboard**

This dashboard shows the following six panels of information under a filter section:

- The top three panels show the number of times that a search was affected by a workload rule or a particular rule classification.
- The middle two panels show graphs. Each panel's title and data dynamically changes, based on the selected **Split by** option. See the following table for panel titles shown for a specific option.
- The bottom panel shows tabular data for scheduled searches.

| Split by option | Panel titles |
|-----------------|--------------|
|                 |              |

|             |                                                                            |
|-------------|----------------------------------------------------------------------------|
| Action      | Searches per Action Triggered over Action<br>Searches per Action           |
| Rule        | Searches per Action Triggered over Rule<br>Searches per Rule               |
| User        | Searches per Action Triggered over User<br>Searches per User               |
| App         | Searches per Action Triggered over App<br>Searches per App                 |
| Search Type | Searches per Action Triggered over Search Type<br>Searches per Search Type |

### ***Investigate your panels***

1. Go to **Cloud Monitoring Console > Workload Management Monitoring**.
2. Refine the displayed data by specifying values for **Role**, **Time Range**, and **Split by**. The selected **Role** determines the information displayed in the top three panels. The specified **Time Range** and **Split by** values determine the information displayed in the bottom three panels.
3. Review the top three panels, which show a numerical value for the following classifications:
  - ◆ **Searches Aborted**
  - ◆ **Searches Reclassified**
  - ◆ **Searches Triggering an Alert**

These classifications are derived from the rules and actions set up in the workload configuration.
4. Review the **Searches per Action Triggered over <variable>** panel. See the table in [Review the Workload Management Monitoring dashboard](#) for panel titles.

This panel is a graph of the information shown in the top three panels. It compares the total number of searches that triggered a specific action against the selected **Split by** option. Use this information to analyze who or what is triggering the greatest number of actions so you can take the appropriate remediation steps.
5. Review the **Searches per <variable> over time** panel. See the table in [Review the Workload Management Monitoring dashboard](#) for panel titles.

This panel is a graph of the information shown in the top three panels. It shows the number of searches executed over time and color-coded by the selected **Split by** option. Use this information to analyze the peak days and times for a particular search and identify patterns of heavy usage. This helps you determine if you need to optimize resource allocations in your workload management configuration.
6. The **Scheduled Searches Triggering Rules Detail** panel shows how many searches assigned the `search_type` of `scheduled` have triggered an alert, or have been reassigned or aborted.

Use this information to monitor your scheduled searches, particularly those that were aborted or triggered an alert. Investigate why a search triggered either the abort or alert rule so you can correct any issues and rerun the search.

## **Monitor your deployment with the splunkd health report**

The splunkd health report is a REST-based monitoring tool that lets you view and investigate the health status of your deployment directly inside the Splunk Cloud Platform UI. Individual Splunk Cloud Platform features report their health



status through a tree structure that provides a continuous, real-time view of the health of your service, with no impact on search loads or ingest latency.

By default, the splunkd health report in Splunk Cloud Platform lets the `sc_admin` role view the health status of Search Scheduler features only. See [Supported features](#).

## Access the splunkd health report

To access the splunkd health report:

1. In Splunk Web, find the health report icon on the Splunk bar. The icon changes color from green to yellow or red, based on the health status of features in the health report.
2. Select the health report icon to open the health report.
3. In the health status tree, select any feature to view information about that feature's status.

**Health of Distributed Splunk Deployment** Local Distributed X

- ✓ splunkd
  - ✓ Search Scheduler
    - ✓ Search Lag
    - ✓ Searches Delayed
    - ✓ Searches Skipped In 1

**How to interpret this health report:**

This health report displays information from the `/health/splunkd/details` endpoint. There are three potential states for a feature:

- ✓ Green: The feature is functioning properly.
- ⚠ Yellow: The feature is experiencing a problem. The feature's status might automatically improve, or it might worsen over time. For details, see Root Cause.
- ! Red: The feature has severe issues and is negatively impacting the functionality of your deployment. For details, see Root Cause.
- ? Grey: Health report is disabled or snoozed for the feature.

To manage red and yellow threshold values for the individual features, go to [Health Report Manager](#).

For more information on this health report, see [Learn more](#).

## How the splunkd health report works

The health report records the health status of Splunk Cloud Platform features in a tree structure, where leaf nodes represent particular features, and intermediary nodes categorize the various features. Feature health status is color-coded in four states as follows:

- Green: The feature is functioning properly.
- Yellow: The feature is experiencing a problem.
- Red: The feature has a severe issue and is negatively impacting the functionality of your deployment.
- Grey: Health report is disabled for the feature.

### *The health status tree structure*

The health status tree has the following nodes:

| Health status tree node   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>splunkd</b>            | The top level node of the status tree shows the overall health status (color) of <code>splunkd</code> . The status of <code>splunkd</code> shows the least healthy state present in the tree. The REST endpoint retrieves the instance health from the <code>splunkd</code> node.                                                                                                                                                                                                                                          |
| <b>Feature categories</b> | Feature categories represent the second level in the health status tree. Feature categories are logical groupings of features. For example, "Search Lag", "Searches Delayed", and "Searches Skipped" are features that form a logical grouping with the name "Search Scheduler". Feature categories act as buckets for groups and reflect the status of the least healthy feature within the category. For example, if the health status of the "Search Lag" feature is red, the "Search Scheduler" category displays red. |
| <b>Features</b>           | The next level in the status tree is feature nodes. Each node contains information on the health status of a particular feature. Each feature contains one or more indicators that determine the status of the feature. The overall health status of a feature is based on the least healthy color of any of its indicators.                                                                                                                                                                                               |
| <b>Indicators</b>         | Indicators are the fundamental elements of the <code>splunkd</code> health report. These are the lowest levels of functionality that are tracked by each feature, and change colors as functionality changes. Indicator values are measured against red or yellow threshold values to determine the status of the feature. See <a href="#">What determines the status of a feature?</a>                                                                                                                                    |

### *What determines the health status of a feature?*

The health status of a feature depends on the current value of its associated indicators. For example, the status of the Search Scheduler: Search Skipped feature depends on the following two indicators:

- `percent_searches_skipped_high_priority_last_24h`
- `percent_searches_skipped_non_high_priority_last_24h`

Each indicator has a configurable threshold for yellow and red. When an indicator's value meets the threshold condition, the feature's status changes from green to yellow or yellow to red.

For instructions on configuring indicator thresholds, see [Edit feature indicator thresholds](#).

## View the splunkd health report

The `splunkd` health report in Splunk Web provides two options for viewing the health status of your deployment: a local health report view and a distributed health report view.

## Local health report view

The local health report view shows the health status of your deployment from the viewpoint of the local instance on which you are monitoring. The local health report is the default view.

## Distributed health report view

The distributed health report view shows the health status of features across a distributed deployment. The distributed health report aggregates health status information from connected instances on a single central instance. In Splunk Cloud Platform, search heads and search head cluster members act as central instances. The distributed health report option appears only on central instances of a distributed deployment.

The distributed health report is enabled by default on all search heads and search head cluster members. You can disable individual features in either the local health report or distributed health report to prevent the feature from reporting health status information to the splunkd health status tree. For more information, see [Disable a health report feature](#).

## Configure the splunkd health report

The splunkd health report displays the status of a pre-defined set of Splunk Cloud Platform features. You can modify some health report settings, including feature indicator thresholds, using the health report manager page in Splunk Web.

### Supported features

The splunkd health report lets the `sc_admin` role monitor these features by default:

| Feature Category | Features                                       |
|------------------|------------------------------------------------|
| Search Scheduler | Searches Skipped, Searches Delayed, Search Lag |

To view and edit thresholds for Search Scheduler features in the splunkd health report, a role must have the `list_health_subset` and `edit_health_subset` capabilities. The `sc_admin` role has these capabilities by default.

For information on additional health report features, see Supported features in the *Monitoring Splunk Enterprise* manual.

### Edit feature indicator thresholds

Each feature in the health status tree has one or more indicators. Each indicator reports a value against a pre-set threshold, which determines the status of the feature. When the indicator value meets the threshold condition, the health status of the feature changes, for example, from green to yellow, or yellow to red. Changing threshold values for any feature applies to all associated search heads or search head cluster members.

You can edit the threshold value for any feature indicator using Splunk Web, as follows:

1. In Splunk Web, select **Settings > Health report manager**.
2. Find the feature you want to modify and select **Edit Thresholds**.  
The Edit Threshold modal opens showing a detailed description of each feature indicator.
3. Set new indicator threshold values. For example, to modify thresholds for the Search Scheduler: Searches Skipped feature, you can set new Red or Yellow threshold values for the `percent_searches_skipped_high_priority_last_24h` and `percent_searches_skipped_non_high_priority_last_24h` indicators:

## Edit Thresholds for Search Scheduler Searches Skipped



|             |                                                                                                                                                                                                                                                                                       |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Indicator   | percent_searches_skipped_high_priority_last_24h                                                                                                                                                                                                                                       |
| Description | This indicator tracks the skip rate for high priority scheduled searches. These are scheduled searches where the priority field is set to "higher" or "highest". By default, this indicator is yellow if the skipped search ratio over the last 24 hours is 5%, and red if it is 10%. |
| Red         | <input type="text" value="10"/> ^<br>v                                                                                                                                                                                                                                                |
| Yellow      | <input type="text" value="5"/> ^<br>v                                                                                                                                                                                                                                                 |

4. Select **Save**.

### **Disable a health report feature**

You can disable any feature in the local or distributed health report using Splunk Web. Disabling a feature stops the feature from reporting health status information to the splunkd health status tree, which can be useful for removing noisy or irrelevant features from the health report view.

To disable a feature in the local health report using Splunk Web:

1. Log in to the local instance on which you want to disable the feature.
2. In Splunk Web, select **Settings > Health report manager**.
3. Find the feature you want to disable and select **Edit**.
4. In the modal, set the **Include in local report** switch to Enabled or Disabled.

The feature is disabled and no longer impacts the overall health status of `splunkd`. The feature now appears greyed out in the local health report.

To disable a feature in the the distributed health report using Splunk Web:

1. Log in to the central instance of the distributed health report, such as the cluster manager, search head, or search head cluster.
2. In Splunk Web, select **Settings > Health report manager**.
3. Find the feature you want to disable, and select **Edit**.
4. In the modal, set the **Include in distributed report** switch to Enabled or Disabled.

The feature is disabled and no longer impacts the overall health status of `splunkd`. The feature now appears greyed out in the distributed health report.

To temporarily disable alerts for a feature, select **Snooze** and specify a time duration. At the end of the specified duration alerts for the feature are re-enabled.


## Example: Investigate search scheduler health status changes

The splunkd health report lets you view the current health status of Search Scheduler features, including Searches Skipped, Searches Delayed, and Search Lag. You can use the report to identify and investigate Search Scheduler issues that can impact search performance.

The following example shows how you can use the splunkd health report to investigate Search Scheduler health status changes.

### 1. Check the health report status

1. In Splunk Web, check the color of the health report icon in the main menu. A red or yellow icon indicates that one or more search scheduler features have a problem.
2. Select the health report icon to open the health report. The following health report indicates that the Skipped Searches feature has a severe problem, and that the Search Lag feature might also have a problem.



### 2. Examine root cause and related messages

1. Select the Searches Skipped feature to view diagnostic information about the current health status of the feature.
2. Review the information under **Root Cause**. In this case, the percentage of high priority searches skipped is 44% over the last 24 hours, which exceeds the red threshold of 10% and causes the feature's health status to change to red.
3. Review the **Last 50 related messages**. These log entries include warning messages showing that some scheduled searches cannot be executed. For example:

```
09-15-2020 16:11:00.324 +0000 WARN SavedSplunker - cannot execute scheduled searches that live at the system level (need an app context).
```

Among explanations for this type of warning message is the possibility that the number of high-priority searches running exceeds the maximum concurrent search limit, which can cause searches to be skipped.

### 3. Confirm the cause of feature status change

After you review root cause and log file information, which suggest that maximum search concurrency limits caused the Searches Skipped feature's status change, you can use the Cloud Monitoring Console to check search scheduler activity and confirm if the suspected cause is correct.

1. In Splunk Web, select **Apps > Cloud Monitoring Console**.
2. Select **Search > Scheduler Activity**.

The **Count of Scheduler Executions** panel shows that 43.62 % of searches have been skipped over the last 4 hours, which approximates the percentage of skipped searches reported under root cause in the health report.

Total: 1733

| Status ↕  | Count ↕ | Percent of Total ↕ |
|-----------|---------|--------------------|
| completed | 977     | 56.38 %            |
| skipped   | 756     | 43.62 %            |

### 3. Select **Search > Skipped Scheduled Searches**.

The **Count of Skipped Scheduled Searches** panel shows that 756 searches have been skipped over the last 4 hours because "The Maximum number of concurrent historical searches on this instance has been reached." This confirms that the cause of the `Skipped Searches` status change is that the maximum concurrent search limit has been reached on the system.

| Reason ↕                                                                                         | Count ↕ | Percent of Total ↕ |
|--------------------------------------------------------------------------------------------------|---------|--------------------|
| The maximum number of concurrent historical scheduled searches on this instance has been reached | 778     | 100.00 %           |

4. You can now take steps to remedy this issue, by decreasing the total number of concurrent scheduled searches running, and increasing the relative concurrency limit for scheduled searches, which can bring the number of concurrent searches below the maximum concurrent search limit, and return the Searches Skipped feature to the green state.

For information on relative concurrency limits for scheduled searches, see [Set limits for concurrent scheduled searches](#).

## How Splunk monitors Splunk Cloud Platform

The Splunk Cloud Service Level Schedule describes Splunk's service level commitment for Splunk Cloud Platform. This topic describes some of the monitoring efforts that Splunk performs in support of that service level commitment. Splunk monitors the service with the following goals:

- Detect issues
- Restore service as quickly as possible
- Keep customers and their stakeholders informed about outages

Splunk Cloud Platform is monitored 24x7 worldwide by our Network Operations Center (NOC). During U.S. business hours, specialized teams work to resolve and identify causes for novel issues.

### Splunk Network Operations Center

The NOC takes action in response to automated alerts. For consistency and repeatability, the NOC uses runbooks to respond to alerts and files proactive incidents when novel issues occur.

The NOC manages more than 100 priority-one automated alerts that monitor the following components in particular. See the Splunk Cloud Platform Monitoring Matrix for more detail.

- Disk usage
- Indexers, search heads, cluster manager, KV store, or Inputs Data Managers (IDMs) down

- User Interface unresponsive
- Search head synchronization issues

## Specialized Teams

Specialized teams monitor critical product functionality and resolve issues during U.S. business hours. These specialized teams investigate to determine and remediate root causes and feed back into the development process to improve code resilience.

Specialized teams work on critical functions such as the following:

- Search
- Ingest
- Login

See Splunk Cloud Platform Service Details for more information.

## Splunk Cloud Platform Monitoring Matrix

The following table lists Splunk Cloud features that are monitored and the Splunk response when issues are detected. This is a representative list of Splunk Cloud Platform monitoring features. It is not exhaustive and is subject to change without notice. This document does not describe Splunk Cloud Platform for FedRAMP Moderate, Splunk Cloud Platform for FedRAMP High, or Splunk Cloud Platform for DoD IL5.

| Feature                                      | Issue                                      | Support             | Splunk Action                                                                                                                               |
|----------------------------------------------|--------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Federated Search</b>                      | Federated search issues                    | U.S. business hours | Investigate cause.                                                                                                                          |
| <b>Inputs Data Manager (IDM)</b><br>(Note 1) | IDM requires upsizing                      | 24x7                | Schedule maintenance window to upsize IDM.                                                                                                  |
| <b>Indexing</b>                              | Indexer down                               | 24x7                | Check bundle push and detention status. Verify available disk space and potentially restart service. Create proactive incident if required. |
|                                              | Indexing latency >5 minutes (Note 2)       | U.S. business hours | Investigate cause.                                                                                                                          |
|                                              | Indexing queues blocked                    | U.S. business hours | Investigate cause.                                                                                                                          |
| <b>Infrastructure</b>                        | Disk space full                            | 24x7                | Rotate logs to clear old backups or expand disk space (Note 3). Create proactive incident if required.                                      |
| <b>Ingestion</b>                             | Splunk-to-Splunk (S2S) ingestion port down | 24x7                | Check bundle push and detention status. Verify available disk space and potentially restart service. Create proactive incident if required. |
|                                              | Ingestion HTTP Event Collector             | U.S. business hours | Investigate cause.                                                                                                                          |
|                                              | Ingestion S2S connection acceptance        | U.S. business hours | Investigate cause.                                                                                                                          |
| <b>KV store</b>                              | KV store down                              | 24x7                | Check data store health, certificates, disk space, and potentially restart service. Create proactive incident if required.                  |
| <b>Login</b>                                 | Splunk native authentication               | U.S. business hours | Investigate cause.                                                                                                                          |
|                                              | Identity provider authentication           | U.S. business hours | Investigate cause.                                                                                                                          |

| Feature       | Issue                                                                                  | Support             | Splunk Action                                                                                                                                                                          |
|---------------|----------------------------------------------------------------------------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Search</b> | Search Head Cluster (SHC) out of sync                                                  | 24x7                | Check knowledge object replication and potentially re-sync cluster members. Create proactive incident if required.                                                                     |
|               | Search peer isolated                                                                   | 24x7                | Check for unavailable or stuck search peers, potentially restart service or remove unresponsive peer. Create proactive incident if required.                                           |
|               | Search initiation                                                                      | 24x7                | Check health of the indexer running searches, bundle synchronization, down peers, and cluster manager health, and potentially restart services. Create proactive incident if required. |
|               | Search execution                                                                       | U.S. business hours | Investigate cause.                                                                                                                                                                     |
|               | Search performance                                                                     | U.S. business hours | Note search performance reductions relative to customer's historical performance. Investigate cause.                                                                                   |
|               | Skipped search percentage                                                              | U.S. business hours | Investigate cause.                                                                                                                                                                     |
|               | API unavailable                                                                        | 24x7                | Check for system overload and disk space issues, and potentially restart the service. Create proactive incident if required.                                                           |
|               | Splunk Web user interface unavailable (Search Head or Enterprise Security Search Head) | 24x7                | Check certificates and potentially restart processes or instances. Create proactive incident if required.                                                                              |

**Notes**

1. IDM applies to Classic Experience only.
2. Indexing latency applies to Victoria Experience only.
3. Disk expansion limited by entitlement.

## Manage your Splunk Cloud Platform capacity

You can reallocate or adjust your system capacity to suit your organization's needs with the Splunk Cloud Platform Cloud Flex program. Cloud Flex is an Early Access release program.

In the Early Access release stage, Splunk products may have limitations on customer access, features, maturity and regional availability. For additional information about the Cloud Flex Early Access program, contact your Splunk representative.



# Optimize indexing and search processes

## Optimize indexing and search processes

Optimizing search and indexing processes can improve your system performance and Splunk Virtual Compute (SVC) utilization. Because SVC usage is based on processes performed by the search heads and indexers, optimizing these processes for efficiency can positively impact on your SVC usage.

However, SVC usage is not a direct measurement of the health and performance of your deployment. Improving a search or indexing process might not decrease your SVC usage but could improve your system performance. For a better understanding of your system health, see [Use the Health dashboard](#) in the *Splunk Cloud Platform Admin Manual*.

To learn more about SVCs, how you can monitor them using the Cloud Monitoring Console (CMC), and the workload pricing model, see the following documentation:

- [Monitor current SVC usage of your workload-based subscription](#) in the *Splunk Cloud Platform Admin Manual*.
- Performance considerations in the *Splunk Cloud Platform Service Description*.

The following tips and resources can help you improve search and indexing processes and potentially improve SVC usage and system performance.

## Optimize search processes

The following are ways you can optimize search processes so that they're more resource efficient:

| Method                                 | Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Review data models                     | <p>You can use the Common Information Model (CIM) add-on, which contains preconfigured data models that can accelerate key data. Turn on data acceleration and use CIM filters to exclude data from searches so that your searches use less resources. Make sure to include index definitions to reduce the data scanned during data model acceleration.</p> <p>See the following documentation from the <i>Common Information Model Add-on Manual</i>:</p> <ul style="list-style-type: none"><li>• <a href="#">Overview of the Splunk Common Information Model</a></li><li>• <a href="#">Accelerate CIM data models</a></li><li>• <a href="#">Use the CIM Filters to exclude data</a></li></ul> |
| Review skipped searches                | <p>Get more details on skipped searches using the following CMC dashboards in the <i>Splunk Cloud Platform Admin Manual</i>:</p> <ul style="list-style-type: none"><li>• Review the <a href="#">Health dashboard</a></li><li>• Review the <a href="#">Skipped Scheduled Searches dashboard</a></li></ul> <p>See the following resources to learn more about reducing skipped searches:</p> <ul style="list-style-type: none"><li>• <a href="#">Splunk Blogs post: Are You Skipping? Please Read</a></li><li>• <a href="#">Splunk Lantern article: Reducing skipped searches</a></li></ul>                                                                                                        |
| Review searches that run over all time |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Method                                                                               | Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                      | Searches that run over all time might use many resources, especially if they're event searches without tokens or indexed fields that filter the data. However, some searches that run over all time, such as API calls, don't use a lot of resources.                                                                                                                                                                                                                                               |
| Review long time running searches and optimize SPL                                   | <p>Improve your searches so that they're less resource intensive. Prioritize improving the most expensive searches. See the following documentation to learn more:</p> <ul style="list-style-type: none"> <li>• <a href="#">Analyze expensive searches</a> in the <i>Splunk Cloud Platform Admin Manual</i>.</li> <li>• About search optimization in the <i>Splunk Cloud Platform Search Manual</i>.</li> <li>• Write better searches in the <i>Splunk Cloud Platform Search Manual</i>.</li> </ul> |
| Turn off unused scheduled searches, report acceleration, and data model acceleration | <p>Unused scheduled searches, report acceleration, and data model acceleration take up resources unnecessarily. This is especially true for out-of-the-box saved searches and accelerations.</p> <p>You can use the Splunk app for Redundant or Inefficient Search Spotting to identify redundant searches.</p>                                                                                                                                                                                     |
| Remove unused apps and technical add-ons (TAs)                                       | Unused apps and TAs take up resources unnecessarily. This is especially true if you have unused CIM data models, out-of-the-box saved searches, and accelerations.                                                                                                                                                                                                                                                                                                                                  |

## Optimize indexing processes

You can improve indexing processes by investigating data quality issues, and following HTTP Event Collector (HEC) best practices.

| Method                                             | Details                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Investigate data quality issues                    | <p>Review the CMC Data Quality dashboard and see <a href="#">Verify data quality</a> in the <i>Splunk Cloud Platform Admin Manual</i> to investigate data quality issues.</p> <p>Address line breaking, event breaking, and time stamp issues to improve data quality. See the following Splunk Lantern articles to learn more:</p> <ul style="list-style-type: none"> <li>• Solving data quality issues</li> <li>• Configuring new source types</li> </ul> |
| Review your HTTP Event Collector (HEC) performance | To gain more insight on your HEC status, review the CMC HTTP Event Collector (HEC) dashboard and see <a href="#">Check the status of HTTP event collection</a> in the <i>Splunk Cloud Platform Admin Manual</i> .                                                                                                                                                                                                                                           |

# Manage your Indexes and Data in Splunk Cloud Platform

## Manage Splunk Cloud Platform indexes

Splunk Cloud Platform administrators create indexes to organize data, [apply role-based access permissions to indexes that contain relevant user data](#), fine-tune data, specify how long to retain data in indexes, and so on.

Indexes store the data you have sent to your Splunk Cloud Platform deployment. To manage indexes, Splunk Cloud Platform administrators can perform these tasks:

- Create, update, delete, and view properties of indexes.
- Monitor the size of data in the indexes to remain within the limits of a data plan or to identify a need to increase the data plan.
- Modify data retention settings for individual indexes to control when Splunk Cloud Platform automatically deletes data or moves it to storage.
- Optimize search performance by managing the number of indexes and the data sources that are stored in specific indexes.
- Delete indexes. **Caution:** This function deletes all data from an index and removes the index. The operation is final and can't be reversed.
- Move expired data from indexes to self storage or a Splunk-supported archive (Dynamic Data Active Archive). Data from the index is not deleted until it is successfully moved to the storage location. Archived data can be restored to Splunk Cloud Platform for searching. Data from a self storage location can no longer be searched from Splunk Cloud Platform. However, it can be restored to a Splunk Enterprise instance for searching if necessary.

## Event indexes and metrics indexes

You can create two types of indexes:

- Events indexes, for event-based log data. Events indexes impose minimal structure and can accommodate any kind of data, including metrics data.
- Metrics indexes, for metric data points. Metrics indexes use a highly structured format to handle the higher volume and lower latency demands associated with metrics data. Putting metrics data into metrics indexes results in faster performance and less use of index storage, compared to putting the same data into events indexes.

Events indexes are the default index type. To create events indexes, see [Create a Splunk Cloud Platform events index](#).

To create metrics indexes, see [Create a Splunk Cloud Platform metrics index](#).

For more information about the metrics data format see *Metrics*.

## Best practices for creating indexes

Consider these best practices when creating indexes:

- Create separate indexes for long-term and short-term data. For example, you might need to keep security logs for one year but web access logs for only one month. Using separate indexes, you can set different data retention times for each type of data.

- Apply logical or role-based boundaries for indexes. For example, create separate indexes for different departments.
- Devise a naming convention to easily track, navigate, and organize indexes.
- To configure your data retention settings, see the best practice listed here: [Manage Data Retention Settings](#).

## The Indexes page

To view the **Indexes** page, select **Settings > Indexes**. The **Indexes** page lists the indexes in a Splunk Cloud Platform deployment and lets administrators create, update, delete, and modify the properties of indexes. To modify settings for an index, click its name.

From this page you can:

- Create an index.
- View index details such as the following.
  - ◆ **Index name:** The name specified when the index was created.
  - ◆ **Index type:** Whether the index is an events index or a metrics index.
  - ◆ **App:** The app to which the index belongs.
  - ◆ **Current size:** The approximate amount of uncompressed raw data currently stored in the index.
  - ◆ **Max size:** The maximum amount of uncompressed raw data (in TB, GB, or MB) that can be retained in the index.
  - ◆ **Event count:** The number of events in the index.
  - ◆ **Earliest event:** The time of the earliest event found in the index.
  - ◆ **Latest event:** The time of the most recent event found in the index.
  - ◆ **Searchable Retention:** The maximum age of events retained in the index.
  - ◆ **Storage Type:** The storage settings for expired data from a given index. Can be self storage, archive, or no additional storage.
  - ◆ **Status:** Enabled or disabled. Data in a disabled index is ignored in searches.
  - ◆ Delete an index. **Caution:** Deletes all data from an index and removes the index. The operation is final and can't be reversed.

## Create a Splunk Cloud Platform events index

To create an events index:

1. Select **Settings > Indexes**.
2. Click **New Index**
3. In the **Index name** field, specify a unique name for the index. Index names can contain only lowercase letters, numbers, underscores, or hyphens. They must begin with a lowercase letter or number.
4. Set **Index Data Type** to **Events**.
5. In the **Max raw data size** field, specify the maximum amount of raw data allowed before data is removed from the index. Set this value to zero to specify an unlimited maximum raw data size. This is a data retention setting.
6. In the **Searchable time (days)** field, specify the number of days before an event is removed from an index. This is a data retention setting.
7. In the **Dynamic Data Storage** field, select **Splunk Archive** to send data to the Splunk Dynamic Data Active Archive, or choose **Self Storage** to move expired data to your own self-storage area. If you don't want to maintain expired Splunk data, leave **No additional storage** selected.
8. If you enabled data self storage, select a location for data self storage. Or, click **Edit self storage locations** to add a new self storage location. For more information about data self storage and instructions for configuring a data self storage location, see [Manage your Indexes and Data in Splunk Cloud](#).

9. If you enabled Dynamic Data Active Archive, configure retention settings for the archive. For more information, see [Archive expired Splunk Cloud Platform data](#).
10. Click **Save**.
11. Required step for [Classic Experience](#) customers: If this new index must be available to data collection apps on your IDM, contact Splunk Support and request they sync the index with your IDM. This ensures communication between the new index and any data collection apps running on the IDM. If you have a support contract, log in and file a new case using the Splunk Support Portal. Otherwise, contact Splunk Customer Support.

The events index appears after you refresh the page. Retention settings are applied to individual indexes, and data retention policy settings apply to all of the data that is stored in your Splunk Cloud deployment. Monitor and verify that the data retention settings for all indexes do not meet or exceed the values set in the data retention policy. For more information, see [Data retention](#).

## Create a Splunk Cloud Platform metrics index

To create an metrics index:

1. Select **Settings > Indexes**.
2. Click **New Index**
3. In the **Index name** field, specify a unique name for the index. Names must begin with a lowercase letter or a number and can include uppercase letters, hyphens, and underscores.
4. Set **Index Data Type** to **Metrics**.
5. (Optional) Set **Timestamp Resolution** to **Milliseconds** if you want the metrics index to store metric data points at that increased level of granularity. Metrics indexes with millisecond timestamp resolution have decreased search performance. See [Metrics indexes with millisecond timestamps](#).
6. In the **Max raw data size** field, specify the maximum amount of raw data allowed before data is removed from the index. Set this value to zero to specify an unlimited maximum raw data size. This is a data retention setting.
7. In the **Searchable time (days)** field, specify the number of days before an event is removed from an index. This is a data retention setting.
8. In the **Dynamic Data Storage** field, select **Splunk Archive** to send data to the Splunk Dynamic Data Active Archive, or choose **Self Storage** to move expired data to your own self-storage area. If you don't want to maintain expired Splunk data, leave **No additional storage** selected.
9. If you enabled data self storage, select a location for data self storage. Or, click **Edit self storage locations** to add a new self storage location. For more information about data self storage and instructions for configuring a data self storage location, see [Manage your Indexes and Data in Splunk Cloud](#).
10. If you enabled Dynamic Data Active Archive, configure retention settings for the archive. For more information, see [Archive expired Splunk Cloud Platform data](#).
11. Click **Save**.
12. Required step for [Classic Experience](#) customers: If this new index must be available to data collection apps on your IDM, contact Splunk Support and request they sync the index with your IDM. This ensures communication between the new index and any data collection apps running on the IDM. If you have a support contract, log in and file a new case using the Splunk Support Portal. Otherwise, contact Splunk Customer Support.

The metrics index appears after you refresh the page. Retention settings are applied to individual indexes, and data retention policy settings apply to all of the data that is stored in your Splunk Cloud deployment. Monitor and verify that the data retention settings for all indexes does not meet or exceed the values set in the data retention policy. For more information, see [Data retention](#).

### **Metrics indexes with millisecond timestamps**

By default, metrics indexes are only searchable at a second-by-second precision. This is unlike events indexes, which can be searched with subsecond precision by default.

If you are dealing with a high volume source of metric data, such as a utility grid that has the potential to generate millions of metric data points per second, this means that the metric index is populated with sample metric data points or metric data points that are aggregated views of the raw metric data, taken at regular intervals.

If you are concerned about high index volume, this can be a good thing. Having second precision metrics indexes keeps your indexes lean and saves you from having to search through huge numbers of events over relatively short time ranges. But this also means that you cannot run time-based metrics searches that have subsecond precision. Similarly, you cannot set up `mstats` searches that group by subsecond `span` values.

If you need the capability to perform metric searches with subsecond precision, give your new metric index a **Timestamp Resolution of Milliseconds**. Metrics indexes with millisecond timestamp resolution can have decreased search performance in comparison to metrics indexes that have the default second timestamp precision.

Metrics indexes set to millisecond precision might incur more license usage than similar metrics indexes set to second precision. The license cost per metric data point remains the same, but millisecond-precision indexes can index more data points than second-precision indexes ingesting data from the same source.

### **About changing timestamp resolutions of metrics indexes**

You can change the timestamp resolution of a metrics index after you create it. However, if you change the timestamp resolution of a metrics index from millisecond to second, it may look like data loss to people who regularly run searches against that metrics index. This is because the index won't ingest data at millisecond resolution after the change.

When your index is at millisecond timestamp resolution, your indexed metric data points might have timestamps like this.

| <b>_timestamp (seconds)</b> |
|-----------------------------|
| 1.000                       |
| 1.001                       |
| 1.002                       |
| 2.000                       |
| 2.435                       |
| 3.123                       |
| 3.651                       |
| 4.000                       |

After four seconds, if you change the timestamp resolution from millisecond timestamp resolution to second timestamp resolution, your index is restricted to indexing one metric data point per second:

| <b>_timestamp (seconds)</b> |
|-----------------------------|
| 5.000                       |
| 6.000                       |

| <b>_timestamp (seconds)</b> |
|-----------------------------|
| 7.000                       |
| 8.000                       |
| 9.000                       |

Some users may perceive this as a data loss when in fact they are just seeing their data with a less granular timestamp resolution.

Similarly, users of a metrics index that is switched from a second timestamp resolution to a millisecond timestamp resolution may be surprised to see their indexes ingesting more events than they did before the switch.

As an administrator of a Splunk Cloud Platform deployment it is up to you to communicate this change and its implications to your users.

## **Review Splunk Cloud Platform data policies**

Splunk Cloud Platform administers your data according to the policies described in the following sections.

### ***Data ingestion and daily license usage***

Splunk Cloud Platform administers your data based on your subscription type. For more information, see Data policies in the "Subscription types" section of the Splunk Cloud Platform Service Description.

### ***Data backup policy***

Splunk Cloud continuously maintains and monitors your deployment. For more information, see the "Ensures Splunk Cloud Platform uptime and security" section of Splunk maintenance responsibilities in the Splunk Cloud Platform Service Description.

### ***Data retention***

When you send data to Splunk Cloud Platform, it is stored in indexes. Splunk Cloud Platform retains data based on index settings that enable you to specify when data is to be deleted or moved to self storage. To configure different data retention settings for different sources of data, store the data in separate indexes according to the desired retention policy.

You can configure the number of days for data to be searchable by configuring the Searchable time (days) setting for an individual index. For more information, see [Manage data retention settings](#).

Index data is stored in directories called buckets. Data is deleted by deleting entire buckets, not individual events. When the maximum age or size of the Index is reached, buckets are deleted or moved starting with the oldest buckets first. Buckets are removed until the index no longer exceeds the configured limit. If you use data self storage or archiving, buckets are not deleted until the data is successfully moved to your self storage or archive location.

Data retention is based on your subscription type. For more details about data retention policies, see the Storage section of the Splunk Cloud Platform Service Description.

## **Manage data retention settings**

Each index uses two settings on the **New Index** page to determine when to delete data:

- The maximum size of the raw index data (MB, GB, or TB, specified in the **Max raw data size** field)
- The maximum age of events in the index (specified in the **Searchable time (days)** field)

When the index reaches the specified maximum size or events reach the specified maximum age, the oldest data is deleted or is moved to your self-storage location (depending on your configuration).

For example, the system ingests data from a particular datasource at an approximate rate of 10 GB per day, and you want to retain and search against the last 90 days worth of data. This would cause the system to normally ingest approximately 900 GB over the configured 90 day searchable time period.

Based on your search and data retention requirements, set these values so that the **Searchable time (days)** value is reached before the **Max raw data size** threshold is reached. A best practice is to set the maximum raw data size to a significantly larger value than the normal total ingestion amount. Doing this allows for unanticipated bursts of data that might otherwise cause the system to start deleting data before reaching the desired retention limit.

Given the above parameters, you might configure the retention settings as follows:

- **Max raw data size** set to 1800 GB (double the example 900 GB normal total ingestion amount)
- **Searchable time (days)** set to 90

### New Index ✕

---

Index name

Index Data Type 📄 Events 📊 Metrics  
The type of data to store (event-based or metrics).

Max raw data size  GB ▾  
Maximum aggregated size of raw data (uncompressed) contained in index. Set this to 0 if you want unlimited.

Searchable time (days)   
Number of days the data is searchable

Dynamic Data Storage  Splunk Archive [?](#)  
 Self Storage [?](#)  
 No Additional Storage  
Learn more about Dynamic Data Storage options.

---



These values together account for both your ingestion rate and the time you want to retain the data. You need to consider these factors for each index that you create.

The new data retention settings appear after you click **Save** and refresh the page.

Check your data retention in the [Cloud Monitoring Console](#) to ensure you estimated your ingestion rate correctly and your storage consumption is within your entitlement. If you did not correctly estimate your ingestion rate, you might have a shorter retention period than expected.

Splunk Cloud Platform administrators can specify the settings that determine when data is removed from a specific index. For more information, see the following:

- Store expired Splunk Cloud Platform data for information about data self storage and instructions for configuring a data self storage location.
- Archive expired Splunk Cloud Platform data for information about archiving data.

Splunk Cloud Platform includes several internal indexes that are named starting with an underscore (\_). The data retention period for these internal indexes cannot be modified.

## Delete index data and the index from Splunk Cloud Platform

Splunk Cloud Platform administrators can delete an index.

This function deletes all data from an index and removes the index. The operation is final and can't be reversed.

1. Select **Settings > Indexes**.
2. Identify the index and click **Delete** from the **Action** column.
3. Click **OK** to confirm that you want to delete the data and index from Splunk Cloud Platform.

The data and index are deleted from Splunk Cloud and can't be restored.

You can't delete default indexes and third-party indexes from the Indexes page.

## Store expired Splunk Cloud Platform data in your private archive

Dynamic Data Self Storage (DDSS) lets you move your data from your Splunk Cloud Platform indexes to a private storage location in your AWS or GCP environment. You can use DDSS to maintain access to older data that you might need for compliance purposes.

You can configure Splunk Cloud Platform to move data automatically from an index when the data reaches the end of the Splunk Cloud Platform retention period that you configure. You can also restore the data from your self storage location to a Splunk Enterprise instance. To ensure there is no data loss, DDSS maintains your data in the Splunk Cloud Platform environment until it is safely moved to your self storage location.

## Requirements for Dynamic Data Self Storage

To configure a self storage location in your AWS or GCP environment, you must have sufficient permissions to create Amazon S3 or Google Cloud Storage buckets and apply policies or permissions to them. If you are not the AWS or GCP

administrator for your organization, make sure that you can work with the appropriate administrator during this process to create the new bucket.

After you move the data to your self storage location, Splunk Cloud Platform does not maintain a copy of this data and does not provide access to manage the data in your self storage environment, so make sure that you understand how to maintain and monitor your data before moving it to the bucket.

If you intend to restore your data, you also need access to a Splunk Enterprise instance.

## Performance

DDSS is designed to retain your expired data with minimal performance impact. DDSS also ensures that the export rate does not spike in the case of a large volume of data. For example, if you reduce the retention period from one year to ninety days, the volume increases, but the export rate does not spike. This ensures that changes in data volume do not impact performance.

For more information on DDSS performance and limits, see [Service limits and constraints](#).

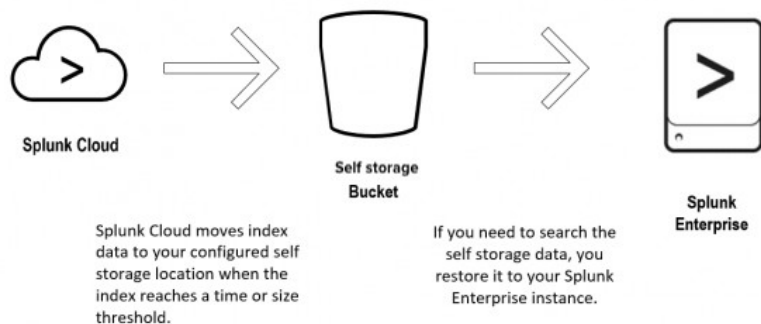
## How Dynamic Data Self Storage works

Splunk Cloud Platform moves data to your self storage location when the index meets a configured size or time threshold. Note the following:

- If an error occurs, a connection issue occurs, or a specified storage bucket is unavailable or full, Splunk Cloud Platform attempts to move the data every 15 minutes until it can successfully move the data.
- Splunk Cloud Platform does not delete data from the Splunk Cloud Platform environment until it has successfully moved the data to your self storage location.
- Data is encrypted by SSL during transit to your self storage location. Because Splunk Cloud Platform encryption applies only to data within Splunk buckets, you might want to encrypt data in the target bucket after transit. For Amazon S3 buckets, to ensure your data is protected, enable AES256 SSE-S3 on your target bucket so that data encryption resumes immediately upon arrival at the SSE-S3 bucket. Enabling AES256 SSE-S3 provides server-side encryption with Amazon S3 Managed keys (SSE-S3) only. This feature does not work with KMS keys. For GCP buckets, data encryption is enabled by default.

After Splunk Cloud Platform moves your data to your self storage location, you can maintain the data using your cloud provider's tools. If you need to restore the data so that it is searchable, you can restore the data to a Splunk Enterprise instance. The data is restored to a thawed directory, which exists outside of the thresholds for deletion you have configured on your Splunk Enterprise instance. You can then search the data and delete it when you finish.

When you restore data to a thawed directory on Splunk Enterprise, it does not count against the indexing license volume for the Splunk Enterprise or Splunk Cloud Platform deployment.



## Configure self storage locations

You can set up one or more buckets in Amazon S3 or Google Cloud Storage (GCS) to store your expired data. For configuration details, see:

- [Configure self storage in Amazon S3.](#)
- [Configure self storage in GCP.](#)

To manage self storage locations in Splunk Cloud Platform your role must hold the `indexes_edit` capability. The `sc_admin` role holds this capability by default. All self storage configuration changes are logged in the `audit.log` file.

For information on how to configure DDSS self storage locations programmatically without using Splunk Web, see Manage DDSS self storage locations in the Admin Config Service Manual.

## Configure self storage in Amazon S3

To configure a new self storage location in Amazon S3, you must create an S3 bucket in your AWS environment, and configure the S3 bucket as a new self storage location in the Splunk Cloud Platform UI. When you configure the S3 bucket as a new storage location, Splunk Cloud Platform generates a resource-based bucket policy that you must copy/paste to your S3 bucket to grant Splunk Cloud the required access permissions.

For information on how to create and manage Amazon S3 buckets, see the AWS documentation. For information on the differences between AWS identity-based policies and resource-based policies, see [https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_identity-vs-resource.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_identity-vs-resource.html)

### Create an Amazon S3 bucket in your AWS environment

When creating an Amazon S3 bucket, follow these important configuration guidelines:

- **Region:** You must provision your Amazon S3 bucket in the same region as your Splunk Cloud Platform environment.
- **Object Lock:** Do not activate AWS S3 Object Lock when creating a bucket. Locking the bucket prevents DDSS from moving data to the bucket. For more information, see <https://docs.aws.amazon.com/AmazonS3/latest/user-guide/object-lock.html>
- **Naming:** When you name the S3 bucket, it must include the Splunk prefix provided to you and displayed in the UI under the **AWS S3 bucket name** field. Enter the prefix *before* the rest of the bucket name. This prefix contains

your organization's Splunk Cloud ID, which is the first part of your organization's Splunk Cloud URL, and a 12-character string. The complete S3 bucket name has the following syntax:

```
Splunk Cloud ID-{12-character string}-{your bucket name}
```

For example, if you administer Splunk Cloud Platform for Buttercup Cloudworks, and your organization's Splunk Cloud URL is `buttercupcloudworks.splunkcloud.com`, then your Splunk Cloud ID is `buttercupcloudworks`. The image shows the following example prefix you'd see when configuring an S3 bucket using the New Self Storage Location dialog box:

```
buttercupcloudworks-rs73hfjie674-{your bucket name}
```

**New Self Storage Location** X

Configure a new storage location for your expired data. [Learn more](#)

Title

Description   
Optional

**Amazon Web Services (AWS) Configuration**

AWS S3 bucket name   
An AWS S3 bucket in the same region as your Splunk Cloud environment. The S3 bucket must have \*buttercupcloudworks-rs73hfjie674\* as the prefix in the name.

AWS S3 bucket folder

AWS S3 bucket path   
Path is <Bucket Name>/<Bucket Folder>

AWS S3 bucket policy

Copy and apply this bucket policy to your S3 bucket in the AWS Management Console. [Learn more](#)

Test bucket policy

If you do not use the correct prefix, Splunk cannot write to your bucket. By default, your Splunk Cloud Platform instance has a security policy applied which disallows write operations to S3 buckets that do not include your Splunk Cloud ID. This security policy allows the write operation only for those S3 buckets that you create for the purpose of storing your expired Splunk Cloud data.

If you have enabled AES256 SSE-S3 on your target bucket, the data will resume encryption at rest upon arrival at the SSE-S3 bucket.

### ***Configure a self storage location for the Amazon S3 bucket***

To configure your Amazon S3 bucket as a self storage location in Splunk Cloud Platform:

1. In Splunk Web, click **Settings > Indexes > New Index**.
2. In the **Dynamic Data Storage** field, click the radio button for **Self Storage**.
3. Click **Create a self storage location**.  
The Dynamic Data Self Storage page opens.
4. Give your location a **Title** and an optional **Description**.
5. In the **Amazon S3 bucket name** field, enter the name of the S3 bucket that you created.
6. (Optional) Enter the bucket folder name.
7. Click **Generate**. Splunk Cloud Platform generates a bucket policy.
8. Copy the bucket policy to your clipboard. Note: Customers with an SSE-S3 encrypted bucket must use the default policy and not modify the policy in any way.
9. In a separate window, navigate to your AWS Management console and apply this policy to the S3 bucket you created earlier.
10. In the Self Storage Locations dialog, click **Test**.  
Splunk Cloud writes a 0 KB test file to the root of your S3 bucket to verify that Splunk Cloud Platform has permissions to write to the bucket. A success message displays, and the **Submit** button is enabled.
11. Click **Submit**.
12. In the AWS Management Console, verify that the 0 KB test file appears in the root of your bucket.

You cannot edit or delete a self storage location after it is defined, so verify the name and description before you save it.

### **Configure self storage in GCP**

To configure a new self storage location in GCP, you must create a Google Cloud Storage (GCS) bucket in your GCP environment and configure the GCS bucket as a new self storage location in the Splunk Cloud Platform UI. For detailed information on how to create and manage GCS buckets, see the GCP documentation.

#### ***Create a GCS bucket in your GCP environment***

When creating a GCS bucket, follow these important configuration guidelines:

Bucket configurations that deviate from these configuration guidelines can incur unintentional GCS charges and interfere with DDSS successfully uploading objects to your GCS bucket.

- **Region:** You must provision your GCS buckets in the same GCP region as your Splunk Cloud Platform deployment. Your GCP region depends on your location. For more information, see Available regions.

- **Bucket lock/bucket retention policy:** Do not set a retention policy for your GCS bucket. The bucket lock/bucket retention policy feature is not compatible with the GCS parallel composite upload feature DDSS uses to transfer files to GCS buckets and can interfere with data upload. For more information on parallel composite uploads, see <https://cloud.google.com/storage/docs/parallel-composite-uploads>.
- **Default storage class:** Make sure to use the Standard default storage class when you create your GCS bucket. Using other default storage classes can incur unintentional GCS charges.
- **Permissions:** You must configure permissions for the 2 GCP service accounts associated with your Splunk Cloud Platform deployment. These service accounts are shown under **GCP service account** in the New Self Storage Location modal when you configure a new self storage location in Splunk Web.

To configure permissions for the GCP service accounts, you must assign the following predefined GCP roles to the 2 GCP service accounts using the GCP console:

```
Storage Legacy Bucket Writer
Storage Legacy Object Reader
```

For more information on GCP roles, see IAM roles for Cloud Storage in the GCP documentation.

- **Naming:** Your GCS bucket name must include the prefix that Splunk Cloud Platform provides and displays in the UI under the **GCP bucket name** field. The following image shows an example of this prefix. This prefix contains your Splunk Cloud Platform ID, which is the first part of your Splunk Cloud Platform URL, and a 4-character string. The complete GCS bucket name has the following syntax:

```
Splunk Cloud ID-{4-character string}-{your bucket name}
```

**New Self Storage Location** ×

Configure a new storage location for your expired data. [Learn more](#)

Title

Description  Optional

**Google Cloud Platform (GCP) Configuration**

GCP bucket name

A GCP bucket in the same region as your Splunk Cloud environment. The bucket must have "gcp-to-28078-test-4-0697" as the prefix in the name.

GCP bucket folder

GCP bucket path

Path is <Bucket Name>/<Bucket Folder>

GCP service account

- Cluster Master:  
gcp-to-28078-test-4-c0m1@gcp-to-28078-test-4-0697.iam.gserviceaccount.com
- Indexers:  
gcp-to-28078-test-4-idx@gcp-to-28078-test-4-0697.iam.gserviceaccount.com

You must configure proper permissions for the two GCP service accounts. [Learn more](#)

Test bucket policy

## Configure a self storage location for the GCS bucket

To configure your GCS bucket as a self storage location:

1. In Splunk Web, select **Settings > Indexes > New Index**.
2. Under **Dynamic Data Storage**, select the **Self Storage** radio button.
3. Select **Create a self storage location**.  
The Dynamic Data Self Storage Locations page opens.
4. Select **New Self Storage Location**.  
The New Self Storage Location modal opens.
5. Give your new storage location a **Title** and a **Description** (optional).
6. In the **GCP bucket name** field, enter the name of the GCS bucket you created.
7. (Optional) In the **GCP bucket folder** field, enter the name of the GCS bucket folder.
8. Under **GCP service account**, note the 2 service account strings. In your GCP console, make sure that each service account is assigned the proper GCP roles of `Storage Legacy Bucket Writer` and `Storage Legacy Object Reader`, as discussed under "permissions" in the previous section.
9. Select **Test**.  
Splunk Cloud Platform writes a 0 KB test file to the root of your GCS bucket to verify that Splunk Cloud Platform has permissions to write to the bucket. A success message appears, and the **Submit** button is activated.
10. Select **Submit**.

## Manage self storage settings for an index

Enable Dynamic Data Self Storage on any Splunk Cloud Platform index to allow expired data to be stored to an Amazon S3 or GCP bucket.

Managing self storage settings requires the Splunk `indexes_edit` capability. All self storage configuration changes are logged in the `audit.log` file.

### Enable self storage for an index

#### Prerequisite

You must have configured a self storage location. See [Configure Self Storage Locations](#) for details.

1. Go to **Settings > Indexes**.
2. Click **New Index** to create a new index, or click **Edit** in the Actions column for an existing index.
3. In the **Dynamic Data Storage** field, click the radio button for **Self Storage**.
4. Select a storage location from the drop-down list.
5. Click **Save**.

### Disable self storage for an index

If you disable self storage for an index, expired data is deleted.

1. Go to **Settings > Indexes**.
2. Click **Edit** in the Actions column for the index you want to manage.
3. In the **Dynamic Data Storage** field, click the radio button for **No Additional Storage**.
4. Click **Save**. Self storage is disabled for this index. When data in this index expires, it is deleted.

Disabling self storage for an index does not change the configuration of the external location, nor does it delete the

external location or the data stored there. Disabling self storage also does not affect the time or size of the data retention policy for the index.

## Verify Splunk Cloud Platform successfully moved your data

To verify that your data was successfully moved to your self storage location, you can search the `splunkd.log` files.

You must have an `sc_admin` role to search the `splunkd.log` files.

Follow the appropriate procedure for your deployment's Splunk Cloud Platform Experience. For how to determine if your deployment uses the Classic Experience or Victoria Experience, see [Determine your Splunk Cloud Platform Experience](#).

### ***Classic Experience procedure***

1. Search your `splunkd.log` files to view the self storage logs. You search these files in Splunk Web by running the following search:

```
index="_internal" component=SelfStorageArchiver
```

2. Search to see which buckets were successfully moved to the self storage location:

```
index="_internal" component=SelfStorageArchiver "Successfully transferred"
```

3. Verify that all the buckets you expected to move were successfully transferred.

### ***Victoria Experience procedure***

1. Search your `splunkd.log` files to view the self storage logs. You search these files in Splunk Web by running the following search:

```
index="_internal" "transferSelfStorage Successfully transferred raw data to self storage"
```

The results of this search also show you which buckets were successfully moved to the self-storage location.

2. Verify that all the buckets you expected to move were successfully transferred.

## Monitor changes to self storage settings

You might want to monitor changes to self storage settings to ensure that the self storage locations and settings meet your company's requirements over time. When you make changes to self storage settings, Splunk Cloud logs the activity to the `audit.log`. You can search these log entries in Splunk Web by running the following search.

```
index="_audit"
```

Note that Splunk Cloud Platform cannot monitor the settings for the self storage bucket on AWS or GCP. For information about monitoring your self storage buckets, see the following:

- AWS: Access the Amazon S3 documentation and search for "Monitoring Tools".
- GCP: Access the Google Cloud Storage documentation and search for Working with Buckets and Cloud Monitoring.



The following examples apply to AWS and GCP and show the log entries available for monitoring your self storage settings.

### **Log entry for a new self storage location**

Splunk Cloud Platform logs the activity when you create a new self storage location. For example:

```
10-01-2017 11:28:26.180 -0700 INFO AuditLogger - Audit:[timestamp=10-01-2017 11:28:26.180,
user=splunk-system-user, action=self_storage_enabled, info="Self storage enabled for this index.",
index="dynamic_data_sample" ][n/a]
```

You can search these log entries in Splunk Web by running the following search.

```
index="_audit" action=self_storage_create
```

### **Log entry when you remove a self storage location**

Splunk Cloud Platform logs the activity when you remove a self storage location. For example:

```
10-01-2017 11:33:46.180 -0700 INFO AuditLogger - Audit:[timestamp=10-01-2017 11:33:46.180,
user=splunk-system-user, action=self_storage_disabled, info="Self storage disabled for this index.",
index="dynamic_data_sample" ][n/a]
```

You can search these log entries in Splunk Web by running the following search.

```
index="_audit" action=self_storage_disabled
```

### **Log entry when you change settings for a self storage location**

Splunk Cloud Platform logs the activity when you change the settings for a self storage location. For example:

```
09-25-2017 21:14:21.190 -0700 INFO AuditLogger - Audit:[timestamp=09-25-2017 21:14:21.190,
user=splunk-system-user, action=self_storage_edit, info="A setting that affects data retention was
changed.", index="dynamic_data_sample", setting="frozenTimePeriodInSecs", old_value="440", new_value="5000"
][n/a]
```

The following table shows settings that might change.

| Field                                  | Description                                                                                                                        |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| info="Archiver index setting changed." | Notification that you successfully changed self storage settings for the specified index.                                          |
| index="dynamic_data_sample"            | Name of the index for which self storage settings were modified.                                                                   |
| setting="frozenTimePeriodInSecs"       | The number of seconds before an event is removed from an index. This value is specified in days when you configure index settings. |
| old_value="440"                        | Value before the setting was updated.                                                                                              |
| new_value="5000"                       | Value after the setting has been updated.                                                                                          |

You can search these log entries in Splunk Web by running the following search.

```
index="_audit" action=self_storage_edit
```

## Restore indexed data from a self storage location

You might need to restore indexed data from a self storage location. You restore this data by moving the exported data into a thawed directory on a Splunk Enterprise instance, such as `$SPLUNK_HOME/var/lib/splunk/defaultdb/thaweddb`. When it is restored, you can then search it.

You can restore one bucket at a time. Make sure that you are restoring an entire bucket that contains the rawdata journal, and not a directory within a bucket.

An entire bucket initially contains a rawdata journal and associated tsidx and metadata files. During the DDSS archival process, only the rawdata journal is retained.

For more information on buckets, see How the indexer stores indexes in Splunk Enterprise *Managing Indexers and Clusters of Indexers*.

Data in the `thaweddb` directory is not subject to the server's index aging scheme, which prevents it from immediately expiring upon being restored. You can put archived data in the thawed directory for as long as you need it. When the data is no longer needed, simply delete it or move it out of the thawed directory.

As a best practice, restore your data using a 'nix machine. Using a Windows machine to restore indexed data to a Splunk Enterprise instance might result in a benign error message. See [Troubleshoot Dynamic Data Self Storage](#).

### Restore indexed data from an AWS S3 bucket

1. Set up a Splunk Enterprise instance. The Splunk Enterprise instance can be either local or remote. If you have an existing Splunk Enterprise instance, you can use it.

You can restore self storage data only to a Splunk Enterprise instance. You can't restore self storage data to a Splunk Cloud Platform instance.

2. Install the AWS Command Line Interface tool on your local machine. The AWS CLI tool must be installed in the same location as the Splunk Enterprise instance responsible for rebuilding.
3. Configure the AWS CLI tool with the credentials of your AWS self storage location. For instructions on configuring the AWS CLI tool, see the [Amazon Command Line Interface Documentation](#).
4. Use the recursive copy command to download data from the self storage location to the `thaweddb` directory for your index. You can restore only one bucket at a time. If you have a large number of buckets to restore, consider using a script to do so. Use syntax similar to the following:

```
aws s3 cp s3://<self storage bucket>/<self_storage_folder(s)>/<index_name>
/SPLUNK_HOME/var/lib/splunk/<index_name>/thaweddb/ --recursive
```

Make sure you copy all the contents of the archived Splunk bucket because they are needed to restore the data. For example, copy starting at the following level: `db_timestamp_timestamp_bucketID`. Do not copy the data at the level of raw data (`.gz` files). The buckets display in the `thaweddb` directory of your Splunk Enterprise instance.

5. Restore the indexes by running the following command:

```
./splunk rebuild <SPLUNK_HOME>/var/lib/splunk/<index_name>/thaweddb/<bucket_folder> <index_name>
```

When the index is successfully restored, a success message displays and additional bucket files are added to the thawed directory, including `tsidx` source types.

6. After the data is restored, go to the Search & Reporting app, and search on the restored index as you would any

other Splunk index.

When you restore data to the thawed directory on Splunk Enterprise, it does not count against the indexing license volume for the Splunk Enterprise or Splunk Cloud Platform deployment.

### **Restore indexed data from a GCP bucket**

1. Set up a Splunk Enterprise instance. The Splunk Enterprise instance can be either local or remote. If you have an existing Splunk Enterprise instance, you can use it.

You can restore self storage data only to a Splunk Enterprise instance. You can't restore self storage data to a Splunk Cloud Platform instance.

2. Install the GCP command line interface tool, `gsutil`, on your local machine. The GCP CLI tool must be installed in the same location as the Splunk Enterprise instance responsible for rebuilding.
3. Configure the `gsutil` tool with the credentials of your GCP self storage location. For instructions on configuring the `gsutil` tool, see the `gsutil` tool documentation.
4. Use the recursive copy command to download data from the self storage location to the `thaweddb` directory for your index. You can restore only one bucket at a time. If you have a large number of buckets to restore, consider using a script to do so. Use syntax similar to the following:

```
gsutil cp -r gs:<self storage bucket>/<self_storage_folder(s)>/<index_name>/  
/SPLUNK_HOME/var/lib/splunk/<index_name>/thaweddb/
```

Make sure you copy all the contents of the archived Splunk bucket because they are needed to restore the data. For example, copy starting at the following level: `db_timestamp_timestamp_bucketID`. Do not copy the data at the level of raw data (`.gz` files). The buckets display in the `thaweddb` directory of your Splunk Enterprise instance.

5. Restore the indexes by running the following command:  

```
./splunk rebuild <SPLUNK_HOME>/var/lib/splunk/<index_name>/thaweddb/<bucket_folder> <index_name>
```

When the index is successfully restored, a success message displays and additional bucket files are added to the thawed directory, including `tsidx` source types.
6. After the data is restored, go to the Search & Reporting app, and search on the restored index as you would any other Splunk index.

When you restore data to the thawed directory on Splunk Enterprise, it does not count against the indexing license volume for the Splunk Enterprise or Splunk Cloud deployment.

## **Troubleshoot DDSS with AWS S3**

This section lists possible errors when implementing DDSS with AWS S3.

### ***I don't know the region of my Splunk Cloud Platform environment***

I received the following error when testing my self storage location:

Your S3 bucket must be in the same region as your Splunk Cloud Platform environment <AWS Region>.

## Diagnosis

Splunk Cloud Platform detected that you created your S3 bucket in a different region than your Splunk Cloud Platform environment.

## Solution

If you are unsure of the region of your Splunk Cloud Platform environment, review the error message. The `<AWS Region>` portion of the error message displays the correct region to create your S3 bucket. After you determine the region, repeat the steps to create the self storage location.

### *I received an error when testing the self storage location*

When I attempted to create a new self storage location, the following error occurred when I clicked the **Test** button:

```
Unable to verify the region of your S3 bucket, unable to get bucket_region to verify. An error occurred (403) when calling the Headbucket operation: Forbidden. Contact Splunk Support.
```

## Diagnosis

You might get an error for the following reasons:

- You modified the permissions on the bucket policy.
- You pasted the bucket policy into the incorrect Amazon S3 bucket.
- You did not paste the bucket policy to the Amazon S3 bucket, or you did not save the changes.
- An error occurred during provisioning.

## Solution

1. Ensure that you did not modify the S3 bucket permissions. The following actions must be allowed: `s3:PutObject`, `s3:GetObject`, `s3:ListBucket`, `s3:ListBucketVersions`, `s3:GetBucketLocation`.
2. Verify that you applied the bucket policy to the correct S3 bucket, and that you saved your changes.
3. If you created the S3 bucket in the correct region, the permissions are correct and you applied and saved the bucket policy to the correct S3 bucket, contact Splunk Support to further troubleshoot the issue.

To review the steps to create the S3 bucket, see [Configure self storage in Amazon S3](#) in this topic.

To review how to apply a bucket policy, see the Amazon AWS S3 documentation and search for "how do I add an S3 bucket policy?".

### *I'm using Splunk Cloud Platform for a US government entity, and received an error message that the bucket couldn't be found.*

I received the following error message:

```
Cannot find the bucket '{bucket_name}', ensure that the bucket is created in the '{region_name} region.
```

## Diagnosis

For security reasons, S3 bucket names aren't global for US government entities using Splunk Cloud Platform because Splunk can only verify the region of the stack. Buckets with the same name can exist in the available AWS regions that Splunk Cloud Platform supports. For more information, see the Available regions section in the Splunk Cloud Platform Service Description. For more information about AWS GovCloud (US), see the AWS GovCloud (US) website and the AWS GovCloud (US) User Guide.

## Solution

If buckets that share the same name must exist in both regions, add the missing bucket to the appropriate region.

## Troubleshoot DDSS with GCP

### *I received a region error*

I received one of the following errors when testing my GCP self storage location:

#### **Error 1**

```
Unable to verify the region of your bucket=<bucket-id>.404 GET
https://storage.googleapis.com/storage/v1/b/<bucket-id>?projection=noAcl&prettyPrint=false: Not Found.
Contact Splunk Support.
```

#### **Error 2**

```
Unable to verify the region of your bucket=<bucket-id>.403 GET
https://storage.googleapis.com/storage/v1/b/<bucket-id>?projection=noAcl&prettyPrint=false:
<gcp-cm-serviceaccount> does not have storage.buckets.get access to the Google Cloud Storage bucket.
Contact Splunk Support.
```

#### **Error 3**

```
Your bucket in US-EAST4 is NOT in the same region as your Splunk Cloud environment: US-CENTRAL1.
```

#### **Diagnosis**

Splunk Cloud detected region or permissions issues with your GCP self storage location that must be resolved.

## Solution

#### **Error 1**

This error may indicate that the bucket may not exist, or that the bucket does not have read access. Confirm that the bucket exists and that it has the correct permissions.

#### **Error 2**

This indicates an issue with the DDSS bucket access. Splunk Cloud GCP CM and IDX service accounts must have access to the bucket. Ensure that both service accounts have `Storage Legacy Bucket Writer` role access to the DDSS bucket.

#### **Error 3**

This error indicates that the assigned region for the GCP bucket does not match the assigned region for your Splunk Cloud Platform environment.

When using DDSS with GCP, Splunk Cloud Platform does not support multi-region buckets.

To review the steps to create the GCP bucket in your GCP environment and then configure it for Splunk Cloud Platform, see [Configure self storage in GCP](#) in this topic.

## ***I received an error when testing the self storage location***

When I attempted to create a new GCP self storage location, I received one of the following errors when I clicked the Test button.

- The GCP CM service account doesn't have `create objects` access.

```
Something went wrong with bucket access. Check that the bucket exists and that the service account is granted permission. Error details: 403 POST
https://storage.googleapis.com/upload/storage/v1/b/<bucket-id>/o?uploadType=multipart: { "error": { "code": 403, "message": "<gcp-cm-serviceaccount> does not have storage.objects.create access to the Google Cloud Storage object.", "errors": [ { "message": "<gcp-cm-serviceaccount> does not have storage.objects.create access to the Google Cloud Storage object.", "domain": "global", "reason": "forbidden" } ] } } : ('Request failed with status code', 403, 'Expected one of', <HTTPStatus.OK: 200>)
```

- The GCP CM service account doesn't have `delete objects` access.

```
Something went wrong with bucket access. Check that the bucket exists and that the service account is granted permission. Error details: 403 DELETE
https://storage.googleapis.com/storage/v1/b/<bucket-id>/o/splunk_bucket_policy_test_file1619122815.2109005?generation=1619122815419331&prettyPrint=false: <gcp-cm-serviceaccount> does not have storage.objects.delete access to the Google Cloud Storage object.
```

### **Diagnosis**

You might get an error for the following reasons:

- The GCP CM and IDX service accounts must have CRUD access to the bucket. Ensure that both service accounts have the same permissions to the bucket.
- You did not assign the correct GCP role in the GCP service account field. The correct role is `Storage Legacy Bucket Writer`.
- You did not save the changes.
- An error occurred during provisioning.

### **Solution**

1. Verify that you assigned the correct GCP role to the correct GCP bucket, and that you saved your changes.
2. If you created the GCP bucket in the correct region, the permissions are correct and you applied and saved the bucket policy to the correct GCP bucket, contact Splunk Support to further troubleshoot the issue.

To review the steps to create the GCP bucket in your GCP environment and then configure it for Splunk Cloud Platform, see [Configure self storage in GCP](#) in this topic. For more information on managing GCP service accounts, see the Google Cloud documentation [Creating and managing service accounts](#).

## **Troubleshoot DDSS when restoring data using Windows**

This section lists possible errors when restoring data using a Windows machine.

### ***I received an error when using Windows to restore data.***

I attempted to restore data using a Windows machine, but the following error occurred:

```
Reason='ERROR_ACCESS_DENIED'. Will try to copy contents
```

## Diagnosis

This error occurs only on Windows builds and is benign. Splunk Cloud Platform bypasses this error by copying the content. You can safely ignore the error and continue with the restore process.

## Solution

This error is benign. You can ignore it and continue with the restore process. See [Restore indexed data from a self storage location](#).

## Store expired Splunk Cloud Platform data in a Splunk-managed archive

Dynamic Data Active Archive (DDAA) lets you move your data from your Splunk Cloud Platform indexes to a Splunk-managed archive. You can use DDAA to maintain access to older data for compliance purposes. You specify archiving at the index level by creating an archiving rule for a specified index. This gives you the flexibility to archive only the specific data that you need to maintain.

You can configure Splunk Cloud Platform to automatically archive the data from an index when the data either reaches a specified maximum size or the end of the Splunk Cloud Platform searchable retention period for an index. You can restore archived data to your Splunk Cloud Platform environment for searching within the configured archival retention time period.

You can manually clear restored data or let it auto-expire from searchable storage after 30 days. You can also track archived and restored data storage consumption, as well as the growth and expiration of your archived data.

Dynamic Data Active Archive moves data from your Splunk Index to a Splunk-maintained archive, and subsequently back from the Splunk-maintained archive to the Splunk Index in a secure and tamper-resistant manner.

## How Dynamic Data Active Archive works

Data is moved to the archive when the index meets a configured size or time threshold. When that threshold is met, Splunk Cloud Platform attempts to move the data to the archive location. If an error occurs, if there are connection issues, Splunk Cloud Platform attempts to move the data every 15 minutes until it can successfully move it.

It can take up to 48 hours from the archive initiation for the archiving process to complete.

If an error occurs, the error is logged to the `splunkd.log`. Splunk Cloud Platform does not delete data from the Splunk Cloud Platform environment until it has successfully moved the data to the archive. If you need to restore the data so that it is searchable, you can restore the data to your Splunk Cloud Platform environment. You can then search the data and delete it when you have finished.

When you restore archived data to Splunk Cloud Platform, it does not count against the indexing license volume for the Splunk Cloud Platform deployment.

## Dynamic Data Active Archive Performance

Restoring large amounts of archived data can impact performance. Splunk Cloud Platform has checks in place to help you determine if the amount of data you want to restore is too large, and it provides a warning when the data size may impact

performance. Splunk Cloud Platform will block you from restoring large amounts of data that could potentially have an extremely negative impact on performance. If this occurs, select a smaller time range.

## Configure archive settings for an index

This section shows you how to configure archive settings for a specific index.

Managing archive settings requires the `indexes_edit` capability. All archive changes appear in the `audit.log` file.

Setting incorrect or inadequate data retention values can result in a loss of data. If you have any questions about correctly setting the searchable and archive retention values for your Splunk Cloud Platform deployment, contact your Splunk account representative.

For more information on Splunk Cloud Platform data retention settings and policies and the DDAS and DDAA subscription options, see:

- [Manage data retention settings](#)
- Storage section in the *Splunk Cloud Platform Service Description*

### Configure archiving for an index

1. In Splunk Cloud, go to **Settings > Indexes**.
2. Click **New Index** to create a new index or click **Edit** in the Actions column for an existing index.
3. In the **Max raw data size** field, specify the maximum amount of raw data allowed before data is removed from the index and archived.
4. In the Dynamic Data Storage field, select **Splunk Archive**.
5. Set the **Searchable retention (days)** and **Archive Retention Period** values. Note the following:
  1. **Searchable retention (days)** holds the Dynamic Data Active Searchable (DDAS) or *searchable storage* value. This is the searchable retention period, and is considered *warm storage*.
  2. **Dynamic Data Storage > Splunk Archive > Archive Retention Period** holds the Dynamic Data Active Archive (DDAA), or *archive storage* value, and is considered *cold storage*. You can specify this value in years, months, or days. The maximum archive retention period is 3650 days (10 years). Specify a value within this range.
  3. The archive retention period is the total amount of time that Splunk retains your data. The archive retention period includes the searchable retention period. For example, if you want Splunk Cloud Platform to retain your data for a total of 365 days, but you want that data searchable for the first 90 days, set the searchable retention period to 90 days and the archive retention period to 365 days (not 365-90 days).
  4. When specifying the archive retention period value, you must specify a value that is greater than the searchable retention period. For example, if you set **Searchable retention (days)** to 90 days, you must set the **Archive Retention Period** to a value greater than 90 days, such as 180 days.
6. Click **Save**.

You cannot enable both DDAA and DDSS at the same time for the same index. If you enable DDAA for an index, then later decide to change the index settings to use either DDSS or no storage, you must contact Splunk Support if you want to retain the archived data.

### Disable archiving for an index

1. Go to **Settings > Indexes**.
2. Click **Edit** in the Actions column for the index you want to manage.



3. In the Dynamic Data Storage field, select **Self Storage** to move data to self-storage location when it expires or **No Additional Storage** to delete data as it expires.
4. Click **Save**. When data in this index expires, it is deleted.

Disabling archiving for an index marks the existing archived data with a status of delete. Deleted archive data will be permanently erased 30 days after the deletion date. Be aware that disabling archiving for an index does not affect the time or size of the data retention policy for the index. If you disable archiving for an index in error, contact Splunk Support as soon as possible. If you have a support contract, file a new case using the Splunk Support Portal. Otherwise, contact Splunk Customer Support.

## Restore archived data to Splunk Cloud Platform

DDAA lets you restore indexed data from the Splunk archive. Data in DDAA can be restored to Dynamic Data Active Searchable (DDAS) to be searched. After restoring data, you can search it like any other data.

You restore data based on the time period for the data you want to search. For example, you might want to restore data for a period of one day. When you pick a date from the date-picker, DDAA treats it as 12 AM UTC of the selected date. So, if you want to restore one day's worth of archived data, (for example, on 07/10/2018) you must specify 07/10/2018 in the 'from' field and 07/11/2018 in the 'to' field.

By default, restored data is searchable for a period of one month. Splunk automatically removes the data after this period. Splunk does not remove data from the archive.

The archival process can take up to 48 hours to complete and the restoration process can take up to 24 hours to complete. Because the complete archival and restoration cycle can take up to 72 hours to complete, be sure to plan any data restoration processes accordingly.

### *How restoring data works*

When you restore data to Splunk Cloud Platform from the archive, a copy of the archived data is moved back to the Splunk Cloud Platform environment. To ensure your data is safe, Splunk Cloud Platform never moves or deletes the original archived data. This method of temporary data restoration ensures that you can never mistakenly delete your archived data.

When you restore archived data to an index in your Splunk Cloud Platform instance, it does not count against the retention periods configured for data in your index. Restored data exists outside of the constraints of retention periods and size limits and does not affect the retention of your existing index data.

When you restore data, Splunk Cloud Platform checks several conditions to ensure that you do not experience performance issues and that you do not duplicate data and cause your queries to return incorrect results:

- **Check for overlapping data.** Splunk Cloud Platform does not restore data if you have already restored data in that same time range. This is to ensure you do not restore duplicate data, which would cause inaccurate search results. For example, if you specify that you want to restore data from 07/01/2018-07/03/2018, but you have already restored data from 07/01/2018-07/02/2018, Splunk Cloud Platform will prevent your data restore. In this case, it is recommended you restore the data that falls outside of the range of the data you have already restored. In this example, you would restore data from 07/03/2018-07/04/2018.
- **Check to ensure data is not likely to cause performance issues.** Splunk Cloud Platform checks the size of the data you want to restore and presents you with a warning if the size of the data may cause performance issues. If the size of that data is very likely to cause performance issues, Splunk Cloud Platform will prevent you from

restoring the data.

During the data restoration process, the Splunk platform retrieves all buckets that contain events necessary for the specified search period. For certain restoration scenarios, this can result in the total size of the restored data being much greater than the total number of restored events. This behavior is normal and to be expected.

After you have restored data, you may notice that events appear in your index that are older than your configured retention period specifies. This restored data will remain in your index for 30 days or until you clear it.

If your attempt to restore archived data fails, verify that the data was not recently archived. Because there is a time period during which data is being transitioned from Splunk Cloud Platform to the archive, you will not be able to restore that data during the processing period. Generally, data moved to the archive is available in approximately 48 hours.

If you want to restore data archived within the last 48 hours, you must explicitly disable the default "Exclude" option for Recently Archived Data in the Restore Archive modal. When set to "Exclude", DDAA skips restoration of data archived within the last 48 hours. See [Steps to restore archived data to Splunk Cloud Platform](#). In addition, ensure that the data is fully archived and the timestamps are correct, or data restoration will fail with the following error message: "You cannot restore data that was archived less than 48 hours ago. Please try again later". If you receive this message, you must set the Recently Archived Data mode to "Exclude" to proceed with data restoration. For more information, see [Troubleshoot Dynamic Data Active Archive](#).

#### **What happens when you are finished searching the restored data**

After the data is temporarily restored to your Splunk Cloud Platform environment it is available for searching for 30 days. Restored data is a copy of the archived data so you never need to move the data back to the archive, but for best performance, you should remove the temporarily restored data when you have finished searching it.

Temporarily restored data is available only for 30 days. This 30-day time period can't be modified in any way, meaning reduced or extended. Also, this time period restriction applies to all temporarily restored data, regardless of the configuration settings for your deployment's indexers.

#### ***Steps to restore archived data to Splunk Cloud Platform***

1. In Splunk Cloud, go to **Settings > Indexes**.
2. For the index where you want to restore data, click **Restore**. The menu displays the restore history for the specified index. You can see the history of data restoration and file size for the data restored.
3. Use the date picker to select a time range to retrieve.
4. Click **Check size**. Splunk Cloud Platform checks to see if the size of the file might impact performance. If the file size is too large, Splunk Cloud Platform blocks you from restoring data. If there is a potential performance impact, Splunk Cloud Platform displays a warning. Splunk Cloud Platform also prevents you from restoring data that overlaps with existing restored data.
5. Enter an email address to send job status notifications. Splunk Cloud Platform will notify you when the restoration is complete.
6. (Optional) If your time range includes data archived within the last 48 hours, toggle the **Recently Archived Data** switch to disable the default **Exclude** mode. When set to "Exclude" mode, DDAA skips restoration of data archived within the last 48 hours. Note that attempting to restore data that is not fully archived can cause data restoration to fail. For more information, see [Troubleshoot Dynamic Data Active Archive](#).
7. Click **Restore** when you have refined the file size or date range to acceptable limits.

After you initiate data restoration, it can take up to 24 hours before data is restored. If it takes longer than 24 hours, contact Splunk Technical Support.

8. To check the status of your data restoration, click **Splunk Archive** in the **Storage Type** field to open the Archive page. To view the restore status, click the **Restore** tab. In the **JobStatus** field, you can see the status of your job:

- **Pending:** The job has been submitted, but has not begun processing.
- **In progress:** The job has been started, and is progressing.
- **Success**
- **Cleared:** You've successfully deleted the temporary archive from your index.
- **Expired:** The restored data has passed the 30 day retention period and has been deleted from the index.
- **Failed:** If you receive a Failed status, click the > button for the archive to display more details about why the restoration failed.

### **Steps to remove restored data from Splunk Cloud Platform**

Splunk recommends you manually remove restored data when you are finished searching it.

Restored data is a copy of the archived data, so you never need to move the data back to the archive, but for best performance, you should remove the temporarily restored data when you are done searching it.

To remove restored data:

1. In Splunk Cloud, go to **Settings > Indexes**.
2. Select the index with data you want to remove and click **Restore** to open the Restore Archive page.
3. For the range of data you want to remove, select **Clear** in the Actions column.

When the data is successfully removed, the **Jobstatus** column displays a **Cleared** status.

### **Monitor logs during archiving**

Splunk generates logs when you archive data and when you restore archived data. You may want to monitor these logs to check for errors during these processes.

#### **Archiving logs**

To check for error messages that occur when you are archiving data, you can view the coldstoragearchiver entries in the splunkd.log. You can find these entries by running the following search:

```
index=_internal source=*/splunkd.log component=coldstoragearchiver
```

#### **Data restoration logs**

To check for error messages that occur when you restore archived data, you can view entries in the splunk\_archiver\_restoration.log, restoration.log, and python.log. You can find these entries by running the following search:

```
index=_internal source=*/splunk_archiver_restoration.log
```

```
index=_internal source=*/restoration.log
```

```
index=_internal source=*/python.log
```

## Manage your archives

You might want to review the status of your archived indexes or understand how much of your entitlement has been used. You can review the status of your archived indexes on the Archived Indexes page.

### **Steps to review the overall status of your restore requests for the last 90 days**

1. From Splunk Web, go to **Settings > Indexes**.
2. From the Indexes page, click on a value in the **Archive Retention** column.
3. Click the **Restore** tab to open the Restore page.
4. Review the **Restore Summary (90 days)** table to see the overall status of your restored data.

| Field                    | Description                                                                                                                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Total Restored Data (GB) | The total amount of raw data (uncompressed) that has been restored. This value is updated nightly.                          |
| Total Cleared Data (GB)  | The total amount of raw data (uncompressed) that has been deleted from the restored archive. This value is updated nightly. |
| Total Expired Data (GB)  | The total amount of raw data (uncompressed) that has expired from the restored archive. This value is updated nightly.      |

You can view the details for restored archived data from the last 90 days in the table below. For each index, you can see the following details:

| Field              | Description                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index Name         | The name of the restored index.                                                                                                                                 |
| Restored Count     | The total number of restoration requests, including both successful and failed restore requests. This value also includes cleared and expired restore requests. |
| Restored Size (GB) | The total amount of raw data (uncompressed) that has been restored.                                                                                             |
| Cleared Count      | The total number of restored index requests that have been manually deleted.                                                                                    |
| Cleared Size (GB)  | The total amount of raw data (uncompressed) that has been manually deleted.                                                                                     |
| Expired Count      | The total number of restored index requests that have aged out.                                                                                                 |
| Expired Size       | The total amount of restored raw data (uncompressed) that has aged out.                                                                                         |

### **Steps to review the status of individual restore requests**

1. From Splunk Web, go to **Settings > Indexes**.
2. From the Indexes page, click on a value in the Archive Retention column.
3. Click the **Restore** tab to open the Restore page.
4. Go to the **Restore Request History (Last 50 requests)** table.

From here, you can see the start time, end time, time of the request, data volume in GB, and the expiration date. To understand the status for each job, check the **Job Status** field for each index. The following table shows the possible values.

| Field   | Description                                                            |
|---------|------------------------------------------------------------------------|
| Pending | The request for restoration has been initiated, but has not yet begun. |

| Field       | Description                                                                                               |
|-------------|-----------------------------------------------------------------------------------------------------------|
| In progress | The restoration process has started, but it has not been completed.                                       |
| Success     | The data has been successfully restored to your index.                                                    |
| Failure     | The restoration failed. Click the > button next to the archive to display more details about the failure. |
| Cleared     | You have successfully cleared the temporarily restored data.                                              |
| Expired     | The restored data has passed the 30 day retention threshold.                                              |

After you have reviewed the archived indexes, you can determine what actions you want to take for each archived or restored index. You may want to clear archived data or stop archiving an index. Or you may see that a restoration or archive operation failed and chose to troubleshoot the issue.

### **Steps to review the overall size and growth of your archived indexes**

You might want to review the size and growth of your archived indexes to better understand how much of your entitlement you are consuming. This can help you predict usage and expenses for your archived data.

1. From Splunk Web, go to **Settings > Indexes**.
2. From the Indexes page, click on a value in the **Archive Retention** column.

The Archive Summary page displays the following information:

| Field                                   | Description                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Archive Usage                     | The total amount of raw data (uncompressed) that is stored in the archive. This number turns red when total archive usage exceeds the total entitlement. This value is updated nightly.                                                                                                                                                                               |
| Total Entitlement                       | Your total entitlement as determined in your service agreement.                                                                                                                                                                                                                                                                                                       |
| Total Archive Data Growth (90 Days)     | The total amount of raw data (uncompressed) that has been added to the archive in the past 90 days. This value is updated nightly.                                                                                                                                                                                                                                    |
| Total Archive Data Expiration (90 Days) | The total amount of raw data (uncompressed) that has aged out of the archive within the past 90-day window. This value is updated nightly. Note that each index has an archive retention setting and the data ages out over time. For example, index A has 2-year archive retention. Every night for that index, Splunk ages out the data that is older than 2 years. |

### **Steps to review the size and growth of each archived index**

You might want to review the size and growth of each index to understand how much it grows over time.

1. From Splunk Web, go to **Settings > Indexes**.
2. From the Indexes page, click on a value in the **Archive Retention** column.

The Archive Summary page displays the following information:

| Field             | Description                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------|
| Index Name        | Name of the index.                                                                          |
| Current Size (GB) | The current amount of raw data (uncompressed) that is stored in the archive for each index. |
| Earliest Event    | The earliest event in the archived index.                                                   |
| Latest Event      | The latest event in the archived index.                                                     |

| Field                       | Description                                                                                                  |
|-----------------------------|--------------------------------------------------------------------------------------------------------------|
| 90-Day Data Growth (GB)     | The amount of raw data (uncompressed) that has been added to the archive in the past 90 days for each index. |
| 90-Day Data Expiration (GB) | The amount of raw data (uncompressed) that has been removed from the archive after 90 days for each index.   |

## Troubleshoot Dynamic Data Active Archive

### *I received an error when attempting to restore data*

If an error occurs, the error is logged to the `splunkd.log`. When you review the Archive page, if you experience errors, you may want to review the `splunkd.log` and specify the `coldstoragearchiver` component here: `index=_internal source=*/splunkd.log component=coldstoragearchiver`

### *I clicked the Check Size button and nothing happened*

When restoring data, I clicked the **Check Size** button multiple times and nothing happened.

#### Diagnosis

When restoring a large amount of data, it may take some time for Splunk to verify that the size of the data can be restored without causing performance issues. If you click the **Check Size** buttons multiple times, it may trigger AWS to block the check process.

#### Solution

Do not click the **Check Size** button multiple times if you don't immediately receive feedback.

### *I archived some of my data. When I attempted to restore it a few hours later, an error message appeared.*

When I archived data and attempted to restore it soon after, I received an error message.

#### Diagnosis

Data can take up to 48 hours to archive. If you attempt to restore the data before this time period completes, the restoration will fail and the following error message appears: *You cannot restore data that was archived less than 48 hours ago. Please try again later.*

#### Solution

Make sure the **Recently Archived Data** switch in the Restore Data modal is set to **Exclude**. When set to Exclude, DDAA skips restoration of all data archived within the last 48 hours. See [Steps to restore archived data to Splunk Cloud Platform](#). OR

Wait until the 48 hour threshold has been met, and then attempt to restore the data. You can check the status of the archival process as follows:

- Run the following search against the internal log to determine when Splunk Cloud Platform started the archival process:

```
index=_internal component=ColdStorageArchiver "Successfully executed archiving script"
```

### ***I'm trying to restore fully archived data, but I'm still receiving an error message about data archival.***

I'm trying to restore data and the archival process completed more than 48 hours ago, but I'm receiving the following error message: *You cannot restore data that was archived less than 48 hours ago. Please try again later.*

#### **Diagnosis**

Receiving this error message after the archival process is complete indicates that there are incorrect timestamps in the data.

#### **Solution**

Contact Splunk Technical Support for help with correcting the timestamp data.

## **Manage indexes on Splunk Cloud Platform Classic Experience**

Splunk Cloud Platform Classic Experience now provides full support for managing indexes programmatically using the ACS (Admin Config Service) API. For more information, see [Manage indexes in Splunk Cloud Platform in the Admin Config Service Manual](#).

If your Splunk Cloud Platform deployment is on Classic Experience, you can manage your indexes programmatically using the Splunk REST API `cluster_blaster_indexes/sh_indexes_manager` endpoint. To determine if your Splunk Cloud Platform deployment is on Classic Experience:

1. In Splunk Web, click **Support & Services > About**.
2. In the About panel, under Splunk Cloud, find your Experience: Classic or Victoria.

For more information on Splunk Cloud Platform Experiences, see [Determine your Splunk Cloud Platform Experience](#).

### **Requirements**

- Splunk Cloud Platform version 8.0.2007 or higher.
- You must have the `sc_admin` (Splunk Cloud Administrator) role.

#### ***Authentication and authorization***

There are two methods you can use to authenticate and authorize endpoint requests:

- Authentication token. You can create a valid JWT authentication token in the Splunk Cloud Platform UI.
- Username and password.

For more information on authentication and authorization, see [Basic Concepts about the Splunk platform REST API in the Splunk Enterprise REST API User Manual](#).

### **Manage indexes**

You can perform the following index management actions on Splunk Cloud Platform deployments running on Classic Experience.

## List all indexes

To list all indexes, send an HTTP GET request to the following endpoint:

```
localhost:8089/services/cluster_blaster_indexes/sh_indexes_manager
```

For example:

```
curl -k -H "Authorization: Bearer ${TOKEN}"  
https://agile-albatross-4ej.stg.splunkcloud.com:8089/services/cluster_blaster_indexes/sh_indexes_manager?output  
_mode=json
```

The request output lists existing indexes, including configuration parameter values for each index. For example:

```
{ "links": { "create": "/services/cluster_blaster_indexes/sh_indexes_manager/_new", "disabled": "/services/cluster  
_blaster_indexes/sh_indexes_manager/disabled", "origin": "https://agile-albatross-4ej.stg.splunkcloud.com:8089  
/services/cluster_blaster_indexes/sh_indexes_manager", "updated": "2021-10  
-06T22:44:52+00:00", "generator": { "build": "2e4da17c2b37", "version": "8.2.2107" }, "entry": [ { "name": "christian", "id  
": "https://agile-albatross-4ej.stg.splunkcloud.com:8089/servicesNS/nobody/cloud_administration/cluster_blaster  
_indexes/sh_indexes_manager/christian", "updated": "1970-01-01T00:00:00+00:00",  
  
...  
  
"content": { "archiver.enableDataArchive": "0", "archiver.maxDataArchiveRetentionPeriod": "315360000", "datatype": "e  
vent", "disabled": "0", "eai:acl": null, "eai:acl.app": "cloud_administration", "eai:acl.appDisplayName": "_cluster  
_admin", "frozenTimePeriodInSecs": "30000", "isS2Mode": "true", "isVirtual": "0", "maxGlobalDataSizeMB": "0", "maxGloba  
lRawDataSizeMB": "5000", "maxTotalDataSizeMB": "0", "metric.timestampResolution": "s", "totalEventCount": "0", "totalR  
awSizeMB": "0" }, { "name": "cryt", "id": "https://agile-albatross-4ej.stg.splunkcloud.com:8089/servicesNS/nobody/cloud  
_administration/cluster_blaster_indexes/sh_indexes_manager/cryt", "updated": "1970-01-01T00:00:00+00:00",  
  
...  
  
"messages": [] }
```

## Create a new index

To create a new index, send an HTTP POST request to the `cluster_blaster_indexes/sh_indexes_manager` endpoint, specifying the following index parameters: `name`, `maxTotalDataSizeMB`, `frozenTimePeriodInSecs`, and `maxGlobalRawDataSizeMB`. For example:

```
curl -k -H "Authorization: Bearer ${TOKEN}"  
https://agile-albatross-4ej.stg.splunkcloud.com:8089/services/cluster_blaster_indexes/sh_indexes_manager -d  
name=my_name -d maxTotalDataSizeMB=500 -d frozenTimePeriodInSecs=30000 -d maxGlobalRawDataSizeMB=6000
```

Sample request output:

```
...  
<title></title>  
<id>https://agile-albatross-4ej.stg.splunkcloud.com:8089/services/cluster_blaster_indexes/sh_indexes  
_manager</id>  
<updated>2021-10-06T22:38:13+00:00</updated>  
<generator build="2e4da17c2b37" version="8.2.2107"/>  
<author>  
  <name>Splunk</name>  
</author>  
<link href="/services/cluster_blaster_indexes/sh_indexes_manager/_new" rel="create"/>  
<link href="/services/cluster_blaster_indexes/sh_indexes_manager/disabled" rel="disabled"/>
```



```
<opensearch:totalResults>0</opensearch:totalResults>
<opensearch:itemsPerPage>30</opensearch:itemsPerPage>
<opensearch:startIndex>0</opensearch:startIndex>
<s:messages/>
```

### ***View an individual index***

To view an individual index, send an HTTP GET request to the `cluster_blaster_indexes/sh_indexes_manager/{name}` endpoint, specifying the name of the index. For example:

```
curl -k -H "Authorization: Bearer ${TOKEN}"
https://agile-albatross-4ej.stg.splunkcloud.com:8089/services/cluster_blaster_indexes/sh_indexes_manager/christian
```

Sample request output:

```
...
  <title>christian</title>
  <id>https://agile-albatross-4ej.stg.splunkcloud.com:8089/servicesNS/nobody/cloud_administration/cluster
_blaster_indexes/sh_indexes_manager/christian</id>
  <updated>1970-01-01T00:00:00+00:00</updated>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="alternate"/>
  <author>
    <name>system</name>
  </author>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="list"/>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="edit"/>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="remove"/>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian/clean"
rel="clean"/>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian/disable"
rel="disable"/>
  <content type="text/xml">
    <s:dict>
      <s:key name="archiver.enableDataArchive">0</s:key>
      <s:key name="archiver.maxDataArchiveRetentionPeriod">315360000</s:key>
      <s:key name="datatype">event</s:key>
      <s:key name="disabled">0</s:key>
      <s:key name="eai:acl">
        <s:dict>
          <s:key name="app">cloud_administration</s:key>
          <s:key name="can_list">1</s:key>
          <s:key name="can_write">1</s:key>
          <s:key name="modifiable">0</s:key>
          <s:key name="owner">system</s:key>
          <s:key name="perms"/>
          <s:key name="removable">0</s:key>
          <s:key name="sharing">system</s:key>
        </s:dict>
      </s:key>
      <s:key name="eai:acl.app">cloud_administration</s:key>
      <s:key name="eai:acl.appDisplayName">_cluster_admin</s:key>
```

```

<s:key name="eai:attributes">
  <s:dict>
    <s:key name="optionalFields">
      <s:list>
        <s:item>datatype</s:item>
        <s:item>maxGlobalDataSizeMB</s:item>
        <s:item>maxTotalDataSizeMB</s:item>
        <s:item>metric.timestampResolution</s:item>
      </s:list>
    </s:key>
    <s:key name="requiredFields">
      <s:list>
        <s:item>frozenTimePeriodInSecs</s:item>
        <s:item>maxGlobalRawDataSizeMB</s:item>
      </s:list>
    </s:key>
    <s:key name="wildcardFields">
      <s:list>
        <s:item>archiver\..*</s:item>
      </s:list>
    </s:key>
  </s:dict>
</s:key>
<s:key name="frozenTimePeriodInSecs">30000</s:key>
<s:key name="isS2Mode">>true</s:key>
<s:key name="isVirtual">0</s:key>
<s:key name="maxGlobalDataSizeMB">0</s:key>
<s:key name="maxGlobalRawDataSizeMB">5000</s:key>
<s:key name="maxTotalDataSizeMB">0</s:key>
<s:key name="metric.timestampResolution">s</s:key>
<s:key name="totalEventCount">0</s:key>
<s:key name="totalRawSizeMB">0</s:key>
</s:dict>
</content>
</entry>

```

### **Update an index**

To update an index, send an HTTP POST request to the `cluster_blaster_indexes/sh_indexes_manager/{name}` endpoint, specifying the name of the index you want to update, along with the updated index parameter values. The POST request support updates to `maxTotalDataSizeMB`, `frozenTimePeriodInSecs`, and `maxGlobalRawDataSizeMB` parameters. For example:

```

curl -k -H "Authorization: Bearer ${TOKEN}" -X 'POST'
https://agile-albatross-4ej.stg.splunkcloud.com:8089/services/cluster_blaster_indexes/sh_indexes_manager/christian
-d maxTotalDataSizeMB=500 -d frozenTimePeriodInSecs=30000 -d maxGlobalRawDataSizeMB=6000

```

### **Sample request output:**

```

...
<entry>
  <title>christian</title>
  <id>https://agile-albatross-4ej.stg.splunkcloud.com:8089/servicesNS/nobody/cloud_administration/cluster
_blaster_indexes/sh_indexes_manager/christian</id>
  <updated>1970-01-01T00:00:00+00:00</updated>
  <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="alternate"/>
  <author>
    <name>system</name>

```

```

    </author>
    <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="list"/>
    <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="edit"/>
    <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian"
rel="remove"/>
    <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian/clean"
rel="clean"/>
    <link
href="/servicesNS/nobody/cloud_administration/cluster_blaster_indexes/sh_indexes_manager/christian/disable"
rel="disable"/>
    <content type="text/xml">
      <s:dict>
        <s:key name="archiver.enableDataArchive">0</s:key>
        <s:key name="archiver.maxDataArchiveRetentionPeriod">315360000</s:key>
        <s:key name="datatype">event</s:key>
        <s:key name="disabled">0</s:key>
        <s:key name="eai:acl">
          <s:dict>
            <s:key name="app">cloud_administration</s:key>
            <s:key name="can_list">1</s:key>
            <s:key name="can_write">1</s:key>
            <s:key name="modifiable">0</s:key>
            <s:key name="owner">system</s:key>
            <s:key name="perms"/>
            <s:key name="removable">0</s:key>
            <s:key name="sharing">system</s:key>
          </s:dict>
        </s:key>
        <s:key name="eai:acl.app">cloud_administration</s:key>
        <s:key name="eai:acl.appDisplayName">_cluster_admin</s:key>
        <s:key name="frozenTimePeriodInSecs">70000</s:key>
        <s:key name="isS2Mode">>true</s:key>
        <s:key name="isVirtual">0</s:key>
        <s:key name="maxGlobalDataSizeMB">0</s:key>
        <s:key name="maxGlobalRawDataSizeMB">6000</s:key>
        <s:key name="maxTotalDataSizeMB">0</s:key>
        <s:key name="metric.timestampResolution">s</s:key>
        <s:key name="totalEventCount">0</s:key>
        <s:key name="totalRawSizeMB">0</s:key>
      </s:dict>
    </content>
  </entry>

```

### **Delete an index**

To delete an index, send an HTTP DELETE request to the `cluster_blaster_indexes/sh_indexes_manager/{name}` endpoint, specifying the name of the index you want to delete. For example:

```

curl -k -H "Authorization: Bearer ${TOKEN}" -X 'DELETE'
https://agile-albatross-4ej.stg.splunkcloud.com:8089/services/cluster_blaster_indexes/sh_indexes_manager/christian

```

# Manage Apps and Add-ons in Splunk Cloud Platform

## Install apps on your Splunk Cloud Platform deployment

Splunk apps are composed of pre-built dashboards, reports, alerts, and workflows, optimized for a particular purpose such as monitoring Web servers or network security. Splunk add-ons are a type of app that provide specific capabilities to other apps, such as getting data in, mapping data, or providing saved searches and macros. For more information on Splunk apps, see *Apps and add-ons in the Splunk Enterprise Admin Manual*.

You can install most Splunk apps on Splunk Cloud Platform in a self-service manner without assistance from Splunk support. For more information, see [About self-service app installation](#).

To install premium apps, such as Enterprise Security (ES) or IT Service Intelligence (ITSI), you must contact Splunk Support. For information on assisted installation for premium apps, see Splunk premium solutions in the *Splunk Cloud Platform Service Description*.

You can only install approved apps that meet Splunk Cloud Platform requirements. For information on Splunk app validation, see [Validate the quality of Splunk apps using Splunk AppInspect](#).

### About self-service app installation

Splunk Cloud Platform supports self-service installation of both public apps available from Splunkbase and private apps that you create for your deployment. More than 98% of Splunkbase apps are available for self-service app installation. You can install both Splunkbase apps and private apps directly to your Splunk Cloud Platform deployment using the Splunk Web UI.

For instructions on how to install Splunkbase apps using Splunk Web, see [Install a public app from Splunkbase](#).

For instructions on how to install private apps using Splunk Web, see [Manage private apps in your Splunk Cloud deployment](#).

You can also install both Splunkbase apps and private apps programmatically using the Admin Config Service (ACS) API. For more information, see [About the Admin Config Service \(ACS\) API](#) in the *Admin Config Service Manual*.

Self-service app installation behavior differs depending if your deployment is on Victoria Experience or Classic Experience. For more information, see [How self-service app installation works on Victoria Experience](#).

### When to contact Splunk Support

Some app installation and upgrade tasks require help from Splunk Support. Contact Splunk Support and submit a Cloud App Request in the following situations:

- The app is not available for self-service installation.
- The app is not available on Splunkbase.
- The app requires installation on the IDM (Inputs Data Manager).
- The app is a premium solution, such as Enterprise Security (ES) and IT Service Intelligence (ITSI).
- The app requires installation on a Classic Experience premium search head that runs a Splunk Premium App. To determine if your Splunk Cloud Platform deployment is on Victoria Experience or Classic Experience, see [Determine your Splunk Cloud Platform Experience](#).

For more information, see Splunkbase and private apps in the *Splunk Cloud Platform Service Description*.

## Install a public app from Splunkbase

You can install most public apps from Splunkbase directly to your deployment using the Splunk Apps Browser in Splunk Web. You must have the `sc_admin` role to install apps on Splunk Cloud Platform.

To install a public app on Splunk Cloud Platform:

1. In Splunk Web, click the **Apps** gear icon.
2. Click **Browse more apps**.  
The Splunk App Browser opens.
3. Find your app or add-on, then click **Install**.
4. Enter your Splunk.com login credentials (username and password).
5. Click **Agree and Install**.  
This confirms that you accept the app license terms and installs the app on your deployment.
6. Consult the specific app's documentation to determine if you must also install the app on forwarders. If yes, you can download the app package from Splunkbase and deploy it manually to your forwarders.

When you install an app with declared dependencies, Splunk Cloud Platform automatically resolves its dependencies through Splunkbase. To learn more about dependencies, see Splunk Packaging Toolkit.

### ***Install restricted Splunkbase apps***

App developers can control access to certain apps posted on Splunkbase by setting them to either unrestricted or restricted. If an app is set to restricted, only authorized users can download and install the app. To gain access to a restricted app, you must contact the developer specified on the app download page on Splunkbase.

When you install a restricted Splunkbase app using self-service app installation, you must specify the authorized user's Splunk.com credentials (username and password). You can find the username on your account at <https://splunkbase.splunk.com/profile>. Make sure to check the specific app installation instructions on Splunkbase for any additional download or installation requirements.

To install a restricted app that is not supported by self-service app installation, you must contact Splunk Support. For apps not authored by Splunk, you must download the app, specifying the authorized user's credentials, then upload it to a support case with your installation request. For apps authored by Splunk, you do not need to upload the app with your support case. To identify the app author, see the specific app download page on Splunkbase.

For more information on restricted Splunkbase apps, see Manage Splunk Cloud Platform and Splunk Enterprise content on Splunkbase in the *Splunk Developer Guide*.

## Manage apps

You can perform the following app management tasks on the App Management page of a Splunk Cloud Platform deployment.

### ***Update an app***

You can update an app using either the App Management page or the App Browser page in Splunk Web.

To update an app using the App Management page:

1. In Splunk Web, click **Apps > Manage Apps**.
2. Find your app, then click **Update Available** to install the new version.

To update an app using the App Browser page:

1. In Splunk Web, click **Apps > Find More Apps**.
2. Find your app, then click **Update**.

After you update an app, you cannot revert to an earlier version of the app. If a new version of your app is available, but the update action is not available in Splunk Web, contact Splunk Support.

You cannot update Splunk Premium apps using the App Management page in Classic Experience. To update a premium app in Classic Experience, contact Splunk Support. You cannot update Enterprise Security (ES) or ITSI premium apps using the App Management page in either Classic or Victoria Experience. To do so, you must contact Splunk support.

### ***Uninstall an app***

1. In Splunk Web, click **Apps > Manage apps**.
2. Click **Uninstall**. If the Uninstall action is not available for your app, open a support case.

### ***View app install details***

To view app installation details, including install date, install location, and install type (self-service or Splunk):

1. In Splunk Web, click **Apps > Manage apps**.
2. Click the arrow next to the app name to see app installation details.

### ***Configure an app***

To configure an app, consult the specific app's documentation for instructions. Configure apps only on the nodes in your deployment where configuration is required.

On Splunk Cloud Platform deployments, inputs must be configured on forwarders under your control.

## **Manage private apps on your Splunk Cloud Platform deployment**

Private apps are custom apps that you create for your Splunk Cloud Platform deployment. Private apps are not publicly available on Splunkbase.

You can install private apps on your Splunk Cloud Platform deployment in a self-service manner using the app management page in Splunk Web.

When you upload an app, Splunk Cloud Platform automatically runs the app through Splunk AppInspect validation to confirm that it meets Splunk Cloud Platform requirements. For more information, see [Install private apps on Splunk Cloud Platform](#).

You can also install private apps programmatically using the Admin Config Service (ACS) API. For more information, see *Manage private apps in Splunk Cloud Platform* in the *Admin Config Service Manual*.

You must have the `sc_admin` (Splunk Cloud Platform Administrator) role to install apps in Splunk Cloud Platform.

## Create a private app

This section provides an overview of the steps involved in creating a private app for your Splunk Cloud Platform deployment. For detailed information on how to create a private app, see *Develop Splunk Apps for Splunk Cloud Platform* in the *Splunk Developer Guide*.

### Prerequisites

- See the *Building Splunk Apps* documentation on Splunkbase.
- For information about dependencies, see the *Splunk Packaging Toolkit* in the *Splunk Developer Guide*.
- For information on how to validate your app using Splunk AppInspect, see *Validate the quality of Splunk apps using AppInspect* in the *Splunk Developer Guide*.
- If your private app uses Python, make sure to use a supported Python version. For more details, see the *Splunk Cloud* section of the *Python 3 Migration* guide.

### Steps

1. Create an app that conforms to Splunk app standards and requirements.
2. Make sure the app package does not have any static dependencies. Splunk Cloud Platform supports dynamic dependencies only.
3. Package the app using `.spl`, `.tar`, `.tar.gz` or `.tgz` file format. Limit the package size to 128MB.
4. Run the app through AppInspect and make sure it passes all app validation checks. The AppInspect validation report must show 0 Failures, 0 Errors, and 0 Manual Checks before you can install the app using self-service app installation in Splunk Web.
5. You can now install the app on your deployment using Splunk Web. See [Install private apps on Splunk Cloud Platform](#).

If the AppInspect report shows greater than 0 Manual Checks, you must contact Splunk Support to install your app.

## Install private apps on Splunk Cloud Platform

Splunk Cloud Platform supports self-service installation of private apps on search heads and indexers.

In Splunk Cloud Platform version 8.2.2106 and higher, there are two slightly different private app installation workflows that can appear in Splunk Web. Both workflows automatically run your app through AppInspect validation checks and let you view a report that shows the results of the app validation process.

The app installation workflow available to you in Splunk Web depends on your Splunk Cloud Platform Experience: Victoria or Classic. To find your Splunk Cloud Platform Experience, in Splunk Web, click **Support & Services > About**.

After you determine your Splunk Cloud Platform Experience, follow the app installation instructions that apply to your deployment:

- [Install a private app on Victoria Experience](#).

- [Install a private app on Classic Experience.](#)

For more information on Splunk Cloud Platform Experience, see [Determine your Splunk Cloud Platform Experience.](#)

## Install a private app on Victoria Experience

If your Splunk Cloud Platform deployment is on Victoria Experience, you can upload and install your private app using the **Install app from file** workflow on the Apps page in Splunk Web.

Splunk Cloud Platform deployments on Victoria Experience do not require IDM. If your deployment is on Victoria Experience you can run apps and add-ons that contain scripted or modular inputs directly on the search head.

To install a private app on a Victoria Experience deployment:

1. In Splunk Web, click the **Apps** gear icon.
2. Click **Install app from file.**
3. Click **Upload App.**
4. Enter your splunk.com account credentials (username and password). Splunk Cloud Platform uses these credentials to authenticate your AppInspect app validation.
5. Click **Agree and Login.**  
This confirms that you accept the specified license conditions and submits your login credentials.
6. Select your private app package and click **Upload.**  
Your app appears in the Uploaded Apps table. Splunk Cloud Platform automatically runs your app through AppInspect validation.
7. Check the app validation status. Your app must pass all AppInspect checks and be approved before you can install it. While you only need to enter your credentials once per session, if you log out while an app is still in the vetting process, then log back in, you must click **Check Status** and enter your credentials again to complete the vetting process and view the app validation report. For more information, see [Check app validation status.](#)

You can click "Upload App" to upload additional apps while an app is in the "vetting" process.

8. If your app validation status shows approved, click **Install.** If your app validation status shows rejected, click **View Report** to determine the issues you must fix before you can install the app. For more information, see [View app validation report.](#)

For a detailed explanation of self-service app installation behavior on Victoria Experience, see [How self-service app installation works on Victoria Experience.](#)

### **Check app validation status**

The app validation status can be one of the following:

Status	Description
Vetting	The app package is in the validation process.
Approved	The app package has passed all AppInspect checks, or you have chosen to acknowledge the Splunk General Terms regarding potential impact of known issues and proceed with installation.
Installed	The app package is installed on your Splunk Cloud Platform deployment.



Status	Description
Rejected	The app package did not pass AppInspect validation checks. This means that some checks failed, or manual checks were detected that might require a manual review by the Splunk AppInspect team. To see the results of AppInspect validation, click <b>View Report</b> . For assistance with manual app vetting, contact Splunk Support.
Failed message	The app package validation did not complete due to some issues, for example, issues with the AppInspect service. Click <b>More Info</b> to find out why the package failed validation.
Check Status (Victoria Experience only)	This appears if you log out then log back in while an app is in the vetting process, or if the logged in user does not match the user who uploaded the app. In either case, to continue the app vetting process and view the report, click <b>Check Status</b> and enter the appropriate credentials.

### **View app validation report (Victoria Experience)**

1. Click **View Report**.
2. Review the details of the app validation report to determine why AppInspect rejected the package. If the report shows a value greater than zero for Manual Checks, you must contact Splunk Support to install your app.
3. Fix any issues specified in the report and upload your app again.

### **Update a private app (Victoria Experience)**

1. Click **Upload App** and select the updated version of your app.
2. Verify that the app status is **Approved** in the **Uploaded Apps** table. If the app status shows **Rejected**, review the app validation report. If the report shows a value greater than zero for Manual Checks, you must contact Splunk Support to install your app. See [View app validation report \(Victoria Experience\)](#).
3. Click **Install** on the later version of the app to install the later version.
4. Go to the **Apps** tab to see that the later version of your private app is listed in the **Apps** table.

If you want to downgrade to an earlier version of an app, you must first uninstall the app, then upload the earlier version.

Unlike Classic Experience, Victoria Experience does not support download of previously uploaded apps in the Splunk Web UI.

## **How self-service app installation works in Victoria Experience**

When you install an app using self-service app installation in Victoria Experience, the app is automatically installed on all standalone search heads, search head cluster members, and premium search heads running premium apps, such as Splunk IT Service Intelligence (ITSI) and Splunk Enterprise Security (ES). The app is also automatically installed on indexers.

All app configuration files are installed on search heads, while only configuration files that have indexing-related functionality are installed on indexers, including `indexes.conf`, `props.conf`, and `transforms.conf`.

This app installation behavior means that default knowledge objects within an app will be available on all search heads across a deployment. For example, if you install an app containing search-time field extractions, those fields can be used by searches dispatched on any search head.

It also means, in some cases, you might need to enable or disable specific app features, depending on their default settings. For example, if an app ships with lookup-populating scheduled searches disabled by default, you must enable those scheduled searches after installation on search heads where the lookup is required. For more information, see

[Configure self-service apps in Victoria Experience.](#)

### **Post migration app setup changes and considerations in Victoria Experience**

Victoria Experience introduces some changes to the underlying Splunk Cloud Platform architecture that enhance scalability and enable self-service management of nearly all Splunk apps. As a result, there are some changes in self-service app setup post-migration that are important to consider when migrating from Classic Experience to Victoria Experience.

Post-migration app setup changes apply only to deployments with more than one search head group, for example, a deployment that has a standalone search head and a premium search head (or search head cluster). These changes do not apply to deployments that have only one search head or one search head cluster.

The following sections provide a summary of self-service app setup changes and considerations in Victoria Experience.

#### **All apps are visible on all search heads**

After migration from Classic Experience to Victoria Experience, all globally shared apps become visible to all users on all search heads. This means that users on any search head can now see all apps on any other search heads across your deployment, including all apps on standalone search heads, premium search heads, and search head clusters.

To prevent users from seeing an app, you can limit app sharing to specific user roles after migration. For example, you can create a role for premium search head users, another role for search head cluster users, and so on. Then, within each app's permission page, remove read access from "Everyone" and grant read/write access to users' specific roles accordingly.

You can manage roles and app permissions for app visibility in the Splunk Web UI. For more information on roles, see [Create and manage roles with Splunk Web](#). For information on configuring app permissions, see [Managing app and add-on configurations and properties](#).

You can also manage roles and app permissions for app visibility programmatically using the Admin Config Service (ACS) API. See [Manage users, roles, and capabilities in Splunk Cloud Platform](#) and [Manage app permissions in Splunk Cloud Platform](#) in the *Admin Config Service Manual*.

Additionally, you can manage app permissions for app visibility in bulk using the ACS CLI. For more information, see [Run ACS CLI bulk operations](#) in the *Admin Config Service Manual*.

#### **Global app assets are available on all search heads**

After migration to Victoria Experience, all globally shared app assets, such as lookups and other knowledge objects, become available on all search heads. While app asset sharing behavior does not change in Victoria Experience, bringing together globally shared apps on the search head tier during migration can introduce conflicts between shared assets that have the same name.

For example, an older version of a lookup named "LookupA" in one app can interfere with a newer version of "LookupA" in another app, causing recent lookup updates to be missed.

Globally shared assets across apps follow standard Splunk platform configuration context and precedence. For any conflicting assets where two or more assets share the same global context and have the exact same name, Splunk configuration file precedence merges settings based on standard lexicographical ordering. See [Configuration file precedence](#) in the *Splunk Enterprise Admin Manual*.

To avoid post migration app asset conflicts, make sure that all globally shared app assets have unique names. Also, consider limiting app asset sharing to the local app context only, when possible, rather than sharing assets globally, to prevent duplicate app asset naming conflicts. For more information, see [Manage knowledge object permissions](#).

### **Scheduled searches are disabled on non-original search heads**

During a migration from Classic Experience to Victoria Experience, Splunk remembers the search head group on which each app is installed. After migration, Splunk disables any scheduled searches within apps that were not originally on that search head group.

For example, consider a deployment that includes a premium search head with App A, and a search head cluster without App A. After migration, App A now appears on both the premium search head and the search head cluster. However, App A's scheduled searches on the search head cluster are disabled.

To run an app's scheduled searches on the non-original search head group, you must explicitly enable the scheduled searches after migration on that search head group. This does not affect app performance on the original search head group and any existing scheduled searches remain operational there.

You can enable or disable scheduled searches on any search group at any time. Enabling or disabling an app's scheduled searches on one search head group has no affect on other search head groups. For example, enabling App A's scheduled searches on the search head cluster has no affect on the premium search head.

For instructions on how to enable or disable scheduled searches, see [Disable a report](#).

### **Global actions versus local actions**

Most edits made using the Splunk Web UI apply only to the specific search head group on which the change occurs. For example, if you enable or disable a scheduled search on a standalone search head, those changes are written to local files and apply only to that standalone search head. Likewise, if you enable or disable a scheduled search on a premium search head, running a premium app such as Enterprise Security, those changes apply only to that specific premium search head.

Syncing of local changes on a search head cluster occurs by default, as the cluster replicates local changes on one cluster member to all other cluster members. However, syncing of local files does not apply across different search head groups.

For more information on global versus local actions, see [Configure self-service apps in Victoria Experience](#).

### ***Configure self-service apps in Victoria Experience***

While self-service apps in Victoria Experience are installed, updated, enabled, disabled and deleted globally on all search heads and indexers, you configure self-service apps locally on individual search heads. This applies to most features that you configure using the Splunk Web UI or REST endpoints, including knowledge objects, such as dashboards, saved searches, field extractions, macros, eventtypes, tags, and so on, as well as modular and scripted inputs.

After you install an app, you might need to make configuration changes, such as disabling or enabling specific app features, depending on the apps default configuration settings. Some common scenarios that might require you to disable or enable app features on individual search heads include:

- **Alerts (email or modular):** For apps that come with email or modular alerts disabled by default, you must enable the alert on the individual search head most suited to running the workload. Likewise, for apps that come with email or modular alerts enabled by default, you must disable the alert on all but the "best" search head (or search

head cluster) within your deployment. For instructions on how to enable or disable email alerts, see [Define an email notification for an alert or scheduled report](#).

- **Lookups:** For apps that come with lookup-populating scheduled searches disabled by default, you must enable the scheduled searches on each search head where the lookup is required. Likewise, for apps that come with lookup-populating scheduled searches enabled by default, you must disable the lookup-populating scheduled searches on each search head where the lookup is not required. For instructions on how to enable or disable scheduled searches, see [Disable a report](#).
- **Dashboards:** For apps that come with dashboard-related scheduled searches disabled by default, you must enable the scheduled searches on each independent search head where you want to view the dashboard. Likewise, for apps that come with dashboard related scheduled searches enabled by default, you must either accept the performance overhead of scheduled searches running on every independent search head, or disable the scheduled search on each search head where the dashboard is not used. For instructions on how to enable or disable scheduled searches, see [Disable a report](#).
- **Datamodel acceleration:** Apps must ship with datamodels unaccelerated by default to pass AppInspect app validation. As a result, you must enable datamodel acceleration post-install on the particular search head where you want to use acceleration for a given data model. For instructions on how to enable datamodel acceleration, see [Enable data model acceleration](#).
- **Summary indexing:** For apps that come with summary-index-populating scheduled searches disabled by default, you must enable the scheduled search on the independent search head most suited to running the workload. Note that regardless of which search head populates the summary index, all independent search heads will be able to search the summary index. Likewise, for apps that come with summary-index-populating scheduled searches enabled by default, you must disable the scheduled search on all but the "best" search head (or search head cluster) within your deployment. For instructions on how to enable or disable summary indexing, see [Configure summary indexes](#).
- **Scripted inputs:** For apps that come with modular or scripted inputs disabled by default, you must enable the inputs on any search head or search head cluster where you want to run the data input. Likewise, for apps that come with modular or scripted inputs enabled by default, you must disable the inputs on any search head on which you do not want to run the input. For more information on scripted inputs, see [Get data from APIs and other remote data interfaces through scripted inputs](#).

## Install a private app on Classic Experience

If your Splunk Cloud Platform deployment is on Classic Experience, you can upload and install your private app using the **Upload App** workflow in Splunk Web.

Splunk Cloud Platform does not support self-service installation of private apps on IDM. If you are on Classic Experience, and your private app contains modular or scripted inputs that require installation on IDM, you must contact Splunk Support and submit a Cloud App Request to upload your app.

When you install an app using self-service app installation on Classic Experience, the app is automatically installed on all regular search heads and search head cluster members across your deployment. The app is also installed on indexers.

Classic Experience does not support self-service app installation on premium search heads, such as those running IT Service Intelligence (ITSI) or Enterprise Security (ES). If your deployment is on Classic Experience, to install any app on a premium search head, you must contact Splunk Support.

To install a private app on a Classic Experience deployment:

1. In Splunk Web, click the **Apps** gear icon.
2. Open the **Uploaded Apps** tab, and click **Upload App**.
3. Enter your splunk.com account credentials. Splunk Cloud Platform uses these credentials to authenticate your AppInspect app validation.
4. Click **Agree and Login**.  
This confirms that you accept the specified license conditions and submits your login credentials.
5. Select your private app package and click **Upload**.  
Your app appears in the Uploaded Apps table. Splunk Cloud Platform automatically runs your app through AppInspect validation to confirm that it meets Splunk Cloud Platform requirements.

App	Status	Actions	Date Sub
recovered_content.txt	Failed <a href="#">More Info</a>	<a href="#">Delete</a>	11/16/202
Test Upgrade Private App for Noah	Approved	<a href="#">Install</a> <a href="#">Delete</a> <a href="#">View Report</a>	11/16/202
av_error_app.tar.gz	Rejected	<a href="#">Delete</a> <a href="#">View Report</a>	11/16/202
music-app-for-splunk_100.tgz	Vetting		11/16/202

6. In the Uploaded Apps table, check the app validation status. Your app must pass all AppInspect checks and be approved before you can install it. For more information, see [Check app validation status](#).
7. If your app validation status is approved, click **Install**. If your app validation status is rejected, click **View Report** to determine the issues you must fix before you can install the app. For more information, see [View app validation report](#).
8. After you install your app, click the **Apps** tab to confirm that your private app is now listed in the **Apps** table. You can also see that the value for **App Origin** is **Uploaded**.

### **View app validation report (Classic Experience)**

1. Click **View Report**.
2. Review the details of the report to determine why AppInspect rejected the package. If the report shows a value greater than zero for Manual Checks, you must contact Splunk Support to install your app.
3. Fix the issues specified in the report and upload your app again.

### **Update a private app (Classic Experience)**

1. Click **Upload app** and select your updated app.
2. Verify that the app status is **Approved** in the **Uploaded Apps** table. If the app status shows **Rejected**, review the app validation report. If the report shows a value greater than zero for Manual Checks, you must contact Splunk

- Support to install your app. See [View app validation report \(Classic Experience\)](#).
3. Click **Install** to install an earlier version. Click **Update** to replace an installed app with a later version.
  4. Go to the **Apps** tab to see that the later version of your private app is listed in the **Apps** table.

If you want to downgrade to an earlier version of an app, you must first uninstall the app, then upload the earlier version.

### **View app install log (Classic Experience)**

As of version 8.2.2107 the Install Log tab has been removed from the app management page in Splunk Web. However, you can still view the same app install log information by running the following `rest` search command:

```
| rest /services/dmc/changes count=0 output_mode=json state=deployed
```

## **Manage lookups in Splunk Cloud Platform**

In Splunk Cloud Platform, on both Victoria Experience and Classic Experience, you cannot update an existing lookup file when you upgrade an app using the Splunk Web UI, the Admin Config Service (ACS API), or the Splunk REST API.

To update an existing lookup file in an app, you can:

- Edit the lookup file using the Splunk App for Lookup File Editing. See [Splunk App for Lookup File Editing on Splunkbase](#).
- Use the `outputlookup` command to update the lookup. See `outputlookup` in the *Search Reference*.

Lookup file behavior during app upgrade differs on Victoria Experience and Classic Experience deployments. The following sections describe default lookup behaviors during app upgrade, with workarounds you can use to update or retain existing lookup files when upgrading apps.

### **Lookup file behavior during app upgrade on Victoria Experience**

On Victoria Experience 8.2.2106 and higher, when you upgrade an app, Splunk Cloud Platform preserves any existing lookup files as is. If the new version of the app contains an update to an existing lookup, that update is ignored, but if the new version contains a new lookup that doesn't yet exist in the app, it will deploy it. For example, if myapp v1.0 has myLookupA, and myapp v2.0 has an updated myLookupA and a new myLookupB, on upgrade myLookupA is not updated, but myLookupB is added.

When upgrading an app that includes updated lookup files, use one of the following workarounds to update those lookup files:

- Use the Splunk App for Lookup File Editing to upload the updated lookup file.
- Delete the lookup file after app upgrade using the Splunk Web UI. This will update the old lookup to the new lookup.

### **Lookup file behavior during app upgrade on Classic Experience**

Lookup file behavior during app upgrade on Classic Experience differs on single instance, single search head, and search head cluster deployments.

When you upgrade an app on a single instance deployment, by default Splunk Cloud Platform overwrites any existing lookup files with the updated lookup files. In this case, the content of the existing lookup file is replaced with the updated

content, and no new lookup file is created. When upgrading an app on a single instance, you can use the following workaround to retain an existing enriched lookup file (a lookup file modified by the customer):

- Backup the lookup file before upgrade, and upload the backed up lookup to the app after upgrade.

When you upgrade an app on a single search head, by default Splunk cloud platform excludes all lookup file updates. In this case, all updates to lookup files are ignored and no lookup files are added or removed. When upgrading an app that contains updated lookups on a single search head, you can use the following workaround to update the lookup files:

- Use the Splunk App for Lookup File Editing to upload the updated lookup file.

When you upgrade an app on a search head cluster, by default Splunk Cloud Platform preserves any existing lookup files as is. If the new version of the app contains an update to an existing lookup, that update is ignored, but if the new version contains a new lookup that doesn't yet exist in the app, it will deploy it. When upgrading an app that contains updated lookups on a search head cluster, you can use the following workaround to update the lookup files:

- Use the Splunk App for Lookup File Editing to upload the updated lookup file.

For more information on lookups, see About lookups in the *Knowledge Manager Manual*.

## Configuration file reload triggers in app.conf

Splunk apps can contain a combination of Splunk Enterprise core configuration files and custom configuration files, such as those created by app developers for both private apps and public apps on Splunkbase. Whether these configuration files reload when you install an app or make configuration changes depends on reload trigger settings in `app.conf`.

Many Splunk Enterprise core configuration files reload by default on app installation or when configuration updates occur. These files have a reload setting under the `[triggers]` stanza in `$SPLUNK_HOME/etc/system/default/app.conf`, which causes them to reload automatically.

A custom configuration file is by definition any configuration file that does not have a corresponding `.spec` file in `$SPLUNK_HOME/etc/system/README`. This includes custom configuration files found in third party apps, such as `aws_settings.conf`, `service_now.conf`, `eventgen.conf`, and so on.

All custom configuration files reload by default, unless the file has a custom reload trigger in `app.conf`. For example, in the Splunk Security Essentials app, `app.conf` contains the following custom reload trigger: `reload.ssenav = http_get /SSEResetLocalNav`. When you install an app or update configurations for an app that has a custom reload trigger in `app.conf`, Splunk software tries to honor the custom reload trigger setting. If the custom reload trigger fails, then a rolling restart occurs.

If a custom configuration file does not have a reload trigger specified in `app.conf`, the default behavior is to restart for unknown configs. If a restart is not required, you can set the conf level trigger in `app.conf` to `reload.<conf_file_name> = simple`.

For detailed information on how to configure reload trigger settings for configuration files, see `app.conf` in the *Admin Manual*.

For more information on restart vs. reload behavior of Splunk Enterprise core configuration files, see Restart or reload after configuration bundle push? in the Splunk Enterprise documentation.



## Stanza-level reload triggers for inputs.conf

Stanza-level reload triggers enable the reload of only those specific configuration file stanzas that change when a configuration update occurs. This lets admins perform more efficient configuration updates based on which stanzas in the configuration file will change.

Stanza-level reload currently applies to a subset of stanzas in `inputs.conf` only. Any `inputs.conf` stanza that has a `reload.<conf_file_name>.<conf_stanza_prefix>` entry under the `[triggers]` stanza in `app.conf` will reload when changes are made to the specified stanza. Changes made to any `inputs.conf` stanzas that are not specified in a stanza-level reload entry will trigger a rolling restart.

Stanza-level reload for `inputs.conf` applies only when pushing changes to the configuration bundle in the indexer clustering context.

The following stanzas are reloadable in `inputs.conf`:

.conf file name	stanza prefix	Reload or restart
inputs.conf	http	reload
inputs.conf	script	reload
inputs.conf	monitor	reload
inputs.conf	<modular_input>	reload
inputs.conf	batch	reload

For detailed information on stanza-level reload triggers, see `app.conf` in the Splunk Enterprise documentation.

## Disable reload triggers in app.conf

You can disable both `.conf`-level reload triggers and stanza-level reload triggers by specifying the value `never` for any reload trigger entry in `app.conf`. Any reload trigger entry with a value of `never` will trigger a rolling restart when configuration changes occur. This can be useful if for any reason you want a specific configuration change to trigger a rolling restart.

For more information on configuring reload triggers, see `app.conf` in the Splunk Enterprise documentation.

For a listing of restart vs. reload behavior of frequently used apps and configuration files in Splunk Cloud Platform, see [Restart versus reload behavior of common apps and .conf files](#).

## Manage the Splunk Product Guidance app on your Splunk Cloud Platform deployment

Splunk Product Guidance (SPG) is an application that augments your current Splunk Cloud Platform deployment with contextually appropriate guidance for numerous use cases and tasks based on product usage.

### SPG in Splunk Cloud Platform

SPG is available on versions 8.2.2109 and higher as one of the default apps in the Splunk Cloud Platform.



SPG is only available on Splunk Cloud Platform deployments that don't use a premium app like Splunk Enterprise Security or Splunk IT Service Intelligence. For more information about the premium app solution subscriptions, see Splunk premium solutions in the Splunk Cloud Platform Service Description.

## Disable or re-enable the SPG app

Use the **Manage Apps** page to manage the SPG app.

You must be a Splunk Cloud Platform administrator with `sc_admin` permissions to disable and re-enable the SPG app for your Splunk Cloud Platform deployment.

If you are running a search head cluster using the Classic Experience, you must file support ticket to disable the SPG app. If you have a support contract, log in and file a new case using the Splunk Support Portal. Otherwise, contact Customer Support.

Complete the following steps:

1. Access the **Manage Apps** page by one of the following methods:
  1. From the Splunk Home page, click the Manage Apps icon (⚙️) in the **Apps** panel.
  2. From anywhere in Splunk Web, click **Apps** in the top menu bar, then **Manage Apps**.
2. Locate **Splunk Product Guidance** in the list of displayed apps.
3. Click **Disable** in the **Status** column.

To re-enable a disabled SPG app, follow the procedure and click **Enable** in step 3.

## Manage a rolling restart in Splunk Cloud Platform

Some configuration updates can cause the indexers in your Splunk Cloud Platform deployment to begin a process called a rolling restart. To minimize the impact of a rolling restart, deploy these updates during off-peak hours.

### What users experience during a rolling restart

A rolling restart is a sequential restart of Splunk indexers that allows indexing to continue during the restart process.

While indexing remains available at all times during a rolling restart, non-Splunk clients that do not follow best practices for retrying connections and managing backpressure might be impacted by an individual node restarting. Using forwarders or other types of load balancers, rather than network inputs alone, increases the robustness of your indexing during a rolling restart.

Searches still run during a rolling restart, but they might return incomplete results. Users running searches in Splunk Web receive a message warning of incomplete search results.

### What triggers a rolling restart

Deploying certain changes triggers a rolling restart. Examples of changes that trigger a rolling restart include, but are not limited to, the following tasks:

- Deleting the last HEC token (which deletes the app, causing a rolling restart).

- Installing some apps and add-ons. See [Restart versus reload behavior of common apps and .conf files](#).

Deploying a seemingly safe change can indirectly trigger a rolling restart. For example, adding an index doesn't trigger a restart by itself. But if you or another admin has made other changes that trigger a rolling restart and not deployed them, then when you deploy your change that adds an index, you also deploy the previous changes that trigger the rolling restart.

## Restart versus reload behavior of common apps and .conf files

Most configuration files do not trigger a restart when configuration changes occur, but instead trigger a less time-consuming file reload. To minimize service disruptions, before you install apps or deploy configuration updates in Splunk Cloud Platform, consider the restart versus reload behavior of relevant apps and configuration files.

For more information on configuration file reload behavior, see [Configuration file reload triggers in app.conf](#).

The following tables list some common apps and configuration files and show whether they trigger a restart or a reload.

- [Reload or restart behavior of common .conf files](#)
- [Reload behavior of common apps](#)

### **Reload or restart behavior of common .conf files**

Most Splunk configuration files are now reloadable. The following table shows the reload or restart behavior of some frequently used configuration files in Splunk Cloud Platform:

<b>.conf file name</b>	<b>Used for</b>	<b>Reload or restart</b>
authorize.conf	This file is used to configure roles and granular access controls.	reload
collections.conf	This file is used to configure KV store settings for a given app.	reload
distsearch.conf	This file is used to configure attributes and values you can use to configure distributed search.	reload
indexes.conf	This file is used to configure indexes and their properties.  For a list of specific changes to this file that require a restart, see Determine which indexes.conf changes require a restart in the Splunk Enterprise documentation.	reload/restart
inputs.conf	This file is used for HEC CRUD operations, configuring tcp ports for forwarders, configuring scripted inputs for apps, and configuring file system monitoring.  Splunk Cloud Platform supports stanza-level reload for <code>inputs.conf</code> . For more information on stanza-level reload, including a list of reloadable stanzas, see <a href="#">Stanza-level reload triggers for inputs.conf</a> .	reload/restart
multikv.conf	This file is used to configure multikv rules for extracting events from table-like events, such as the output of top, ps, ls, netstat, etc.	reload
props.conf	This file is used to set indexing property configuration, including timezone offset, custom source type rules, and pattern collision properties. Also, map transforms to even properties.	reload
restmap.conf	This file is used to create custom REST endpoints.	reload

.conf file name	Used for	Reload or restart
server.conf	This file is used to configure which settings should be replicated within a search head cluster.  Changes to the <code>[shclustering]</code> stanza require reload only. All other changes to <code>server.conf</code> require a restart.	reload/restart
transforms.conf	This file is used to configure regex transformations to perform on data inputs. Use in tandem with <code>props.conf</code> .	reload
ui-tour	This file is used to configure in-product tours of Splunk software features.	reload
web.conf	This file is used to configure tcp port to listen to incoming connections, <code>appserverports</code> , <code>connectiontimeout</code> .	reload
wmi.conf	This file is used to configure access to Windows Management Instrumentation (WMI).	reload

### **Reload behavior of common apps**

The following table shows the reload behavior of frequently used apps and add-ons in Splunk Cloud Platform:

This list pertains to the specified version of each app. Changes made to an app's configuration settings in subsequent app versions might trigger a restart instead of a reload.

App name	Version	Used for	Reload
Cisco Networks Add-on for Splunk Enterprise	2.5.8	This add-on sets the correct sourcetype and fields for identifying data from Cisco IOS, IOS XE, IOS XR, NX-OS devices in Splunk® Enterprise.	reload
Force Directed App For Splunk	3.0.1	The Force Directed App For Splunk helps you graph out attack paths and review links in your data. Built on D3 this app will allow you to search any form of data that has a source and target.	reload
Lookup File Editor	3.3.2	This app provides an Excel-like interface for editing, importing, and exporting lookup files (both KV store and CSV based lookups)	reload
Palo Alto Networks Add-on for Splunk	6.1.1	This add-on collects and correlates data from Firewalls, Panorama, Traps Endpoints, Aperture SaaS Security, AutoFocus, MineMeld, and WildFire.	reload
Palo Alto Networks App for Splunk	6.1.1	This app combines Palo Alto Networks security platform features with Splunk's investigation and visualization capabilities to provide advanced security reporting and analysis.	reload
Python for Scientific Computing (for Linux 64-bit)	1.4	This add-on contains a Python interpreter bundled with the following scientific and machine learning libraries: <code>numpy</code> , <code>scipy</code> , <code>pandas</code> , <code>scikit-learn</code> , and <code>statsmodels</code> . With this add-on, you can import these powerful libraries in your own custom search commands, custom rest endpoints, modular inputs, and so forth.	reload
Punchcard Custom Visualization	1.3.0	This Punchcard Custom Visualization app provides interactive ways to visualize and investigate cyclical trends in your data.	reload
Qualys Technology Add-on (TA) for Splunk	1.4.3	This add-on provides pre-built inputs for Qualys Cloud Platform data.	reload
Splunk Add-on for Amazon Web Services	4.6.0	This add-on lets Splunk admins collect data from AWS accounts, including configuration details, EC2 instance and EBS metadata, compliance information, CloudWatch log data, performance and billing metrics, S3 bucket stats, and more.	reload
Splunk Add-on for Cisco ASA	3.4.0	The Splunk Add-on for Cisco ASA allows a Splunk software administrator to map Cisco ASA devices, Cisco PIX, and Cisco FWSM events to the Splunk CIM.	reload

App name	Version	Used for	Reload
Splunk Add-on for Microsoft Cloud Services	3.1.0	This add-on lets Splunk admins pull activity logs, service status, operational messages, Azure audit, Azure resource data and Azure Storage Table and Blob data from a variety of Microsoft cloud services using the Office 365 Management APIs, Azure Service Management APIs and Azure Storage API.	reload
Splunk Add-on for Microsoft Office 365	1.1.0	This add-on lets Splunk admins pull service status, service messages, and management activity logs from the Office 365 Management API.	reload
Splunk Add-on for Microsoft Windows	6.0.0	This add-on provides predefined inputs to collect data from Windows systems and maps data to the Common Information Model.	reload
Splunk Add-on for Unix and Linux	6.0.2	The Splunk Add-on for Unix and Linux allows a Splunk software administrator to collect *nix data from *nix hosts.	reload
Splunk App for AWS	5.1.3	This app provides insight into your Amazon Web Services account. The app includes pre-built dashboards, reports, and alerts that provide real-time visibility into your AWS environment, including your AWS Config, CloudWatch, CloudTrail, Billing, S3, VPC Flow Log, Amazon Inspector, and Metadata inputs.	reload
Splunk App for Windows Infrastructure	1.5.2	This app provides pre-built data inputs, searches, reports, and dashboards that let you monitor, manage, and troubleshoot Windows operating systems, including Active Directory elements, from a single location.	reload
Splunk Common Information Model (CIM)	4.13.0	This add-on contains a collection of pre-configured data models that support the consistent, normalized treatment of data for maximum efficiency at search time.	reload
Splunk Dashboard Examples	7.3.0	The Splunk Dashboard app delivers examples that give you a hands-on way to learn the basic concepts and tools needed to rapidly create rich dashboards using Simple XML.	reload
Splunk Datasets Add-on	1.0	This app delivers new SPL commands, custom visualizations, assistants, and examples to explore a variety of machine learning concepts.	reload
Splunk Machine Learning Toolkit	5.2.0	This add-on provides an intuitive interface to build, edit and analyze table datasets (tables) without SPL.	reload
Splunk Sankey Diagram - Custom Visualizations	1.5.0	Sankey diagrams show metric flows and category relationships. You can use a Sankey diagram to visualize relationship density and trends.	reload
Splunk Supporting Add-on for Active Directory	2.2.1	This app provides support functions to the Windows Infrastructure, Active Directory, and Exchange apps that enable you to extract information from an Active Directory database.	reload
Splunk Timeline - Custom Visualization	1.4.0	A timeline visualization shows activity time intervals and discrete events for a resource set.	reload

## Guidance for managing a rolling restart

To minimize impact to users, deploy configuration changes during times that are off peak for both indexing and searching. You can identify off-peak times from the Snapshots in your Splunk Cloud Monitoring Console. See [Monitor your Splunk Cloud Platform Deployment](#).

During a rolling restart, monitor indexing and search performance with the Splunk Cloud Monitoring Console.

## More information

For more information about how a rolling restart works, see Perform a rolling restart of an indexer cluster in the Splunk Enterprise documentation. Note that some of the advanced options are not available by default in Splunk Cloud Platform.

# Configure Search Settings in Splunk Cloud Platform

## Configure hybrid search

An on-premises Splunk Enterprise search head can connect to both a set of on-premises indexers and a Splunk Cloud Platform indexer cluster on Classic Experience. The search head can then run **hybrid searches** that combine on-premises data with data from Splunk Cloud Platform.

End of life for hybrid search is targeted for October 30, 2024. Customers who use hybrid search must migrate to federated search. See [Migrate from hybrid search to Federated Search for Splunk in \*Federated Search\*](#). After migrating to federated search, contact Splunk customer support to disable hybrid search on your Splunk Cloud Platform deployment.

To search across both on-premises and Splunk Cloud Platform data with hybrid search, you must run the search from an on-premises search head. A Splunk Cloud Platform search head can only search data on Splunk Cloud Platform.

## Hybrid search limitations

The following conditions and limitations apply to hybrid search:

- End of life for hybrid search is targeted for October 30, 2024.
- You must run hybrid searches from an on-premises search head. You cannot run a hybrid search from a Splunk Cloud Platform search head.
- The on-premises search head must be compatible with the target Splunk Cloud Platform version. For more information, see [Supported hybrid search versions in the \*Splunk Cloud Platform Service Description\*](#).
- Only ad-hoc searches are supported. Scheduled searches are not supported.
- You cannot install a Splunk Premium Solution on a hybrid search head. However, you can run a hybrid search against a Splunk Cloud Platform stack that includes a premium solution, as long as the hybrid search head running the hybrid search complies with all necessary conditions and limitations. See [Splunk premium solutions in the \*Splunk Cloud Platform Service Description\*](#) for a complete list of premium solutions.
- You cannot initiate searches from an on-premises Splunk Enterprise search head to multiple Splunk Cloud Platform environments.

See also Hybrid search in the [Splunk Cloud Platform Service Description](#).

## Expand your cross-deployment search options with federated search

Federated search is an improvement on hybrid search that expands your ability to search across Splunk deployments. If you are considering a move to hybrid search, consider federated search instead.

- Federated search does not require an on-premises search head. You can configure federated search between Splunk Cloud Platform deployments.
- Federated search can be set up between a single "local" Splunk deployment and multiple "remote" Splunk deployments.
- Federated search supports scheduled searches.
- Federated search supports all search management tier architecture options. This means that it allows search of Splunk Cloud Platform deployments with search head cluster configurations.
- In most cases you can configure federated search between an on-premises deployment and a Splunk Cloud Platform deployment without contacting a Splunk Support representative.

- After you migrate to federated search, open a support ticket to disable hybrid search in your Splunk Cloud Platform deployment.

See Migrate from hybrid search to Federated Search for Splunk in *Federated Search*.

## Enable hybrid search

Complete the following steps to enable hybrid search

This procedure is valid only for an on-premises standalone search head that is not part of either an on-premises indexer cluster or an on-premises search head cluster.

1. Confirm that the on-premises search head is already configured to search across on-premises indexers. To learn how to configure a search head to connect with on-premises indexers, see Deploy a distributed search environment in the Splunk Enterprise *Distributed Search* manual.
2. Confirm that the on-premises search head is compatible with the target Splunk Cloud Platform version, as stated in [Hybrid search limitations](#). If necessary, upgrade the search head to the Splunk Cloud version.
3. Contact your Splunk account representative to enable hybrid search for your Splunk Cloud Platform instance. Be sure to provide the public IP address(es) of the hybrid search head(s) so that access lists in the Splunk Cloud Platform environment can be created. In addition, specify that you need:
  - ◆ A 1 MB Splunk Enterprise license for the on-premises search head that you want to use for hybrid search.
  - ◆ The URI for the manager node of the Splunk Cloud Platform indexer cluster.
  - ◆ The security key for the Splunk Cloud Platform indexer cluster.

If necessary for your deployment, your representative can help you open a case with Splunk Support to enable hybrid search using the Splunk Support portal.

4. Install the 1 MB license on the on-premises search head. See [Install a license](#).
5. Add the following lines to the `server.conf` file on the on-premises search head. The example shows the required TCP port for `manager_uri`: 8089.

```
[general]
site = site0

[clustering]
multisite = true
manager_uri = <manager node URI in the format https://c0m1.<stack name>.splunkcloud.com:8089>
mode = searchhead
pass4SymmKey = <security key>
```

6. Restart the search head.
7. Run a search command like the following, which retrieves Splunk log events and lists the indexers that the events come from:

```
index = *_* | stats count by splunk_server.
```

If hybrid search is configured correctly, indexers from both your Splunk Enterprise and your Splunk Cloud Platform deployments are listed in the results.

## Disable hybrid search

To disable hybrid search:

1. Remove the following lines from the `server.conf` file on the on-premises search head.

```
manager_uri = <manager node URI in the format https://c0m1.<stack name>.splunkcloud.com:8089>
pass4SymmKey = <security key>
```

2. Restart the search head.
3. Run a search command like the following, which retrieves Splunk log events and lists the indexers that the events come from:

```
index = *_* | stats count by splunk_server.
```

If you've correctly disabled hybrid search, the search results show indexers only from your on-premises Splunk Enterprise search head. The results should not include indexers from Splunk Cloud Platform deployments.

Splunk Customer Support will assist you in disabling hybrid search functionality configured for your Splunk Cloud Platform deployment. If you have a support contract, log in and file a new case using the Splunk Support Portal. Otherwise, contact Splunk Customer Support.

## Set limits for concurrent scheduled searches

When you run searches, Splunk limits the number of concurrent searches to preserve the performance for each search. In Splunk Cloud Platform, this concurrent limit is configured for you. Note the following defaults:

- 50% of your concurrent searches are scheduled searches.
- 25% of your concurrent searches are summarization searches.

You can adjust the relative number of concurrent scheduled and summarization searches to meet the needs of your organization. For example:

- If you primarily run scheduled searches, consider making most of your searches scheduled searches.
- If you use primarily data model acceleration searches, consider allocating a larger percentage of searches to summarization searches.

You can adjust limits for concurrent scheduled searches and summarization searches in Splunk Web using the **Settings > Server settings > Search preferences** page. Before adjusting these concurrency limits, note the following:

- You must have the `edit_search_concurrency_scheduled` capabilities to configure these settings.
- If your Splunk Cloud Platform deployment is on the Victoria Experience, you can configure concurrency limits for both search heads and search head clusters.
- If your Splunk Cloud Platform deployment is on the Classic Experience and a search head cluster, you must click **Show All Settings** for server settings to be available. Concurrency limits are replicated between search head cluster members.

For more information about the Experience designations, see [Determine your Splunk Cloud Platform Experience](#).

When you configure limits, be sure to first test your settings in a test environment to ensure they are optimal for your production environment. For example:

- If you set a limit of 70% of your searches to be scheduled searches, ad-hoc searches may run more slowly.
- If you set the limits for scheduled searches too low, certain scheduled searches may be skipped.

There are other concurrency settings which can impact your search resource allocation. For an overview of how searches are prioritized based on all configured concurrency settings, see [Configure the priority of scheduled reports](#). The following sections describe scheduled and summarization searches and how you can modify their default values.

## Scheduled searches

By default, the scheduled searches value is set to 50% in the **Relative concurrency limits for scheduled searches** field on the **Search preferences** page. These are the searches that you schedule using the Search Scheduler or that are created as a part of report acceleration or data model acceleration. Ad-hoc searches are not included in this group. For this limit, you set the value as a percentage of your total searches. Consider the following examples:

- You set the value to 70%. This means that out of available searches, at most scheduled searches (including user-scheduled and summarization searches) can use up to 70%.
- Your concurrent limit of 50% results in 19 scheduled searches and you change the settings to 75%. This results in the number of allowed concurrent scheduled searches changing to 27.

## Summarization searches

By default, the summarization searches value is set to 50% in the **Relative concurrency limits for summarization searches** field on the **Search preferences** page. These are searches that are generated for report acceleration or data model acceleration. For this limit, you set the value as a percentage of the searches allocated for scheduled searches. Consider the following example:

You set the value to 50%. This means that at most acceleration searches can use up to 50% of the searches allocated for scheduled searches. If you configure a scheduled search concurrency limit of 50%, this results in a limit of 18 concurrent searches. If you set the summarization searches limit to 50%, then the summarization searches are allocated roughly 25% of total searches or roughly 9 summarization searches.

### *Configure concurrent scheduled search limits*

1. In Splunk Web, click **Settings > Server settings > Search preferences**.
2. Specify an option for **Default search time range**.
3. In **Relative concurrency limits for scheduled searches**, select a percentage value for the concurrency limit for scheduled searches. This includes user-scheduled searches and summarization searches.
4. In **Relative concurrency limits for summarization searches**, select a percentage value for the concurrency limit for summarization searches. This value represents a percentage of the total allocated resources for scheduled searches.
5. Click **Save**.

Any changes that you make to the Default search time range, Relative concurrency limits for scheduled searches, and Relative concurrency limits for summarization searches values won't trigger a restart after you click Save.



# Manage Search Workloads in Splunk Cloud Platform

## Workload Management overview

This documentation applies to workload management in Splunk Cloud Platform only. For documentation that applies to workload management in Splunk Enterprise, see the Workload Management manual in the Splunk Enterprise documentation.

Workload management is a rule-based framework that enables the allocation of compute resources (CPU and memory) to search, indexing, and other workloads in Splunk Cloud Platform. You can use workload management to ensure that high priority searches receive adequate resources, while lower priority searches are appropriately restricted.

Workload management lets you:

- Isolate data-ingestion from the search workloads
- Prioritize critical search workloads
- Isolate resource-heavy searches to reduce impact on the overall system

## Requirements

- You must have the `sc_admin` role to see workload pools and workload rules in Splunk Cloud Platform.
- You must have the following capabilities to configure workload management: `list_workload_pools`, `list_workload_rules`, `edit_workload_rules` and `select_workload_pools`.

## How workload management works

Workload management lets you allocate CPU and memory resources to searches in logical containers called workload pools. You then define workload rules to place searches in different workload pools automatically. You can also define workload rules to monitor search runtime and perform automated remediation actions.

For example, you can create a workload rule that places searches from the security team in the high-priority workload pool, and create another rule to move those searches to the standard pool if the search runtime exceeds 2 minutes.

## Workload management concepts

Review the following key concepts and features before using workload management:

### ***Workload pools***

A workload pool is a logical container that allows prioritization of workloads in the pool. Splunk Cloud Platform provides three pre-defined workload pools for searches. Each pool is allocated a percentage of CPU and memory resources:

- **Standard:** All searches are assigned to this pool by default. You must use workload rules to place searches in other pools.
- **HighPriority:** Compared to the Standard pool, this pool is assigned a larger share of system resources. Workloads assigned to this pool are assigned a higher priority compared to executing in the Standard pool when system resources are in contention. However, you might still need to modify the search for better performance. For information about search optimization, see Search Optimization in the *Search* manual.

- **LowPriority:** Compared to the Standard pool, a relatively smaller share of system resources is assigned to this pool. Consequently, workloads assigned to this pool will execute with the lowest priority compared to the other two pools.

The following table shows the default allocation of Search resources among different pools. You cannot modify these values.

Search Category Pools (% of Search Resources):

Pool	CPU	Memory
Standard	35%	100%
HighPriority	60%	100%
LowPriority	5%	100%

**Tips**

1. When migrating to this version of Splunk Cloud Platform, if you do nothing, there is no change in your search priority. All of your searches will run in the Standard pool.
2. Selectively add workloads (by creating workload rules) to the HighPriority pool to ensure higher performance and speed for that workload in your priority pool. The HighPriority pool is intended to serve a few selected high priority searches. Assigning too many searches to the HighPriority pool will degrade the search performance.
3. Using workload pools helps to ensure that your priority searches have high performance. This means that searches in your Standard and LowPriority pools may degrade somewhat by comparison. This is expected behavior, and you may need to monitor and adjust rules to ensure that you get the best performance for the searches that matter most.

**Workload rules**

A workload rule contains a user-defined condition based on a set of predicates. For example, role=security AND search\_type=adhoc. When a search meets the user-defined condition, the rule is triggered and a specified action occurs. You can define workload rules to place searches in workload pools automatically, or create rules to monitor and perform remediation actions on long-running searches.

For more information on workload rules, see [Create workload rules](#).

**Admission rules**

Admission rules filter out searches automatically before they start based on a user-defined predicate (condition).

You can use admission rules to prevent the execution of rogue searches that might consume a large amount of resources and interfere with critical search workloads. You can also use admission rules to limit which roles, apps, and so on, can run searches over specific time ranges, such as peak business days.

For more information on admission rules, see [Create admission rules to prefilter searches](#).

## Configure workload rules

Workload rules provide an automated way to assign searches to workload pools and monitor running searches. Workload management evaluates workload rules in the order in which they are listed. If a search meets the predicate condition defined in a rule, a specified action is taken. Workload rules are evaluated for every new search and reevaluated every 10 seconds.

There are two types of workload rules:

- Search placement rules
- Search monitoring rules

Search placement rules determine the pool in which a search is placed when you start a search. Predicates that you can define to control search placement include `app`, `role`, `user`, `index`, `search_type`, `search_mode`, and `search_time_range`. You can use search placement rules to ensure that high-priority searches are assigned to pools that provide adequate resources, while low-priority searches are restricted.

Search monitoring rules automatically trigger actions on running searches based on the defined rule predicate and the status of the search. When you create a monitoring rule, you must specify a `runtime` value in the predicate. If a search exceeds the `runtime` value, workload management performs the specified action. Supported actions are `Abort search`, `Display in Messages`, and `Move search to alternate Pool`. You can use monitoring rules to manage heavy search loads and prevent rogue processes from monopolizing pool resources.

Workload rules apply to base searches only. Workload rules do not apply to subsearches.

## Create a workload rule in Splunk Web

To create a workload rule in Splunk Web:

1. In Splunk Web, click **Settings > Workload Management**.
2. Click **Add Workload Rule**.
3. Define the following fields to configure a new workload rule:

Field	Action
Name	Specify the name of the workload rule.
Predicate (Condition)	<p>Specify a predicate (condition) that must match to trigger this rule. The predicate syntax is <code>&lt;type&gt;=&lt;value&gt;</code> with optional AND, OR, NOT, (). For example, <code>app=search AND role=power</code> triggers all searches belonging to both the Search app and the power role.</p> <p>Valid predicate types are <code>app</code>, <code>role</code>, <code>index</code>, <code>user</code>, <code>search_type</code>, <code>search_mode</code>, <code>search_time_range</code>, and <code>runtime</code>.</p> <p>For supported predicate values, see <a href="#">Specify predicate values</a> in the next section on this page.</p> <p>In complex predicates, <code>AND</code>, <code>OR</code>, and <code>NOT</code> operators must be upper case. Lower case is not supported.</p>
Schedule	(Optional) Set a schedule for the workload rule. The schedule determines the time period during which the rule is valid.

Field	Action
	<p>If set to <code>Always On</code> (the default), the rule remains valid indefinitely and does not expire.</p> <p>If set to <code>Time Range</code>, the rule is valid during the specified time range only and expires when the time range ends.</p> <p>If set to <code>Every Day</code>, <code>Every Week</code>, or <code>Every Month</code>, the rule becomes valid on a recurring basis during the specified time range every day, on the specified days of the week, or on the specified days of the month.</p> <p>The schedule time for a workload rule is based on the system timezone, regardless of the timezone set for an individual user in the UI.</p>
Action	<p>Specify the action to perform when a search meets the predicate condition.</p> <p>Place search in a <code>Pool</code> (the default) assigns searches that meet the predicate condition to the specified workload pool.</p> <p><code>Abort search</code> kills the search process.</p> <p><code>Display a Message</code> shows a message in the job inspector to users that have all of the following required capabilities: <code>list_workload_pools</code>, <code>edit_workload_pools</code>, <code>list_workload_rules</code>, and <code>select_workload_rules</code>.</p> <p><code>Move search to alternate Pool</code> moves the running search to a different specified pool.</p> <p><code>Abort search</code>, <code>Display a Message</code>, and <code>Move search to alternate Pool</code> actions apply to in-progress searches only. You must specify a <code>runtime</code> condition to enable these actions. For example, the predicate <code>index=_internal AND runtime&gt;1m</code> triggers the specified action on all searches that contain <code>index=_internal</code> and run for more than one minute.</p> <p>The <code>Place search in a pool</code> action is not valid with rules containing a <code>runtime</code> condition. <code>Abort search</code>, <code>Display in Messages</code>, and <code>Move search to alternate Pool</code> actions are valid only when a <code>runtime</code> condition exists.</p>
Workload Pool	<p>Select the workload pool to which this rule applies.</p>
User Message	<p>Enter a custom message that notifies the end user when a search triggers the workload rule action. For example, "Search runtime exceeded 30 seconds. The search was moved to the high_perf pool."</p> <p>A user message is required with the <code>Display a Message</code> action, and is optional for other actions. Messages are limited to a maximum of 140 characters.</p> <p>When a search triggers the rule action, the user message appears in the Jobs manager in Splunk Web: Click <b>Activity &gt; Jobs &gt; Job</b>. It also appears under the <b>Job</b> menu in the Search app.</p>

4. Click **Submit**.

## Specify predicate values

The table describes valid predicate values for each type of workload rule predicate:

Predicate type	Valid values
app	<p>Name of the app. For example, <code>app=search</code></p> <p>The correct name to specify for an app is the name of the app directory located in <code>\$(SPLUNK_HOME)/etc/apps</code>. You can also find the correct name for an app in Splunk Web: Click <b>Apps &gt; Manage Apps</b>. See app names listed under <b>Folder name</b>. App names are case insensitive.</p>
role	<p>Name of the role. For example, <code>role=admin</code>.</p> <p>The <code>role</code> predicate supports a wildcard (*) in the rule definition. For example, a rule that contains <code>role=support_*</code> will match the role names <code>support_team1</code>, <code>support_team2</code> and so on. A rule that contains <code>role=*</code> will match all role name values.</p> <p>For information on how role inheritance impacts <code>role</code> predicate matches, see <i>Searches run by single user can match multiple roles</i> in the <i>Workload Management</i> manual in the Splunk Enterprise documentation.</p>
user	<p>Name of any valid user. For example, <code>user=bob</code>. The reserved internal user "nobody" is invalid. The reserved internal user "splunk-system-user" is valid.</p> <p>The <code>user</code> predicate supports a wildcard (*) in the rule definition. For example, a rule that contains <code>user=*</code> will match all user name values.</p>
index	<p>Name of the index. For example, <code>index=_internal</code>. Value can refer to internal or public index.</p> <p>The <code>index</code> predicate supports a wildcard (*) in the rule definition. For example, a rule that contains <code>index=support_*</code>, will match the index names <code>support_prod</code>, <code>support_test</code>, and so on. A rule that contains <code>index=_*</code> will match any index name that starts with an underscore, such as <code>_internal</code>, <code>_audit</code>, and <code>_introspection</code>.</p> <p><b>Note:</b> <code>index=*</code> is a special case, where * is not treated as a wildcard, but as a literal string match. For example, if a rule contains <code>index=*</code>, it will match the exact string <code>index=*</code> or <code>index=_*</code>. <code>index=*</code> will not match all internal indexes, such as <code>index=_internal</code>, <code>index=_audit</code>, and so on.</p>
search_type	Supports <code>adhoc</code> , <code>scheduled</code> , <code>datamodel_acceleration</code> , <code>report_acceleration</code> , and <code>summary_index</code> search types.
search_mode	Supports <code>realtime</code> and <code>historical</code> search modes.
search_time_range	<p><code>search_time_range</code> predicate values match against the time range of a search. For example, if you specify <code>search_time_range&gt;4h</code>, any search whose time range exceeds 4 hours is subject to the specified action.</p> <p>Supports <code>alltime</code> and numerical values with time units <code>m</code>, <code>minute</code>, <code>minutes</code>, <code>h</code>, <code>hour</code>, <code>hours</code>, <code>w</code>, <code>weeks</code>, <code>d</code>, <code>day</code>, or <code>days</code>. Supports comparison operators <code>=</code>, <code>&gt;</code>, and <code>&lt;</code>.</p> <p><code>search_time_range</code> predicate values match against time ranges selected using the Time Range Picker in the UI and time ranges defined by <code>earliest</code> and <code>latest</code> time modifiers for historical</p>

Predicate type	Valid values
	<p>searches specified in a search string. See <a href="#">Specify time modifiers in your search</a>.</p> <p>Searches that run over the "all time" time range only match against rules that specify the <code>alltime</code> value, and do not match against rules that specify numerical values. For example, an all time search defined by <code>earliest=0</code> and <code>latest=now</code> only matches rules that specify <code>search_time_range=alltime</code>, and does not match rules that specify a numerical value, such as <code>search_time_range&gt;10m</code>.</p> <p><b>Note:</b> <code>search-time_range</code> predicate values that contain an equal sign (=) might not match against search time ranges that use "snap to" time units. For example, while <code>earliest=-24h</code> and <code>latest=now</code> defines a time range of 24 hours, when using "snap to" time units, such as <code>earliest=-24h@h</code> and <code>latest=now</code>, the time range can be evaluated as greater than or equal to 24 hours, depending on what time the search runs. See <a href="#">Relative time modifiers that snap to a time</a>. To ensure <code>search_time_range</code> values match against time ranges with "snap to" time units, avoid using the = operator, and instead specify a value such as <code>search_time_range&gt;24h</code>.</p> <p><code>srchTimeWin</code> and <code>srchTimeEarliest</code> settings in <code>authorize.conf</code>, which let you set maximum time range and earliest event time for searches, respectively, do not apply to workload rules and admission rules. For example, if you run a search with a time range of 7 days, and you have an admission rule that specifies <code>search_time_range&gt;48h</code>, the search will be filtered out, even if the <code>srchTimeWin</code> value for that role is set to 24 hours or greater.</p>
runtime	<p>The amount of time that a search must run in a workload pool to trigger a specified action, such as <code>Abort search</code>, <code>Display in Messages</code>, or <code>Move search to alternate Pool</code>. For example, if you specify <code>runtime&gt;1m</code>, any search in the pool that runs for more than 1 minute is subject to the specified action.</p> <p>By default, <code>runtime</code> applies to both historical and real-time searches. If you want the <code>runtime</code> condition to apply to historical searches only or real-time searches only, you must specify the <code>search_mode</code> in the predicate condition. For example, if you specify <code>runtime&gt;5m AND search_mode=realtime</code>, only those realtime searches in the pool that run for more than 5 minutes are subject to the specified action.</p> <p>Valid time units for <code>runtime</code> values include <code>s</code>, <code>second</code>, <code>seconds</code>, <code>m</code>, <code>minute</code>, <code>minutes</code>, and <code>h</code>, <code>hour</code>, <code>hours</code>.</p> <p><b>Note:</b> <code>runtime</code> applies to workload rules only. It is not a valid predicate type for admission rules.</p>

For workload rule use case examples, see [Workload Management examples](#).

## Enable workload rules

You can enable or disable individual workload rules. This lets you create and save multiple different workload rules and apply them as needed. Individual workload rules are enabled by default when you create them. Disabled workload rules are not evaluated and have no effect on running searches.

To enable or disable an individual workload rule:

1. In Splunk Web, click **Settings > Workload Management > Workload Rules**.
2. In the Status column, toggle the switch to enable or disable the individual workload rule.

The workload management feature must be enabled for workload rules to apply to searches. In Splunk Cloud Platform, the workload management feature is enabled for the `sc_admin` role by default and cannot be disabled.

## Monitor triggered workload rule actions

When a running search triggers a workload rule action, information about the action appears in the Search job inspector. This includes the action that was taken on the search and the timestamp. If a search triggers multiple rules, the information appears in reverse chronological order.

To view details of a workload rule action:

1. In Splunk Web, click **Activity > Jobs**.
2. Find the specific search job and click **Job > Inspect Job > Search job properties**.
3. View details of the workload rule action under the `workload_action_information` property.

<code>sid</code>	1579042838.2401
<code>statusBuckets</code>	300
<code>ttr</code>	600
<code>workload_action_information</code>	Search 1579042838.2401 was moved from default_search_pool to high_priority_pool by rule message_test_rule on Tue Jan 14 15:00:47 2020.
<code>workload_pool</code>	high_priority_pool
Additional info	<a href="#">timeline</a>

To view the `workload_action_information` property, you must have `list_workload_pools` and `select_workload_pools` capabilities.

## Configure admission rules to prefilter searches

Admission rules let you filter out searches automatically before they start, based on a predicate condition that you define. If a search meets the specified condition, the search does not run.

You can use admission rules to prevent the execution of rogue searches, such as poorly written and potentially harmful searches that might consume an excessive amount of resources and interfere with critical search workloads. For example, you can create a rule to filter out wildcard searches that target all indexes, or filter out searches in the `alltime` time range.

You can also use admission rules to set up time-bound access to searches for roles, users, apps and so on. For example, you can create a rule that filters out all ad hoc searches from a certain role during peak business days, but allows the same role to run searches on weekends. And you can create rules that limit the number of concurrent ad hoc searches to retain search capacity for scheduled searches.

Admission rules are enabled by default for the `sc_admin` role. You can create a maximum of 100 admission rules per Splunk Cloud Platform deployment.

Admission rules apply to base searches only. Admission rules do not apply to subsearches.

## Admission rules evaluation behavior

Admission rules have no explicit ordering. Admission rules are evaluated for ad hoc searches at the the time the search is dispatched.

Admission rules for scheduled searches are evaluated asynchronously every 600s by default. As a result, changes to an admission rule or scheduled search can take up to 600s to become active.

If a search meets the conditions of a rule, the rule takes effect before the search is executed. If a search is already running, and you create a new admission rule that applies to that search, the running search is not affected by the new rule.

## Create an admission rule using Splunk Web

To create and edit admission rules, a user's role must have the `list_workload_rules` and `edit_workload_rules` capabilities.

To create an admission rule using Splunk Web:

1. In Splunk Web, click **Settings > Workload Management**.
2. Click **Add Admission Rule**.
3. Define the following fields to configure a new admission rule:

Field	Action
Name	Specify the name of the admission rule.
Predicate (Condition)	<p>Specify a predicate (condition) that must match to trigger this rule. The predicate syntax is <code>&lt;type&gt;=&lt;value&gt;</code> with optional AND, OR, NOT, (). For example, <code>app=search AND role=power</code> triggers all searches belonging to both the Search app and the power role.</p> <p>Valid predicate types are <code>app</code>, <code>role</code>, <code>index</code>, <code>user</code>, <code>search_type</code>, <code>search_mode</code>, <code>search_time_range</code>, and <code>adhoc_search_percentage</code>. For supported predicate values, see <a href="#">Specify predicate values on this page</a>.</p> <p>In complex predicates, AND, OR, and NOT operators must be upper case. Lower case is not supported.</p>
Schedule	<p>(Optional) Set a schedule for the admission rule. The schedule determines the time period during which the rule is valid.</p> <p>If set to <code>Always On</code> (the default), the rule remains valid indefinitely and does not expire.</p> <p>If set to <code>Time Range</code>, the rule is valid during the specified time range only and expires when the time range ends.</p> <p>If set to <code>Every Day</code>, <code>Every Week</code>, or <code>Every Month</code>, the rule becomes valid on a recurring basis during the specified time range every day, on the specified days of the week, or on the specified days of the month.</p> <p>The schedule time for an admission rule is based on the system timezone, regardless of the</p>



Field	Action
	timezone set for an individual user in the UI.
Action	<p>Specify the action to perform when a search meets the predicate condition.</p> <p>If set to <code>Filter search</code>, the rule filters out any search that meets the specified predicate condition.</p> <p>If set to <code>Queue search</code>, the rule places a search that meets the predicate condition in a queue to be re-run at a later time, when the predicate condition allows it. This action is currently allowed for <code>adhoc_search_percentage</code> predicate type rules only.</p>
User Message	<p>Enter a custom message that notifies the end user when a search triggers the admission rule action. For example, "This search meets specified admission rule conditions. The search was not executed."</p> <p>If an ad hoc search triggers the rule action, the custom message appears beneath the search bar in the Search and Reporting app. If a scheduled search triggers the action, a default message appears in <code>scheduler.log</code> only.</p>

4. Click **Submit**.

### Specify predicate values

The table describes valid values for each admission rule predicate type:

Predicate type	Valid values
app	<p>Name of the app. For example, <code>app=search</code></p> <p>The correct name to specify for an app is the name of the app directory located in <code>\$SPLUNK_HOME/etc/apps</code>. You can also find the correct name for an app in Splunk Web: Click <b>Apps &gt; Manage Apps</b>. See app names listed under <b>Folder name</b>. App names are case insensitive.</p>
role	<p>Name of the role. For example, <code>role=admin</code>.</p> <p>The <code>role</code> predicate supports a wildcard (*) in the rule definition. For example, a rule that contains <code>role=support_*</code> will match the role names <code>support_team1</code>, <code>support_team2</code> and so on. A rule that contains <code>role=*</code> will match all role name values.</p> <p>For information on how role inheritance impacts <code>role</code> predicate matches, see Searches run by single user can match multiple roles in the <i>Workload Management</i> manual in the Splunk Enterprise documentation.</p>
user	<p>Name of any valid user. For example, <code>user=bob</code>. The reserved internal user "nobody" is invalid; the reserved internal user "splunk-system-user" is valid.</p> <p>The <code>user</code> predicate supports a wildcard (*) in the rule definition. For example, a rule that contains <code>user=*</code> will match all user name values.</p>

Predicate type	Valid values
index	<p>Name of the index. For example, <code>index=_internal</code>. Value can refer to internal or public index.</p> <p>The <code>index</code> predicate supports a wildcard (*) in the rule definition. For example, a rule that contains <code>index=support_*</code>, will match the index names <code>support_prod</code>, <code>support_test</code>, and so on. A rule that contains <code>index=_*</code> will match any index name that starts with an underscore, such as <code>_internal</code>, <code>_audit</code>, and <code>_introspection</code>.</p> <p><b>Note:</b> <code>index=*</code> is a special case, where * is not treated as a wildcard, but as a literal string match. For example, if a rule contains <code>index=*</code>, it will match the exact string <code>index=*</code> or <code>index=_*</code>. <code>index=*</code> will not match all internal indexes, such as <code>index=_internal</code>, <code>index=_audit</code>, and so on.</p>
search_type	<p>Supports <code>adhoc</code>, <code>scheduled</code>, <code>datamodel_acceleration</code>, <code>report_acceleration</code>, and <code>summary_index</code> search types.</p>
search_mode	<p>Supports <code>realtime</code> and <code>historical</code> search modes.</p>
search_time_range	<p><code>search_time_range</code> predicate values match against the time range of a search. For example, if you specify <code>search_time_range&gt;4h</code>, any search whose time range exceeds 4 hours is subject to the specified action.</p> <p>Supports <code>alltime</code> and numerical values with time units <code>m</code>, <code>minute</code>, <code>minutes</code>, <code>h</code>, <code>hour</code>, <code>hours</code>, <code>w</code>, <code>weeks</code>, <code>d</code>, <code>day</code>, or <code>days</code>. Supports comparison operators <code>=</code>, <code>&gt;</code>, and <code>&lt;</code>.</p> <p><code>search_time_range</code> predicate values match against time ranges selected using the Time Range Picker in the UI and time ranges defined by <code>earliest</code> and <code>latest</code> time modifiers for historical searches specified in a search string. For more information, see <a href="#">Specify time modifiers in your search</a>.</p> <p>Searches that run over the "all time" time range only match against rules that specify the <code>alltime</code> value, and do not match against rules that specify numerical values. For example, an all time search defined by <code>earliest=0</code> and <code>latest=now</code> only matches rules that specify <code>search_time_range=alltime</code>, and does not match rules that specify a numerical value, such as <code>search_time_range&gt;10m</code>.</p> <p><b>Note:</b> <code>search-time_range</code> predicate values that contain an equal sign (=) might not match against search time ranges that use "snap to" time units. For example, while <code>earliest=-24h</code> and <code>latest=now</code> defines a time range of 24 hours, when using "snap to" time units, such as <code>earliest=-24h@h</code> and <code>latest=now</code>, the time range can be evaluated as greater than or equal to 24 hours, depending on what time the search runs. See <a href="#">Relative time modifiers that snap to a time</a>. To ensure <code>search_time_range</code> values match against time ranges with "snap to" time units, avoid using the <code>=</code> operator, and instead specify a value such as <code>search_time_range&gt;24h</code>.</p> <p><code>srchTimeWin</code> and <code>srchTimeEarliest</code> settings in <code>authorize.conf</code>, which let you set maximum time range and earliest event time for searches, respectively, do not apply to workload rules and admission rules. For example, if you run a search with a time range of 7 days, and you have an admission rule that specifies <code>search_time_range&gt;48h</code>, the search will be filtered out, even if the <code>srchTimeWin</code> value for that role is set to 24 hours or greater.</p>
adhoc_search_percentage	<p>[0-100]. The percentage of the total search concurrency limit that you want to allocate to ad hoc searches.</p> <p>Specify an <code>adhoc_search_percentage</code> value to set a limit on the number of ad hoc searches that can run concurrently on the local search head. Use this to prevent skipped scheduled searches on systems running large numbers of ad hoc searches. See <a href="#">Example 4: Ad hoc</a></p>

Predicate type	Valid values
	<p><a href="#">search quota control</a>.</p> <p>When specifying <code>adhoc_search_percentage</code>, you must also specify <code>search_type=adhoc</code> in the predicate statement. For example, <code>search_type=adhoc AND adhoc_search_percentage&gt;10</code>.</p> <p>In a search head cluster environment, the <code>adhoc_search_percentage</code> value applies to each individual cluster member.</p> <p>The <code>adhoc_search_percentage</code> predicate does not work with other predicate types, with the exception of <code>search_type=adhoc</code>. Specifying any another predicate type returns an error.</p>

## Enable admission rules

The admission rules feature must be enabled for any existing admission rules to apply to searches. In Splunk Cloud Platform, the admission rules feature is enabled for the `sc_admin` role by default and cannot be disabled.

### Enable individual admission rules

You can enable or disable individual admission rules. This lets you create and save multiple different admission rules and deploy them as needed. Individual admission rules are enabled by default when you create them. Disabled admission rules are not evaluated and have no impact on search execution.

To enable or disable an individual admission rule:

1. In Splunk Web, click **Settings > Workload Management > Admission Rules**.
2. In the Status column, toggle the switch to enable or disable the individual admission rule.

## Example admission rules

The following examples show how you can define admission rules to filter out searches based on your use case objectives.

### Example 1: Stop wildcard searches

The following rule excludes data model acceleration searches that use `index=*` from the filter:

```
index=* AND (NOT search_type=datamodel_acceleration)
```

### Example 2: Stop alltime searches

Some monitoring searches use the `alltime` time range. The following rule excludes those searches from the filter:

```
search_time_range=alltime AND (NOT role=sc_admin) AND (NOT app=splunk_instance_monitoring)
```

The "View index inheritance for roles" feature also uses an `alltime` search. You can exclude users or roles that need to view indexes from the filter. The above rule excludes the `sc_admin` role.

REST API searches default to the `alltime` time range, unless you explicitly define a different time range in the search query.

### **Example 3: Do not allow certain users to run ad hoc searches at peak hours**

```
search_type=adhoc AND role=new_users
```

After you define the admission rule predicate, set the schedule for the rule, specifying the time range that corresponds to your peak hours.

### **Example 4: Ad hoc search quota control**

You can define admission rules to limit the number of concurrent ad hoc searches running on your system.

The following rule limits the number of concurrent ad hoc searches to 50% of the total search concurrency limit:

```
search_type=adhoc AND adhoc_search_percentage>50
```

After you specify the predicate statement, set the action for the rule to `Filter search` or `Queue search`.

If set to `Filter search`, the rule filters out any ad hoc searches over the 50% concurrency limit. If set to `Queue search`, the rule places any ad hoc searches over the 50% concurrency limit in a queue to be run again when the number of concurrent ad hoc searches falls below the 50% threshold.

For more information on search concurrency limits in Splunk Cloud Platform, see [Set limits for concurrent scheduled searches](#) in the *Splunk Cloud Platform Admin Manual*.

For more admission rule use case examples, see [Scenario 3: Create admission rules to prefilter searches](#).

## **Manually assign searches to workload pools**

Using workload rules to assign searches to workload pools automatically is the recommended method for allocating resources. You can however also manually assign searches to workload pools.

This page shows you how to assign searches to workload pools manually. For detailed instructions on how to assign searches to workload pools automatically using workload rules, see [Configure workload rules](#).

To assign searches to workload pools manually, you must have `list_workload_pools` and `select_workload_pools` capabilities.

### **Assign a scheduled search to a workload pool manually**

You can assign a scheduled search to a workload pool using Splunk Web, as follows:

1. Click **Settings > Searches, Reports, and Alerts**.
2. Find the specific saved search, and click **Edit > Advanced Edit**.
3. In the **Workload Pool** field, enter the name of the pool.
4. Click **Save**.

## Assign an ad hoc search to a workload pool manually

You can assign an ad hoc search to a workload pool using Splunk Web, as follows:

1. In the Search bar, enter your ad hoc search string.
2. Select a workload pool from the menu.
3. Run the search.

The ad hoc search job runs in the specified workload pool.

The screenshot shows the Splunk Web interface for a search job. At the top, the navigation bar includes 'splunk>cloud', 'App: Search & Reporting', '24 Messages', 'Settings', and 'Activity'. Below this is a secondary navigation bar with 'Search', 'Analytics', 'Datasets', 'Reports', 'Alerts', and 'Dashboards'. The main content area is titled 'New Search' and contains a search bar with the query 'index="\_internal" error'. Below the search bar, it indicates '159,553 events (10/10/19 10:00:00.000 PM to 10/11/19 10:26:56.000 PM)' and 'No Event Sampling'. A 'Job' dropdown menu is visible on the right. Below the search bar, there are tabs for 'Events (159,553)', 'Patterns', 'Statistics', and 'Visualization'. The 'Events' tab is active, showing a timeline view with green bars representing event counts over time. Below the timeline, there are controls for 'Format Timeline', 'Zoom Out', 'Zoom to Selection', and 'Deselect'. At the bottom, there is a table with columns for 'Time' and 'Event'. The table shows two event entries for 10/11/19 at 10:26:50.494 PM and 10:26:50.368 PM. On the left side of the table, there are field selection options: 'Hide Fields', 'All Fields', and 'SELECTED FIELDS' including 'host 7', 'search 100+', 'source 23', and 'sourcetype 18'. There is also an 'INTERESTING FIELDS' section at the bottom left.

4. Click **Job > Inspect Job > Search job properties.**
5. Confirm that the ad hoc search ran in the specified pool. For example:

 Search job inspector | Splunk 8.0.0 - Google Chrome

csms-3jvl10-8892.stg.splunkcloud.com/en-US/manager/search/job\_inspector?sid=1570832816.2893

searchCanBeEventType	true
searchEarliestTime	1570744800
searchLatestTime	1570832816
searchProviders	[ [-] idx-i-04f401827b8b0a766.csms-3jvl10-8892.stg.splunkcloud.com idx-i-097b4637b6ae257e0.csms-3jvl10-8892.stg.splunkcloud.com idx-i-0d4c1a709877ffda1.csms-3jvl10-8892.stg.splunkcloud.com sh-i-0c67b60d997a297f5.csms-3jvl10-8892.stg.splunkcloud.com ]
searchTotalBucketsCount	379
searchTotalEliminatedBucketsCount	16
sid	1570832816.2893
statusBuckets	300
ttl	600
workload_pool	standard_perf
Additional info	<a href="#">timeline</a> <a href="#">field summary</a> <a href="#">search.log_ ( idx-i-04f401827b8b0a766.csms-3jvl10-8892.stg.splunkcloud.com . 8892.stg.splunkcloud.com idx-i-0d4c1a709877ffda1.csms-3jvl10-8892.stg.splunk</a>

Server info: Splunk 8.0.0, csms-3jvl10-8892.stg.splunkcloud.com, Fri Oct 11 15:28:43 2019 User: admin

## Assign accelerated reports to workload pools manually

You can assign any report that qualifies for acceleration to a workload pool manually.

Assigning an accelerated report to a workload pool with ample CPU and memory resources can help you minimize performance issues that can occur during report acceleration, which can be resource-intensive.

To assign an accelerated report to a workload pool using Splunk Web:

1. Click **Settings > Searches, Reports, and Alerts**.
2. Find the report you want to accelerate and click **Edit > Edit Acceleration**.
3. Select the **Accelerate Report** checkbox.
4. Select the **Summary Range** for the report acceleration.
5. Select a workload pool from the menu.
6. Click **Save**.

For more information on report acceleration, see Accelerate reports in the *Splunk Enterprise Reporting Manual*.

## Assign accelerated data models to workload pools manually

You can assign an accelerated data model to a workload pool using Splunk Web, as follows:

1. Click **Settings > Data models**.
2. Find the data model you want to accelerate and click **Edit > Edit Acceleration**.
3. Select the **Accelerate** checkbox.
4. Select the **Summary Range** for the data model acceleration.
5. Select a workload pool from the menu.
6. Click **Save**.

For more information on accelerated data models, see Accelerate data models in the *Splunk Enterprise Knowledge Manager Manual*.

## Workload Management examples

The following scenarios provide some guidance on how to use workload management in Splunk Cloud Platform. These are hypothetical examples only. The exact steps will depend on your specific objectives and requirements.

### Scenario 1: Prioritize Security team searches

#### Use cases:

- Provide a high priority resource pool for all searches run by the security team.
- Put all index=\* and all time range searches in low priority pool.
- Abort all real-time searches after 1m.
- Move all long-running searches (>5m) that are not from the security team or sc\_admin into a low priority pool.
- Abort all long-running searches (>10m) that are not from the security team or sc\_admin.

#### Steps:

1. From Splunk Web, go to **Settings > Workload Management**.

2. Create the following workload rules by clicking **Add Workload Rule**.

The order of the workload rules is important. Workload rules are evaluated in order from top to bottom. If a search triggers a rule, corresponding action is taken and none of the rules below are evaluated. For example, if Rule #2 were ordered above Rule #1 in the table below, Rule #2 will be triggered after 5 minutes and the search will be moved to alternate pool. On next evaluation, again Rule #2 will be triggered. Rule #1 will never trigger and the search will not be aborted even after 10 minutes.

Order	Condition	Action
1	NOT (role=security OR role=sc_admin) AND runtime>10m	Abort
2	NOT (role=security OR role=sc_admin) AND runtime>5m	Move search to alternate pool: LowPriority
3	search_mode=realtime AND runtime>1m	Abort
4	index=* OR search_time_range=alltime	Place search in pool: LowPriority
5	role=security	Place search in pool: HighPriority

The rules are created and placed in a certain order to achieve the use cases. The rules are evaluated every few seconds and when a new search is started. If a rule is matched, the corresponding action is taken, and rules below that are not evaluated.

## Scenario 2: Create a high priority pool for scheduled searches

### Use Cases:

- Provide high priority pool for all scheduled searches from users in role=privileged but move these searches to the standard pool if they run for more than 2m.
- Move all adhoc searches running for more than 5m to low priority pool.
- Put all index=\* and all time range searches in low priority pool.
- Abort all searches running for more than 15m except searches from the sc\_admin.

### Steps:

1. From Splunk Web, go to **Settings > Workload Management**.
2. Create the following workload rules by clicking **Add Workload Rule**.

Order	Condition	Action
1	NOT (role=sc_admin) AND runtime>15m	Abort



Order	Condition	Action
2	search_type=adhoc AND runtime>5m	Move search to alternate pool: LowPriority
3	role=privileged AND search_type=scheduled AND runtime>2m	Move search to alternate pool: Standard
4	index=* OR search_time_range=alltime	Place search in pool: LowPriority
5	role=privileged AND search_type=scheduled	Place search in pool: HighPriority

### Scenario 3: Create admission rules to prefilter searches

#### Use cases:

- Filter out a rogue search acting on all indexes or in the `alltime` time range.
- Filter out a rogue search acting on all indexes and in the `alltime` time range and not from the Enterprise Security app.
- Filter out an ad hoc search from a role (e.g. `role=non_essential`) during peak business days.
- Filter out any search acting on the `security_events` index whose time range exceeds 24 hours, except for `role=security_users`.

#### Steps:

1. In Splunk Web, click **Settings > Workload Management**.
2. Click the **Admission Rule** tab.
3. Create the following admission rules by clicking **Add Admission Rule**.

Condition	Action	Schedule
index=* OR search_time_range=alltime	Filter search	always_on
index=* AND search_time_range=alltime AND NOT app=SplunkEnterpriseSecuritySuite	Filter search	always_on
search_type=adhoc AND role=non_essential	Filter search	Every Week On Monday, Tuesday, Wednesday, Thursday, Friday
index=security_events AND (NOT role=security_users) AND search_time_range>24h	Filter search	always_on

For more examples of admission rules, see [Example admission rules](#).

# Enable Automatic UI Updates

## Enable automatic UI updates

Enable automatic UI updates for Splunk Cloud Platform to receive the most recent UI features and bug fixes for supported Splunk Web UI pages. Enable automatic UI updates from Splunk Observability Cloud to see previews of observability data from Splunk Observability Cloud while using the Search app.

You can disable automatic UI updates for Splunk Cloud Platform and return to the default Splunk Web UI in the current version of Splunk Cloud Platform at any time. You can disable automatic UI updates from Splunk Observability Cloud to turn off the Related Content feature for your users at any time. To learn about related observability data, see Preview observability data in the Splunk Cloud Platform *Search* manual

### **Requirements**

To enable automatic UI updates a user's role must hold the `edit_auto_ui_updates`, `rest_properties_get`, and `rest_properties_set` capabilities.

### **Supported UI pages**

Splunk Cloud Platform currently supports automatic UI updates for the following Splunk Web UI pages:

- Splunk Home
- Uploaded apps
- IP allow list
- Webhook allow list
- Configure limits
- Dashboards Trusted Domains List
- Triggered Alerts
- Datasets

### **Enable automatic UI updates in Splunk Web**

After you enable Automatic UI updates on a search head, you receive automatic updates on this search head. To receive automatic updates on all search heads, enable the feature on each. The updates are not replicated across search heads.

To enable automatic UI updates, take these steps:

1. In Splunk Web, select **Settings** then select **Automatic UI updates**.
2. In the dialog box, toggle the switch next to Splunk Cloud Platform to get the latest updates for Splunk Cloud Platform. Toggle the switch next to Splunk Observability Cloud to use the Related Content feature, allowing your users to see previews of observability data in the Search app. See Preview observability data in the Splunk Cloud Platform *Search* manual for more information.

If you toggled the Splunk Cloud Platform switch, your Splunk Cloud Platform deployment will now display the most recent UI updates for the supported Splunk Web UI pages. If you toggled the Splunk Observability Cloud switch, your users can now see previews of related observability data when they perform searches in the Search app.

Enabling automatic UI updates does not require a restart.