Splunk Cloud™
Splunk Cloud Service Description 8.0.2004

Splunk Cloud Service Details

Generated: 6/29/2020 1:35 am

Splunk Cloud Service Details

Splunk Cloud delivers the benefits of award-winning Splunk® Enterprise as a cloud-based service. Using Splunk Cloud, 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 service uniformly, so all customers of Splunk Cloud receive the most current features and functionality.

Ingest-based subscription pricing for Splunk Cloud is based on the volume of uncompressed data that you want to index on a daily basis. The subscription pricing also includes access to Splunk support and a fixed amount of data storage. You can optionally add subscriptions for additional storage capacity to store more data, encryption service to maintain privacy of data at rest, HIPAA or PCI cloud environment to assist you with meeting your compliance needs, and add new use cases for Splunk Cloud with Splunk premium solutions such Enterprise Security and IT Service Intelligence. Optionally, Splunk offers infrastructure-based subscription as an alternate pricing option. For more information, see Infrastructure pricing.

Splunk Cloud is available in the following global regions:

AWS Data Centers:

- US (Oregon, Virginia, GovCloud)
- EU (Dublin, Frankfurt, London)
- Asia Pacific (Singapore, Sydney, Tokyo, Seoul)
- Canada (Central)

GCP Data Centers:

- US (Iowa)
- EU (London)
- Asia Pacific (Singapore)

For information regarding the availability of specific features and service components, see Differences between Splunk Cloud regions.

Ensure Operational Contacts listed in your Splunk.com support portal are regularly updated. Operational Contacts are notified when your Splunk Cloud environment undergoes maintenance, requires configuration awareness, or experiences a performance-impacting event.

For commonly asked questions about Managed Splunk Cloud, see the FAQ for Splunk Cloud.

For more information about self-service Splunk Cloud, including free trials, see the Self-service Splunk Cloud FAQ.

For more information about the terms of service, see the Splunk General Terms.

Data collection

Splunk Cloud provides software and APIs that enable you to ingest data from your applications, cloud services, servers, network devices, and sensors into the service. You can send data to Splunk Cloud as follows:

Using Splunk forwarders: 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 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. Your Splunk Cloud subscription includes a deployment server license for centralized configuration management of your Splunk forwarders. You can request the deployment server license from Splunk support. Setup, enablement, transformation, and sending data from forwarders to your Splunk Cloud environment is your responsibility. This means you are responsible for installing, configuring, and managing your forwarders, including maintaining version compatibility (see Supported Forwarder Versions for details). You are responsible for installing the data collection components of any app you wish to use in Splunk Cloud on a Splunk forwarder.

Splunk Cloud supports scripted and modular inputs either via the Inputs Data Manager (IDM) that is included in your Splunk Cloud subscription or via heavy forwarders that you manage and maintain. When you require an app installed on the IDM, you to open a support ticket and Splunk support will install the app on your behalf.

For more information about Inputs Data Manager, see Features of Splunk Cloud in the *Splunk Cloud User* manual.

For more information, see Upload Data in the *Getting Data In* manual.

Using HTTP Event Collector (HEC): HEC lets you send data and application events using a token-based authentication mode to Splunk Cloud 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. HEC is enabled by default for your Splunk Cloud environment with a 1 MB size limit on the maximum content length. You are responsible for setup, enablement, transformation, and sending data to your Splunk Cloud environment via HEC. You are also responsible for monitoring and remediation of any HEC error codes that are received from Splunk Cloud to ensure no interruption of your data ingestion. For more information, see Use the HTTP Event Collector in the *Getting Data In* manual.

Using AWS Kinesis Data Firehose: AWS Kinesis Data Firehose is a fully managed, scalable, and serverless option for streaming data from various AWS services directly into Splunk Cloud. Setup, enablement, transformation, and sending data to your Splunk Cloud environment is your responsibility. If you choose to use the Kinesis Data Firehose service for data ingestion, you are responsible for enabling and configuring AWS Kinesis Data Firehose, and for paying AWS for this service. For more information, see Install and configure the Splunk Add-on for Amazon Kinesis Firehose on a managed Splunk Cloud deployment in the *Splunk Add-on for Amazon Kinesis Firehose* manual.

Additional information about data collection

Data compression: Forwarders and HTTP Event Collectors compress data when sending over TLS protocol. The amount of compression varies based on the content, generally at a ratio between 1:8 and 1:12.

Encryption in transit: For security, data in transit is TLS 1.2+ encrypted. Senders and receivers authorize each other, and HTTP-based data collection is secured using token-based authentication.

IP Whitelisting: You can request to restrict data collection from only whitelisted IP addresses by filing a support ticket.

Ingestion

The amount of data that your Splunk Cloud environment can collect daily is determined by the ingest-based subscription that you purchase, and you can always choose a higher-level ingest-based subscription to increase the amount of data that you can collect. You can see current and past daily data ingestion information using the Cloud Monitoring Console (CMC) app that is included with your Splunk Cloud environment. If you consistently exceed your subscription entitlement, contact Splunk Sales to purchase an appropriate ingest-based subscription plan to handle your volume. If you have an

infrastructure-based subscription, your subscription does not meter ingest.

During ingestion, Splunk Cloud 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. Splunk Cloud allows you to self-service manage your indexes across multiple tasks such as the following:

- Creating, updating, deleting, and viewing properties of indexes
- Modifying the retention settings for individual indexes
- Deleting data from indexes
- Optimizing search performance by managing the number of indexes and the data sources that are stored in specific indexes

For details about limits on data collection, see Splunk Cloud data policies in the *Splunk Cloud User Manual*.

For best practices for creating indexes, see Manage Splunk Cloud indexes in the *Splunk Cloud User Manual*.

For service limits relating to indexes, see Splunk Cloud service limits and constraints.

Storage

Storage space in your Splunk Cloud service is based on the volume of uncompressed data that you want to index on a daily basis. Your ingest-based Splunk Cloud subscription comes with sufficient storage to allow you to store up to 90 days of your uncompressed data. For example, if your daily volume of uncompressed data is 100 GB, your Splunk Cloud environment will have 9000 GB (9 TB) of storage. You can optionally purchase additional storage for your Splunk Cloud environment in 500 GB increments. In addition, you can choose to have your data encrypted at rest using AES 256-bit encryption for an additional charge. If you choose encryption at rest, Splunk manages the encryption keys on your behalf by default and you can request to manage the keys instead.

When you send data to Splunk Cloud, it is stored in indexes and you can self-manage your Splunk Cloud indexes settings using the Indexes page in Splunk Web. Splunk Cloud 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.

When the index reaches the specified maximum size or events reach the specified maximum age, the oldest data is deleted. If you require a lower cost option for storage of data beyond 90 days, you can optionally augment Splunk Cloud with Dynamic Data Active Archive (DDAA). As data ages from searchable storage based on your index retention setting, the aged data is automatically moved to DDAA before deletion. Data remains in DDAA until the DDAA retention setting that you specify expires. Your DDAA subscription enables you to perform restores, subject to the amount of storage you have purchased as part of your Splunk Cloud subscription. An additional 10% of searchable storage is included with your DDAA subscription to assist with restores. Note that multiple restores that overlap within a 30 day period will accrue against the additional 10% of searchable storage included with your DDAA subscription.

If you enable Dynamic Data Self-Storage (DDSS) to export your aged ingested data, the oldest data is moved to your Amazon S3 account in the same region as your Splunk Cloud before it is deleted from the index. You are responsible for AWS payments for your use of Amazon S3. When data is deleted from the index, it is no longer searchable by Splunk Cloud.

For more information about export of your aged ingested data, see Store expired Splunk Cloud data.

For more information about archiving your aged ingested data, see Archive expired Splunk Cloud data.

You can review your storage consumption in the Cloud Monitoring Console app included in your Splunk Cloud environment. The app provides information such as the amount of data stored and the number of days of retention for each index.

For more information about managing indexes, see Manage Splunk Cloud indexes in the *Splunk Cloud User Manual*.

For more information about the Cloud Monitoring Console, see Monitor your Splunk Cloud Deployment in the *Splunk Cloud Admin Manual*.

DDAA and DDSS plus the option for you to manage the encryption keys are not available for Splunk Cloud on GCP. For information regarding the availability of specific features and service components, see Differences between Splunk Cloud regions.

Search

Splunk Cloud 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 the form of visualizations, reports, and alerts.

If you enable Dynamic Data Self-Storage to export of your aged ingested data prior to deletion, any data moved from these indexes to your AWS S3 account will no longer be searchable by Splunk Cloud. If you augment Splunk Cloud with Dynamic Data Active Archive (DDAA), restored DDAA data is searchable within 24 hours of it being restored and is searchable for up to 30 days.

To examine data in Splunk Cloud and your on-premises deployment of Splunk Enterprise in a single search, you can configure a Splunk Enterprise search head to connect to a Splunk Cloud indexer cluster. This configuration is called hybrid search.

The following conditions and limitations apply to hybrid search:

Category	Supported	Limitation
Hybrid Search Topology	You can initiate searches from an on-premises Splunk Enterprise search head to a single Splunk Cloud deployment.	You cannot initiate searches from an on-premises Splunk Enterprise search head to multiple Splunk Cloud environments. You cannot initiate searches from a Splunk Cloud search head to an on-premises Splunk Enterprise environment. You cannot initiate searches from a Splunk Cloud search head to another Splunk Cloud environment.
Splunk Version Compatibility	The on-premises search head must run the same version as Splunk Cloud, down to the maintenance level. For example, if your Splunk Cloud is at version x.y.z, then the on-premise	

Category	Supported	Limitation
	Splunk search head must also be at version x.y.z.	
Search Concurrency	Your Splunk Cloud search concurrency limits apply to searches initiated either from the Cloud search tier or from on-premise hybrid search heads. For more information, refer to Splunk Cloud service limits and constraints.	
Search Types	Ad-hoc search is supported.	Scheduled search is not supported from a hybrid search head.
Premium Solution		Hybrid search is not available for use with all Splunk Premium Solutions: Enterprise Security, IT Service Intelligence, Splunk App for Microsoft Exchange, Splunk App for PCI Compliance, Splunk App for VMware.

For more information about hybrid search, see Configure hybrid search in the *Splunk Cloud User Manual*.

In Splunk Cloud, you open a support ticket to enable real-time search. Note that real-time searches are resource-intensive and can impact the overall health and performance of your searches.

You can review the health and performance of your search using the Cloud Monitoring Console (CMC) app that is included in your Splunk Cloud environment. CMC shows information such as long running searches, skipped scheduled searches, and average search run time.

Splunk Cloud has service limits related to search, such as the maximum number of concurrent searches. This service limit and others are listed in the Splunk Cloud service limits and constraints section.

Splunkbase and private apps

Apps and Add-Ons (apps) include features and functionality ranging from the simplification of data ingest to unique and valuable visualizations. To ensure security and minimize effects on performance, only vetted and compatible apps can run on Splunk Cloud. Note the following:

- Splunkbase is the system of record for app vetting and compatibility with Splunk Cloud.
- Splunk provides support and maintenance for *Splunk Supported Apps*. In addition, Splunk Cloud ensures compatibility for any installed *Splunk Supported Apps* before commencing Splunk Cloud upgrades.
- Splunk does not provide support or maintenance for apps published by any third-party developers. For any *Developer Supported or Not Supported Apps*, you need to ensure compatibility with Splunk Cloud.
- Compatibility of *Developer Supported or Not Supported Apps* is asserted by the developers of those apps. Splunk does not perform compatibility testing of third-party apps with specific versions of Splunk Cloud.
- Splunk support will not be able to assist in tailoring the Splunkbase apps to your use case. For apps that grant you the license to customize, you will need to perform the customization yourself or through a Splunk Professional Services engagement.

For more information, refer to the Splunkbase app support types here.

Apps that are Splunk Cloud vetted and compatible are listed in either the app browser in Splunk Web or through Splunkbase. Depending on the nature of the Splunkbase apps, you may be able to self-install because they have been marked so, or you may need to open a support ticket to install. When you require an app installed on the Inputs Data Manager, you to open a support ticket and Splunk Support will install the app on your behalf.

Apps you create to support your business needs are called private apps and these apps can also be self-service installed on Splunk Cloud. During the private app installation, Splunk will automatically validate if your app is compatible with

Splunk Cloud and allow the installation to complete. If the app is deemed incompatible with Splunk Cloud, you will receive an app vetting report that details areas in your app to remediate to make it compatible with Splunk Cloud and enable it to be self-service installable. Private apps that are developed wholly by you are owned by you and any customization of your private app is outside the scope of the Splunk Cloud subscription.

For more information about Apps, see the following topics in the *Splunk Cloud User Manual*:

- Install apps in your Splunk Cloud deployment
- Manage private apps in your Splunk Cloud deployment
- Manage a rolling restart in Splunk Cloud

Splunk premium solutions

You can optionally purchase Splunk apps and premium solutions (premium solutions) subscriptions on Splunk Cloud. As part of the subscription, the Splunk Cloud environment is enhanced to support the premium solution. Splunk will install the premium solution on your behalf and will also upgrade the premium solution when a new version is vetted for Splunk Cloud. Multiple premium solution subscriptions can run concurrently on the same Splunk Cloud environment. Any customization of the premium solution can be done by you or through a Splunk Professional Services engagement. Splunk support will not be able to assist in tailoring the premium solution to your use case. The following premium solution subscriptions are available for Splunk Cloud:

- Splunk Enterprise Security (ES)
- Splunk IT Service Intelligence (ITSI)
- Splunk App for Microsoft Exchange
- Splunk App for PCI Compliance
- Splunk App for VMware

Splunk App for VMware is not available for Splunk Cloud on GCP. For information regarding the availability of specific features and service components, see Differences between Splunk Cloud regions.

Machine Learning Tool Kit (MLTK) is compatible with Splunk Cloud and supports a variety of use cases. Depending on the use case and algorithm used, the MLTK app can be compute intensive. Splunk recommends that you consult with your Splunk technical resource and MLTK documentation prior to installing the MLTK App on Splunk Cloud. In addition, Splunk recommends adding the ML-SPL Performance App for the Machine Learning Toolkit to ensure you know the resource utilization impact of MLTK. These steps ensure the MLTK best practices are implemented on Splunk Cloud.

The following premium solutions are compatible with Splunk Cloud but no subscription is available on Splunk Cloud. Installation and configuration of these premium solutions can be done by you or through a Splunk Professional Services engagement. Splunk support will not be able to assist with installation and configuration of the following premium solutions. For more information on these Splunk premium solutions, contact your Splunk sales representative.

- Phantom
- User Behavior Analytics
- VictorOps

Network connectivity and data transfer

You access your Splunk Cloud environment via public endpoints. By default, for both Splunk Web access and sending your data, traffic from your network is encrypted, sent over the public Internet and then routed to your Splunk Cloud environment in a Virtual Private Cloud (VPC). These endpoints are protected using firewall rules and customers can also

specify additional access control rules. See the Splunk Cloud service limits and constraints section for the maximum number of customer-defined rules.

You can request to restrict data access from only whitelisted IP addresses by filing a support ticket. For any regulated Splunk Cloud environments such as HIPAA and PCI, you must specify at least one address for the IP whitelist.

In addition, forwarders and HTTP Event Collectors compress data when sending over TLS protocol. The amount of compression varies based on the content. For bandwidth planning, assume a compression ratio between 1:8 and 1:12.

If you are using AWS services such as Direct Connect and Kinesis Data Firehose, note the following:

- If you choose to use connectivity services such as AWS Direct Connect to access Splunk Cloud to reduce your overall network costs and increase bandwidth throughput, you are responsible for setup and configuration of AWS Direct Connect plus any associated fees.
- If you choose to use the Kinesis Data Firehose service for data ingestion, you are responsible for any setup and configuration of AWS Kinesis Data Firehose plus any AWS fees. For more information see Install and configure the Splunk Add-on for Amazon Kinesis Firehose on a managed Splunk Cloud deployment in the *Splunk Add-on for Amazon Kinesis Firehose* manual.
- If you enable Dynamic Data Self-Storage to export your aged ingested data to your Amazon S3 account in the same region as your Splunk Cloud, you are responsible for any setup, configuration and AWS payments. For more information, refer to the *Splunk Cloud User Manual*.
- AWS Direct Connect or Kinesis Data Firehose may not be available in all Splunk Cloud regions.

Users and authentication

Splunk Cloud enables you to configure account policies that require unique usernames, minimum password length, and regular password resets. You are responsible for creating and administering your users' accounts, the roles assigned to them, the authentication method they use, and global password policies. To control what your Splunk Cloud users can do, you assign them roles that have a defined set of specific capabilities, access to indexes, and resource use limits.

Roles give Splunk Cloud users access to features in the service, and permission to perform tasks and searches. Each user account is assigned one or more roles. In addition, your Splunk Cloud environment comes with predefined system roles and system users that are used by Splunk to perform essential monitoring and maintenance activities. You should not delete or modify these system users or roles. For the customer's administrator users, Splunk Cloud provides the sc_admin role, which has the capabilities required to administer Splunk Cloud. You can use the Splunk Cloud sc_admin role for your administrator to perform self-service tasks such as installing apps, creating and managing indexes, and managing users and their passwords. Splunk Cloud does not support direct access to infrastructure, so you do not have command-line access to Splunk Cloud. This means that any supported task that requires command-line access is performed by Splunk on your behalf.

You can configure your user accounts to be authenticated using Identity Providers (IdP) such as Lightweight Directory Access Protocol (LDAP) and Active Directory (AD). You can also configure Splunk Cloud to use SAML authentication for single sign-on (SSO). In order to use multifactor authentication for your Splunk Cloud user accounts, you must use a SAML v2 identity provider that supports multifactor authentication. While Splunk Enterprise has built-in support for multifactor authentication such as Duo and RSA, Splunk Cloud does not support these methods of integration.

Only SHA-256 signatures in the SAML message between your IdP and Splunk Cloud are supported. You are responsible for the SAML configuration of your IdP including the use of SHA-256 signatures.

For more information on User and Roles, see Manage Splunk Cloud users and roles in the *Splunk Cloud Admin Manual*.

For more information on Single Sign On, see Configure SAML single sign-on (SSO) to Splunk Cloud in the *Splunk Cloud User Manual*.

Differences between Splunk Cloud and Splunk Enterprise

Splunk Cloud delivers the benefits of Splunk Enterprise as a cloud-based service. Customers who are familiar with Splunk Enterprise architecture should not make assumptions about the architecture or operational aspects of Splunk software deployed in the Splunk Cloud service. Specifically, Splunk Cloud differs from Splunk Enterprise in the following ways:

Area	Difference
Apps	To ensure security and minimize effects on performance, only vetted and compatible apps can run on Splunk Cloud. The app browser in Splunk Web or Splunkbase lists vetted and compatible Splunk Cloud apps. You can install some apps directly through the app browser (self-service installation). When an app cannot be self-installed, including for an IDM, you must open a support ticket and Splunk Support will install the app on your behalf. Your private apps can also be self-service installed. During the private app installation, Splunk automatically validates if your private app is compatible with Splunk Cloud and allows compatible apps to be self-service installed. If the app is deemed incompatible with Splunk Cloud, you receive an app vetting report that details areas in your app to remediate to make it compatible with Splunk Cloud and enable it to be self-service installable.
Cloud Monitoring Console (CMC)	The Cloud Monitoring Console (CMC) app is included in your Splunk Cloud environment. CMC replaces the Monitoring Console that is used in Splunk Enterprise. You use CMC to holistically monitor the data consumption and health of your Splunk Cloud environment.
Command-line interface (CLI) access	Splunk Cloud does not allow direct access to infrastructure by customers. As a result, you do not have command line access to Splunk Cloud. Any supported task that requires command line access is performed by the self-service capabilities of Splunk or by filing a service ticket.
Direct TCP, UDP, file, and syslog inputs	Splunk Cloud does not accept these types of data directly. In order for Splunk Cloud to receive data sources such as TCP, UDP, file, and syslog, you must use Splunk forwarder software as an agent to send data to Splunk Cloud. This ensures reliable, managed, fault-tolerant delivery of your data into Splunk Cloud.
Direct TCP, UDP, file, and syslog outputs	Splunk Cloud does not accept unencrypted outputs at the search head tier, and does not support outputs of any kind at the indexer tier, including custom search commands, such as cefout (bundled with Splunk App for CEF). This ensures reliable and fault-tolerant performance of your Splunk Cloud environment.
Dynamic Data Active Archive	Dynamic Data Active Archive (DDAA) is only available in Splunk Cloud and it is an optional subscription. DDAA offers a lower cost option for long term storage of your ingested data.
Export of your ingested data to Amazon S3 using Dynamic Data Self-Storage	Dynamic Data Self-Storage is only available in Splunk Cloud.
License pooling and exceeding purchased daily index volume	Splunk Cloud does not support licensing pooling. In addition, you can exceed your purchased daily index volume a maximum of five times in a calendar month. For more information, review the data ingestion and daily license usage policy here.
Multifactor authentication	While Splunk Enterprise has built-in support for multifactor authentication such as Duo and RSA, Splunk Cloud does not support these methods of authentication. In order to use multifactor authentication for your Splunk Cloud user accounts, you must configure a SAML v2 identity provider that supports multifactor authentication.
Native alerts	Splunk Cloud does not provide system-level access. As a result, you cannot define alerts that run operating-system scripts or use other system services (although vetted and compatible apps can do so). Alerts can be sent by email or HTTPS POST using Splunk software webhooks. You might be required to set up an endpoint inside your network. If you have both Splunk Enterprise and Splunk Cloud, you can run an

Area	Difference
	on-premises search head to support searches that require alert actions.
Real-time search	In Splunk Cloud, you open a support ticket to enable real-time search. Note that real-time searches are resource intensive and can impact the overall health and performance of your searches.
REST API	Differences in implementation details between Splunk Cloud and Enterprise plus permissions for the `sc_admin` role impact REST API access. In Splunk Cloud, you open a support ticket to enable REST API access. In addition, Splunk Cloud supports a subset of the REST API endpoints available in Splunk Enterprise. You can find more information regarding using the REST API with Splunk Cloud here.
Scripted and Modular Inputs	Splunk Cloud supports scripted and modular inputs either via the Inputs Data Manager (IDM) that is included in your Splunk Cloud subscription or via heavy forwarders that you manage and maintain. Either choice maintains your SLA since no data ingestion is performed on the search tier. For more information about Inputs Data Manager, see Features of the Splunk Cloud Platform in the *Splunk Cloud User* manual.
Search performance	Splunk Cloud leverages a multi-tier storage architecture and manages the movement of data to optimize performance based on user search patterns. Generally, recently processed data (recently ingested, searched, analyzed for machine learning, and so on) will have better performance than data that has not been processed for some time. This behavior applies to all data, including metrics data.
sc_admin role	For the customer's administrator users, Splunk Cloud provides the `sc_admin` role, which has sufficient capabilities to administer Splunk Cloud. You can use the Splunk Cloud `sc_admin` role for your administrator to perform self-service tasks such as installing apps, creating and managing indexes, and managing users and their passwords.
System user roles	Your Splunk Cloud environment comes with predefined system roles and system users that are used by Splunk to perform essential monitoring and maintenance activities. You should not delete or modify these system users or roles.
Token-based authentication	Splunk Cloud does not support authentication with tokens.
Workload Management	Splunk Cloud provides pre-configured workload pools for your use. For details, see Workload Management in the *Splunk Cloud User Guide*.

Service level

Splunk provides an uptime SLA for Splunk Cloud and will use commercially reasonable efforts to make the Services available. You will receive service credits in the event of SLA failures, as set forth in our current SLA schedule. As Splunk Cloud is offered uniformly across all customers, the SLA cannot be modified on a customer by customer basis.

Splunk Cloud is considered available if you are able to log into your Splunk Cloud Service account and initiate a search using Splunk Software. Splunk continuously monitors the status of each Splunk Cloud environment to ensure the SLA. In addition, Splunk Cloud monitors several additional health and performance variables, including but not limited to the following:

- Ability to log into Splunk Cloud (non-SAML)
- Ability to access Splunk Web
- Ability to access a Splunk REST API endpoint
- Ability to perform searches against an internal Splunk index
- Ability to ingest data cluster wide
- Presence of unsupported configurations

Splunk adds predefined system users and system roles to all Splunk Cloud environments. Splunk leverages system users or roles to perform essential monitoring and maintenance activities in managed Splunk Cloud environments. Customers are advised to not delete or edit system users or roles because they are essential to perform monitoring and maintenance activities in managed Splunk Cloud environments.

Splunk Cloud supports scripted and modular inputs either via the Inputs Data Manager (IDM) that is included in your Splunk Cloud subscription or via heavy forwarders that you manage and maintain. Either choice maintains your SLA since no data ingestion is performed on the search tier.

For more information about Splunk Cloud system users, see Manage Splunk Cloud users and roles in the *Splunk Cloud Admin Manual*.

For more information about the SLA for Splunk Cloud, see the Splunk Cloud Service Level Schedule.

Maintenance

Splunk Cloud delivers the benefits of award-winning Splunk® Enterprise as a cloud-based service. Splunk manages and updates the Splunk Cloud service uniformly, so all customers of Splunk Cloud receive the most current features and functionality. Ensure Operational Contacts listed in your Splunk.com support portal are regularly updated. Operational Contacts are notified when your Splunk Cloud environment undergoes maintenance, requires configuration awareness, or experiences a performance-impacting event. These contacts will receive regular notifications of planned and unplanned downtime, including scheduled maintenance window alerts and email updates related to incident-triggered cases.

What Splunk does on your behalf:

- **Gets you started:** When you first subscribe to Splunk Cloud, Splunk sends you a welcome email containing the information required for you to access your Splunk Cloud deployment and get started. This email contains a lot of important details, so keep it handy.
- **Assists you with supported tasks:** Splunk Cloud enables you to customize user, index and app management via Splunk Web. However, there are features in Splunk Cloud that require assistance from Splunk to activate or make changes to your configurations, such as real-time search and enabling AWS Kinesis Data Firehose data to be received. When you file a support ticket, Splunk will enable such features on your behalf.
- **Upgrades and expands your Splunk Cloud:** Splunk Cloud adopts the release that has the most benefits for you as quickly as possible. To ensure efficiency and agility, you will be assigned to an upgrade cohort and as Splunk releases new versions of Splunk Cloud and Premium Apps, your cohort will be notified by Splunk of the upcoming maintenance window. In certain maintenance situations, data egress of Dynamic Data Self-Storage will be paused. In addition, we will enhance Splunk Cloud on your behalf, such as increasing the amount of your daily ingestion, adding storage, enabling Premium App subscriptions and Encryption at Rest.
- **Ensures Splunk Cloud uptime and security:** Splunk continuously monitors the status of your Splunk Cloud environment to ensure uptime and availability. We look at various health and performance variables such as the ability to log in, ingest data, access Splunk Web and perform searches. Splunk maintains a rolling 30 day history of health and utilization data to help ensure uptime and assist troubleshooting of your Splunk Cloud. In addition, Splunk Cloud maintains a rolling seven day daily backup of your ingested data and configuration files to ensure data durability. Splunk also employs system user roles with limited privileges to perform tasks on your cloud. If you purchase an encryption at rest subscription, we manage the encryption keys on your behalf.

What you can self-service:

- **Customize your Splunk Cloud:** Splunk Cloud offers multiple options to ingest your data, so it is your responsibility to ensure the correct data collection method is used for your data sources. For detailed instructions for sending data to your Splunk Cloud deployment, refer to the Getting Data In manual. In addition, Splunk provides a variety of self-service tools to allow you to customize your Splunk Cloud environment, such as user, index and app management. For more information, refer to the Splunk Cloud User Manual.
- **Monitor your Splunk Cloud health and usage:** You can use the Cloud Monitoring Console (CMC) to holistically monitor the data consumption and health of your Splunk Cloud environment. Your license limits the amount of

data per day that you can send to your Splunk Cloud deployment. CMC is designed to help you manage your usage of the service, while all other monitoring is done by Splunk.

Technical support

Ingest-based Splunk Cloud subscriptions include either Standard Success Plan or Premium Success Plan. Infrastructure-based subscriptions include Premium Success Plan only. For more information regarding Splunk Cloud support terms and program options, refer to: https://www.splunk.com/en_us/support-and-services/support-programs.html. You should also note the following:

- Splunk Cloud offers multiple options to ingest your data so it is your responsibility to ensure the correct data collection method is configured for your data sources.
- Splunk Cloud enables you to perform user, index and app management via Splunk Web. Any customization of Splunk Cloud vetted and compatible apps is also your responsibility.
- In order to use multifactor authentication for your Splunk Cloud user accounts, you must use a SAML v2 identity provider that supports multifactor authentication. It is your responsibility to ensure your Splunk Cloud user accounts are properly configured for multifactor authentication.
- You can choose to leverage the optional Admin on Demand Services to quickly request technical adoption assistance from remote Splunk technical consultant. The Splunk technical consultants can assist you with tasks, such as index creation, building lookups and dashboards, assist with data on-boarding plus install Splunk Cloud vetted and compatible apps.
- There are features in Splunk Cloud that require assistance from Splunk to activate or change your configuration, such as real-time search and enabling AWS Kinesis Data Firehose data to be received. When you file a support ticket, Splunk will enable such features on your behalf.

For more information regarding Admin on Demand Services, refer to the Admin On Demand data sheet and catalog.

For more information regarding data collection, refer to *Getting Data In*.

For more information regarding performing user, index and app management, refer to the *Splunk Cloud User Manual*.

Security

The security and privacy of your data is of the utmost importance to you and your organization, and Splunk makes this a top priority. Splunk Cloud service is designed and delivered using key security controls such as:

Instance Security: Every Splunk Cloud deployment runs in a secured environment on a stable operating system and in a network that is hardened to industry standards using a default-deny firewall policy, which permits access only to specific IP addresses and services. Your deployment is regularly scanned for host- and application-level threats.

Isolation of Data and Service: In the cloud, your data is logically isolated from other customers' data, so your performance and data integrity cannot be affected by other customers who are using the Splunk Cloud service.

Data Encryption: All data in transit to and from Splunk Cloud is TLS 1.2+ encrypted. To encrypt data at rest, you can purchase AES 256-bit encryption for an additional charge. Keys are rotated regularly and monitored continuously.

User Authentication and Access: You can configure authentication using Lightweight Directory Access Protocol (LDAP), Active Directory (AD), and single sign-on using any SAML v2 identity provider. To control what your Splunk Cloud users can do, you assign them roles that have a defined set of specific capabilities. Splunk Cloud enables you to configure account policies that require unique user names, minimum password length, and regular password resets with supported SAML v2 identify providers and LDAP. To enable multifactor authentication, customers must configure a SAML

v2 identity provider that supports multifactor authentication. Only SHA-256 signatures in the SAML message between your IdP and Splunk Cloud are supported.

Data Handling: You can store your data in one of the following regions:

AWS Data Centers:

- US (Oregon, Virginia, GovCloud)
- EU (Dublin, Frankfurt, London)
- Asia Pacific (Singapore, Sydney, Tokyo, Seoul)
- Canada (Central)

GCP Data Centers:

- US (Iowa)
- EU (London)
- Asia Pacific (Singapore)

Data is kept in the region you choose. If you need to store your data in more than one region, you can purchase multiple subscriptions. Data is retained in Splunk Cloud according to the volumes, durations, and index configurations you set. Expired data is deleted based on your pre-determined schedule.

For the purposes of disaster recovery, your configuration and recently-ingested data is backed up on a rolling seven-day window. If you unsubscribe, you can take your data with you to an alternative storage container, such as AWS S3 bucket, prior to deletion. Depending on the amount of data and the work involved, we may charge for this service. For more information on Splunk Cloud data management, please review the documentation at Splunk Cloud data policies and Manage Splunk Cloud indexes in the *Splunk Cloud User Manual*.

Security Controls and Background Screening: Splunk security controls are described in our most recent Service Organization Control II, Type II Report (SOC 2/Type 2 Report). For more information about regions for which Splunk does not have SOC2 controls in place, please see the Splunk Cloud Security Addendum. Splunk conducts criminal background checks on its employees prior to hire, as permitted by law.

App Security: All Splunk apps hosted on Splunk Cloud by Splunk are examined by Splunk engineers to ensure that they comply with the Splunk Cloud app requirements and best practices. The Splunk App Certification Program provides a set of best practices for app developers. For details about how to submit an app for evaluation for Splunk Cloud readiness, see the Splunk Developer web page.

For more information about **Splunk Data Privacy, Security and Compliance**, see Splunk Protects.

For information regarding the availability of specific features and service components, see Differences between Splunk Cloud regions.

Subscription expansions, renewals and terminations

You can expand aspects of your Splunk Cloud subscription anytime during the term of the subscription to meet your business needs. You can:

- increase the amount of your daily ingestion
- increase the amount of storage in 500GB increments
- add Premium App subscriptions

- add Encryption at Rest capabilities

You will receive renewal notifications starting 60 days prior to the end date of your current subscription term. For more information on subscription renewals, contact your Splunk sales representative. If your Splunk Cloud subscription expires, it is considered terminated. The policy for terminated Splunk Cloud subscriptions are the following:

- your ability to perform searches stops immediately
- your ability to ingest data stops 7 days following termination
- your data is deleted 31 days following termination

If you require your ingested data to be moved into your control before the termination of your subscription, this is accomplished through a Splunk Professional Services engagement or by enabling Dynamic Data Self-Storage to export your aged data to your Amazon S3 account in the same region. Note that Dynamic Data Self-Storage does not export your configuration data. If you choose to use Dynamic Data Self-Storage to export your aged ingested data, you must do so prior to termination of your subscription. You are responsible for AWS charges you incur for your use of Amazon S3.

Infrastructure pricing

For certain use cases, Splunk offers infrastructure-based subscription as an alternate pricing option. This pricing model is based on the resources allocated rather than data volume ingested. The main differences between infrastructure-based and ingest-based subscriptions are highlighted below. For more information regarding infrastructure-based pricing, contact your Splunk Sales representative.

Ingestion and Search

This pricing model is based on the resource capacity consumed rather than data volume ingested. As a result, infrastructure-based subscription does not meter ingest. In addition, Splunk Cloud deploys the purchased resources. You can increase ingest and/or search load and operate the service to your desired performance objective. As necessary, you can purchase additional resource capacity to increase ingest and search load or to improve performance.

Storage

You purchase units of storage blocks based on your data retention requirements for your infrastructure-based subscription. Unlike ingest-based subscription, storage is not bundled and provides you flexibility to tailor to the variability in your use case. If you ingested far more data than your initial estimate and thus exceeded your purchased storage capacity, the Splunk Cloud service elastically expands the amount of storage to retain your data per your retention settings. At the end of the annual subscription term, Splunk will true-up your storage for any overages. The Cloud Monitoring Console dashboard provides you with the total amount of data retained at any given time.

Success Plans

Infrastructure-based subscriptions include Premium Success Plan. For more information, refer to the Splunk Success Plan.

Splunk Cloud performance considerations

Every Splunk Cloud ingest-based subscription plan is provisioned with adequate compute capacity. Because search workloads can vary considerably, subscription plans with peak daily ingest of 1000 GB and greater are guaranteed allocation of Splunk Virtual Cores as defined below.

A Splunk Virtual Core (SVC) is a unit of capabilities in Splunk Cloud that includes compute, memory, and I/O resources. SVCs are allocated to your subscription plan based on your average daily ingest-based subscription, up to the maximum of 1 SVC for every 10 GB of licensed peak daily ingest. Purchase of Splunk Enterprise Security (ES) Premium Solution provides incremental SVC allocation of 1 SVC for every 20 GB of licensed peak daily ingest. Purchase of Splunk IT Service Intelligence (ITSI) Premium Solution provides incremental SVC allocation of 1 SVC for every 20 GB of licensed peak daily ingest.

The ratio of allocated SVC to licensed peak daily ingest level is subject to change with the evolving infrastructure and architecture of the service. Splunk Cloud establishes SVC performance using a Splunk Search Benchmark to ensure that new ratios continue to provide the same or better levels of performance.

Compliance and certifications

Splunk has attained a number of compliance attestations and certifications from industry-leading auditors as part of our commitment to adhere to industry standards worldwide and part of our efforts to safeguard customer data. The following compliance attestations/certifications are available:

- **SOC 2 Type II**: Splunk Cloud has an annual SOC 2 Type 2 audit report issued. The SOC 2 audit assesses an organization's security, availability, process integrity, and confidentiality processes to provide assurance about the systems that a company uses to protect customers' data. If you require the SOC 2 Type 2 attestation to review, contact your Splunk sales representative to request it.
- **ISO 27001**: Splunk Cloud is ISO/IEC 27001:2013-certified. ISO/IEC 27001:2013 is a standard for an information security management system, specifying the policies and procedures for all legal, physical, and technical controls used by an organization to minimize risk to information. Splunk's auditors ISO certification can be found here.

SOC 2 Type II compliance attestation and certification is not available for Splunk Cloud on GCP. For information regarding the availability of specific features and service components, see Differences between Splunk Cloud regions.

If your data must be maintained in a regulated cloud environment to assist you with meeting your compliance needs, Splunk Cloud provides these optional subscriptions.

- **Health Insurance Portability and Accountability Act (HIPAA)**: Splunk Cloud (HIPAA) is compliant with the HIPAA Security Rule and HITECH Breach Notification Requirements. These regulations establish a standard for the security of any entity that access, processes, transmits, or stores electronic protected health information (ePHI).

- **Payment Card Industry Data Security Standard (PCI DSS)**: Splunk tests Splunk Cloud for compliance with the PCI DSS v3.2 standard. This standard applies to any entity that processes, transmits, or stores payment card data as well as their critical service providers.

HIPAA and PCI DSS subscriptions are not available for Splunk Cloud on GCP. For information regarding the availability of specific features and service components, see Differences between Splunk Cloud regions.

More information for regulated cloud environments is listed below.

Subscription Type	Region Availability	Encryption At Rest	IP Whitelist	Certification Documents
HIPAA	All AWS Data Centers except GovCloud (US-West).	Enabled by default. Splunk manages the encryption keys on your behalf and they are regularly rotated. Alternatively, you can request to manage the encryption keys	Customer must provide IP whitelisting rules.	If you require the HIPAA compliance report to review, contact your Splunk

Subscription Type	Region Availability	Encryption At Rest	IP Whitelist	Certification Documents
		instead.		sales representative to request a copy.
PCI DSS	All AWS Data Centers except GovCloud (US-West).	Enabled by default. Splunk manages the encryption keys on your behalf and they are regularly rotated. Alternatively, you can request to manage the encryption keys instead.	Customer must provide IP whitelisting rules.	If you require the PCI attestation of compliance to review, contact your Splunk sales representative to request a copy.

Differences between Splunk Cloud regions

Service Component	AWS Data Centers	GCP Data Centers
	US (Oregon, Virginia, GovCloud) EU (Dublin, Frankfurt, London) Asia Pacific (Singapore, Sydney, Tokyo, Seoul) Canada (Central)	US (Iowa) EU (London) Asia Pacific (Singapore)
Storage: DDAA	Available	Not currently available
Storage: DDSS	Available	Not currently available
Storage: Customer Managed Encryption Keys	Available, except GovCloud	Not currently available
Premium Solutions: Splunk App for VMware	Available	Not currently available
Compliance: PCI	Available	Not currently available
Compliance: HIPAA	Available	Not currently available
Compliance: SOC 2 Type II Controls	Available	Not currently available

Service limits and constraints

The following are Splunk Cloud service limits and constraints. You can use this list as guidance to ensure the best Splunk Cloud experience. Keep in mind that some limits depend on configuration, system load, performance, and available resources. Contact Splunk if your requirements are different or exceed what is recommended in this table.

Splunk Cloud service limits and constraints

Category	Service component	Limitation	Additional information
Data Collection	HEC maximum content length size limit	1MB	There is a limit to the HEC payload size in Splunk Cloud.
Data Egress	Dynamic Data Self-Storage export of aged data per index from Splunk Cloud to Amazon S3	No limit to the amount of data that can be exported from your indexes to your Amazon S3 account in the same region.	Dynamic Data Self-Storage is designed to export 1TB of data per hour.

Category	Service component	Limitation	Additional information
Data Egress	Search Results via UI or REST API	Recommend no more than 5% of ingested data	Data extracted as a result of search query, whether from the UI or REST API is limited to 5% of daily ingest for optimal performance. To route data to multiple locations, consider solutions like AWS Kinesis Data Firehose.
Data Egress	Search results to Splunk User Behavior Analytics (UBA)	No limit	Data as a result of search queries to feed into Splunk User Behavior Analytics (UBA).
Ingestion	Active indexes per Splunk Cloud environment	400	The best practice is to maintain no more than the upper limit of active indexes for each Splunk Cloud environment.
Other	Splunk Cloud ID	18 characters	Unique Splunk Cloud name chosen by you that determines your URL at [Splunk Cloud ID].splunkcloud.com. Must not exceed the 18-character count limit.
Search	Search concurrency per Splunk Cloud environment	20	This limit applies to customers with Splunk Cloud and a daily ingestion rate entitlement of less than 50 GB. There is a limit to the number of searches that each Splunk Cloud environment can concurrently process. Your Splunk Cloud search concurrency limits apply to searches initiated either from the Cloud search tier or from on-premise hybrid search heads. If the limit is reached, searches are queued, which can result in delays to producing search results.
Search	Search concurrency per Splunk Cloud environment	38	This starting point applies to customers with Splunk Cloud and a daily ingestion rate entitlement of between 50 GB and 1 TB. This starting point scales up at higher ingestion rates. There is a limit to the number of searches that each Splunk Cloud environment can concurrently process. Your Splunk Cloud search concurrency limits apply to searches initiated either from the Cloud search tier or from on-premise hybrid search heads. If the limit is reached, searches are queued, which can result in delays to producing search results.
Search	Search concurrency per Splunk Cloud or Splunk Cloud and ITSI environment	100	This starting point applies to customers with either Splunk Cloud or Splunk Cloud and ITSI and a daily ingestion rate entitlement of more than 1 TB. This starting point scales up at higher ingest rates. There is a limit to the number of searches that each Splunk Cloud environment can concurrently process. Your Splunk Cloud search concurrency limits apply to searches initiated either from the Cloud search tier or from on-premise hybrid search heads. If the limit is reached, searches are queued, which can result in delays to producing search results.
Search	Search concurrency per Premium Solution	38	When you add the following Premium Apps subscriptions to Splunk Cloud, additional search processes are available for each Premium App. This starting point scales up at higher ingestion rates. These search processes are additive to the Search concurrency per Splunk Cloud environment. • Splunk Enterprise Security • Splunk IT Service Intelligence • Splunk App for Microsoft Exchange

Category	Service component	Limitation	Additional information
			• Splunk App for PCI Compliance • Splunk App for VMware
Search	Join command for subsearch	50,000	The `join` command combines the results of a subsearch with the results of a main search. This limit is the maximum number of result rows in the output of a subsearch that can be joined against a main search. Refer to Splunk Cloud documentation of the join command for more information.
Security	Whitelist IP address rules per Splunk Cloud environment	100	Customers specify the IP address or IP address range that is permitted to access Splunk Cloud, and those from which Splunk Cloud can collect data. These are generically referred to as whitelist IP address rules. There is a limit to the total number of whitelist IP address rules per Splunk Cloud environment.
Workload Management	Workload Rules	100	You can configure up to 100 Workload Rules.

Enterprise Security service limits and constraints

Category	Service component	Limitation	Additional information
Enterprise Security	Correlation Search	60	This was the limit tested for Enterprise Security on Splunk Cloud. Correlation searches are a part of Enterprise Security, and are used to generate notable events or execute other adaptive response actions. Contact your Splunk sales representative if your use case requirements exceed this limit.
Enterprise Security	Data Models	9	This was the limit tested for Enterprise Security on Splunk Cloud. Data Models and Data Model Acceleration are critical components of Enterprise Security. To provide the best experience possible for customers, we suggest a maximum of 9 accelerated models. The most common data models deployed are: Change, Endpoint, Authentication, Intrusion Detection, Network Sessions, Network Resolution, Network Traffic, Web, and Performance.
Enterprise Security	Saved Searches	70	This was the limit tested for Enterprise Security on Splunk Cloud. Enterprise Security on Splunk Cloud uses saved searches for a variety of use-cases, such as search driven lookups. Saved Searches refer to any scheduled searches that run on the ES search tier. Overall ES performance can vary based on search schedule, timespan and search string.

IT Service Intelligence service limits and constraints

Category	Service component	Limitation	Additional information
Event Analytics	Correlation Searches	15	You can configure up to 15 Correlation Searches.
Event Analytics	Notable Event Aggregation Policies	15	You can configure up to 15 Notable Event Aggregation Policies.
Service Insights	Service Templates	500 Services per Service Template	You can configure up to 500 Services per Service Template and with a limit of 5000 services total.

Category	Service component	Limitation	Additional information

Supported Forwarder Versions

The following are the supported forwarder versions for Splunk Cloud. This information is applicable to universal and heavy forwarders that are communicating directly to Splunk Cloud. If you have deployed an intermediate forwarder tier communicating directly to Splunk Cloud, the following information applies to the forwarders in the intermediate tier instead of the forwarders indirectly connected. If you are unable to upgrade forwarders that communicate directly to Splunk Cloud, you accept the risk of continuing to use forwarder versions that have reached their end of support date.

Forwarder version	Supported Splunk Cloud versions	Forwarder version supported until
8.0.x	7.0.3+, 7.1.x, 7.2.x, 8.0.x	October 22, 2021
7.3.x	7.0.3+, 7.1.x, 7.2.x, 8.0.x	June 4, 2021
7.2.x	7.0.3+, 7.1.x, 7.2.x, 8.0.x	October 2, 2020
7.1.x	7.0.3+, 7.1.x, 7.2.x, 8.0.x	April 24, 2020
7.0.3+	7.0.3+, 7.1.x, 7.2.x, 8.0.x	June 30, 2019 - end of support milestone has been reached.
6.6.x or Less (TLS only)	6.6.3, 7.0.2	March 31, 2019 - end of support milestone has been reached.

Current versions

Splunk determines which versions of Splunk Cloud and Premium Apps to make available to Splunk Cloud subscribers. Splunk adopts the release that has the most benefits for customers as quickly as possible. The following are current versions for Splunk Cloud and Premium App subscriptions, as of May 2020.

Subscription	Version
Splunk Cloud	8.0
Splunk Enterprise Security	6.1
Splunk IT Service Intelligence	4.5
Splunk App for Microsoft Exchange	3.5
Splunk App for PCI Compliance	3.7
Splunk App for VMware	3.4

Note: Some Splunk Cloud versions have the following release numbering format that is unique and not available for Splunk Enterprise: [Major Release].[Minor Release]. [Release Date]
The [Release Date] is in the format of YYMM. For example, the 2001 of Splunk Cloud 8.0.2001 denotes a release date of January 2020.

More information

The following links provide information about the terms and policies that pertain to the Splunk Cloud service:

- **Legal**: Terms of Service
- **Level of service**: Splunk Cloud Service Level Schedule
- **Technical support**: Splunk Cloud Support Terms
- **Maintenance**: Service Maintenance Policy
- **Handling of data**: Splunk Cloud Data Policies

- **Splunk Data Privacy, Security and Compliance**: Splunk Protects