



**Splunk® Enterprise**

# **Splunk Enterprise Overview 9.0.4**

Generated: 4/07/2023 9:14 am

# Table of Contents

<b>About Splunk Enterprise.....</b>	<b>1</b>
About Splunk Enterprise.....	1
About Splunk Enterprise users.....	2
About Splunk Enterprise deployments.....	3
<b>Splunk Enterprise Resources and Documentation.....</b>	<b>7</b>
Support and resources for Splunk Enterprise.....	7
Splunk Enterprise administration.....	8
Search and reporting.....	10
Manage Splunk Enterprise knowledge.....	12
Customize and extend Splunk Enterprise.....	13
Troubleshooting.....	14

# About Splunk Enterprise

## About Splunk Enterprise

Splunk Enterprise is a software product that enables you to search, analyze, and visualize the data gathered from the components of your IT infrastructure or business. Splunk Enterprise takes in data from websites, applications, sensors, devices, and so on. After you define the data source, Splunk Enterprise indexes the data stream and parses it into a series of individual events that you can view and search.

Most users connect to Splunk Enterprise with a web browser and use **Splunk Web** to administer their deployment, manage and create knowledge objects, run searches, create pivots and reports, and so on. You can also use the command-line interface to administer your Splunk Enterprise deployment.

You can extend the Splunk Enterprise environment to fit the specific needs of your organization by using apps. An **app** is a collection of configurations, knowledge objects, views, and dashboards that runs on the Splunk platform. A single Splunk Enterprise installation can run multiple apps simultaneously. Browse available apps on Splunkbase or build your own on the Splunk developer site.

## Features of Splunk Enterprise

The following section highlights seven Splunk Enterprise features. You can read about more features on the Splunk Enterprise page at Splunk.com.

### ***Indexing***

Splunk Enterprise processes and stores the data that represents your business and its infrastructure. You can collect data from devices and applications such as websites, servers, databases, operating systems, and more. Once the data is collected, the index segments, stores, compresses the data, and maintains the supporting metadata to accelerate searching. To learn about getting your data into Splunk Enterprise, see *Get started with getting data in* in the *Getting Data In* manual. For more information on the indexing process, see *Indexes, indexers, and indexer clusters* in the *Managing Indexers and Clusters of Indexers* manual.

### ***Search***

Search is the primary way users navigate their data in Splunk Enterprise. You can save a search as a report and use it to power dashboard panels. Searches provide insight from your data, such as:

- Retrieving events from an index
- Calculating metrics
- Searching for specific conditions within a rolling time window
- Identifying patterns in your data
- Predicting future trends

### ***Alerts***

Alerts notify you when search results for both historical and real-time searches meet configured conditions. You can configure alerts to trigger actions like sending alert information to designated email addresses, posting alert information to an RSS feed, and running a custom script, such as one that posts an alert event to syslog.

## **Dashboards**

Dashboards contain panels of modules like search boxes, fields, charts, and so on. Dashboard panels are usually connected to saved searches or pivots. They display the results of completed searches and data from real-time searches that run in the background.

## **Pivot**

Pivot refers to the table, chart, or data visualization you create using the **Pivot Editor**. The Pivot Editor lets users map attributes defined by data model objects to a table, chart, or data visualization without having to write the searches in the **Search Processing Language (SPL)** to generate them. Pivots can be saved as reports and added to dashboards.

## **Reports**

Splunk Enterprise allows you to save searches and pivots as reports, and then add reports to dashboards as dashboard panels. Run reports on an ad hoc basis, schedule them to run on a regular interval, or set a scheduled report to generate alerts when the result meets particular conditions.

## **Data model**

Data models encode specialized domain knowledge about one or more sets of indexed data. They enable Pivot Editor users to create reports and dashboards without designing the searches that generate them.

## **Download the Splunk Enterprise Quick Reference Guide**

The Splunk Enterprise Quick Reference Guide is a 6-page PDF reference card that provides information about Splunk Enterprise features, concepts, search commands, and search examples.

## **About Splunk Enterprise users**

Splunk Enterprise serves different types of users. Here are five main personas that use Splunk Enterprise:

<b>Persona</b>	<b>Industry Role</b>	<b>Activities</b>
Administrator	network engineer, system administrator	Configures, administers, optimizes, and secures the Splunk Enterprise deployment Sets up user accounts and permissions Gets data into Splunk Enterprise
Knowledge Manager	data analyst, system administrator	Oversees knowledge object creation, normalization, and usage across teams, departments, and deployments Gets the data into Splunk Enterprise, or works with the administrator to do so Creates and shares data models
Search User	data analyst, IT professional, network	Uses Search to investigate server problems, understand configurations, monitor user activities, and troubleshoot escalated

Persona	Industry Role	Activities
	engineer, security analyst, system administrator	problems Builds reports and dashboards to monitor the health, performance, activity, and capacity of their IT infrastructure Identifies patterns and trends that are indicators of routine problems
Pivot User	business professional, data analyst, executive, IT professional, manager, system administrator	Uses Pivot to build reports based on data models created by the Knowledge Manager Creates reports and dashboards to monitor their businesses Identifies trends in the health and performance of their businesses
Developer	system integrator, professional developer	Integrates data and functionality of applications with Splunk Enterprise Builds Splunk apps and add-ons with custom dashboards and data visualizations

## About Splunk Enterprise deployments

Splunk Enterprise indexes data from the servers, applications, databases, network devices, and virtual machines that make up your IT infrastructure. As long as the machine that generates the data is a part of your network, Splunk Enterprise can collect the data from anywhere, whether the data is local, remote, or in the cloud.

Splunk Enterprise performs three main functions as it processes data:

1. It ingests data from files, the network, or other sources.
2. It parses and indexes the data.
3. It runs searches on the indexed data.

## Types of deployments

Depending on your needs, you can deploy Splunk Enterprise as a single instance, or you can create deployments that span multiple instances, ranging from just a few to hundreds or even thousands of instances.

### ***Single-instance deployments***

In small deployments, one instance of Splunk Enterprise handles all aspects of processing data, from input through indexing to search. A single-instance deployment can be useful for testing and evaluation purposes and might serve the needs of department-sized environments.

### ***Distributed deployments***

To support larger environments where data originates on many machines, where you need to process large volumes of data, or where many users need to search the data, you can scale the deployment by distributing Splunk Enterprise instances across multiple machines. This is known as a "distributed deployment".

In a typical distributed deployment, each Splunk Enterprise instance performs a specialized task and resides on one of three processing tiers corresponding to the main processing functions:

- Data input tier
- Indexer tier
- Search management tier

You can, for example, create a deployment with many instances that reside on the data input tier and only ingest data, several other instances that reside on the indexer tier and index the data, and one instance that resides on the search management tier and manages searches. These specialized instances are known as "components".

## Splunk Enterprise components and processing tiers

This table lists the processing components and the tiers that they occupy. It also describes the functions that each component performs.

Component	Tier	Description
Forwarder	Data input	A forwarder consumes data and then forwards the data onwards, usually to an indexer. Forwarders usually require minimal resources, allowing them to reside lightly on the machine generating the data.
Indexer	Indexing	An indexer indexes incoming data that it usually receives from a group of forwarders. The indexer transforms the data into events and stores the events in an index. The indexer also searches the indexed data in response to search requests from a search head.  To ensure high data availability and protect against data loss, or just to simplify the management of multiple indexers, you can deploy multiple indexers in indexer clusters.
Search head	Search management	A search head interacts with users, directs search requests to a set of indexers, and merges the results back to the user.  To ensure high availability and simplify horizontal scaling, you can deploy multiple search heads in search head clusters.

You can add components to each tier as necessary to support greater demands on that tier. For example, if you have a large number of users, you can add extra search heads to better service the users.

## Example of a distributed deployment

This diagram illustrates the type of deployment that might support the needs of a small enterprise.

### Small Enterprise Deployment



Less than 100 users

**Search Head**



Single search head

**Indexers**



...



2 to 3 indexers

**Forwarders**



...



100 to 200 forwarders

For more information about components and how to deploy them to scale Splunk Enterprise, see the *Distributed Deployment Manual*.



# Splunk Enterprise Resources and Documentation

## Support and resources for Splunk Enterprise

This topic is an overview of the support, documentation, and other resources available to help you find the information you want about Splunk Enterprise and other Splunk products.

### Support

Get support for Splunk Enterprise:

- Ask questions and get answers through community support at Splunk Answers.
- If you have a support contract, log a case using the Splunk Support Portal.
- If you have a support contract, contact customer support.

### Documentation

This section directs you to finding documentation on a specific product or task.

#### *Splunk Enterprise*

Use the following topics to locate the information you need within the Splunk Enterprise documentation:

- [Splunk Enterprise administration](#)
- [Search and reporting](#)
- [Manage Splunk Enterprise knowledge](#)
- [Customize and extend Splunk Enterprise](#)
- [Troubleshooting](#)

#### *Apps and add-ons*

Typically, an app's documentation is linked from the app's download page or included in the app's download package on Splunkbase. Splunk provides documentation for an app or add-on only if the product is supported by Splunk.

#### *Splunk SDKs*

The Splunk for Developers site provides information, tutorials, and examples on Splunk SDKs. For module libraries and other reference materials, see the Splunk documentation site for SDKs.

### Resources

Access additional resources for Splunk Enterprise:

- Read the Splunk Enterprise Quick Reference Guide for information about Splunk Enterprise features, concepts, search commands, and search examples.
- Join the Splunk user group Slack channel.
- Start a training or certification track on Splunk Education.
- Access more community resources on the Splunk Community page.

## Splunk Enterprise administration

This topic lists common administrator tasks and directs you to the relevant topics within the associated manuals.

### Install and upgrade Splunk Enterprise

The *Installation Manual* describes how to install and upgrade Splunk Enterprise.

Task:	Look here:
Understand installation requirements	Installation overview
Estimate hardware capacity needs	Introduction to capacity planning for Splunk Enterprise
Install Splunk Enterprise	Choose the Windows user Splunk Enterprise should run as Install on Linux Install on Mac OS X
Upgrade Splunk Enterprise	How to upgrade Splunk Enterprise
Perform backups	Back up configuration information Back up indexed data Set a retirement and archiving policy

### Get data into Splunk Enterprise

*Getting Data In* describes the types of Splunk data inputs and how to get data into your Splunk deployment.

Task:	Look here:
Learn how to consume external data	What data can I index?
Configure file and directory inputs	Monitor files and directories
Configure network inputs	Get data from TCP and UDP ports
Configure Windows inputs	Considerations for deciding how to monitor remote Windows data
Configure miscellaneous inputs	Monitor First In, First Out (FIFO) queues Monitor changes to your file system Get data from APIs and other remote data interfaces through scripted inputs
Enhance the value of your data	Overview of event processing How timestamp assignment works About indexed field extraction About hosts

Task:	Look here:
	Why source types matter About event segmentation
See how your data will look after indexing	The Set Sourcetype page
Improve the data input process	Use a test index to test your inputs
Understand the data pipeline	How data moves through Splunk Enterprise: the data pipeline

## Manage indexes and indexers

*Managing Indexers and Clusters* describes how to configure indexes and manage indexers, the components that maintain indexes.

Task:	Look here:
Learn about indexing	Indexes, indexers, and indexer clusters
Manage indexes	About managing indexes
Manage index storage	How the indexer stores indexes
Back up indexes	Back up indexed data
Archive indexes	Set a retirement and archiving policy
Learn about clusters and index replication	About indexer clusters and index replication
Deploy clusters	Indexer cluster deployment overview
Configure clusters	Manager configuration overview
Manage clusters	View the manager dashboard
Learn about cluster architecture	Basic indexer cluster concepts for advanced users

## Scale Splunk Enterprise

The *Distributed Deployment Manual* describes how to distribute Splunk Enterprise functionality across multiple components, such as forwarders, indexers, and search heads.

Task:	Look here:
Learn about Splunk Enterprise distributed deployments	Scale your deployment with Splunk Enterprise components
Perform capacity planning for Splunk deployments	Introduction to capacity planning for Splunk Enterprise
Learn how to forward data	About forwarding receiving
Distribute searches across multiple indexers	About distributed search
Deploy configuration updates across your environment	About deployment server and forwarder management

Associated manuals cover distributed components in detail:

- For information on forwarders, see the *Forwarding Data* manual.
- For information on search heads, see the *Distributed Search* manual.
- To manage your deployment using the deployment server and forwarder management, see the *Updating Splunk Enterprise Instances* manual.

## Secure Splunk Enterprise

*Securing Splunk Enterprise* describes how to secure your Splunk Enterprise deployment.

Task:	Look here:
Authenticate users and edit roles	About user authentication
Secure Splunk data with SSL	About securing Splunk Web
Audit Splunk Enterprise	Use Splunk Enterprise to audit your system activity Audit Splunk activity Use audit events to secure Splunk Enterprise Manage data integrity
Use Single Sign-on (SSO) with Splunk Enterprise	About Single Sign-On using reverse proxy
Use Splunk Enterprise with LDAP	Set up user authentication with LDAP

## Search and reporting

The Search and Reporting app lets you search your data, create data models and pivots, save your searches and pivots as reports, configure alerts, and create dashboards. This app is provided by default.

### Search

The *Search Manual* describes how to search and use the **Search Processing Language (SPL)**. See *Search Reference* for syntax, descriptions, and examples for each search command.

Task:	Look here:
Learn how to search and use the Search Processing Language	About the Search Tutorial
Learn more about the Search Processing Language	Get started with Search About the search language Understanding SPL syntax About transforming commands and searches About real-time searches and reports
Find a specific search command or function	Command quick reference Commands by category Evaluation functions Statistical and charting functions

Task:	Look here:
Manage search jobs	About jobs and jobs management View search job properties

## Create Pivots

The *Knowledge Manager Manual* describes how to design and build data models using the data model editor. The *Pivot Manual* describes how to build pivots tables and charts.

Task:	Look here:
Learn about data models and how to build them	About data models
Learn more about Pivot and how to use the Pivot Editor to design tables and charts	<i>Pivot Manual</i>

## Reports

See more about reports and report management in the *Reporting Manual*.

Task:	Look here:
Use search commands to generate reports	About transforming commands and searches
Learn about types of visualizations	Visualization reference Data structure requirements for visualizations
Save a search or pivot as a report	Create and edit reports
Accelerate a report Understand requirements for report acceleration	Accelerate reports
Schedule a report	Schedule reports
Generate a PDF of your report	Generate PDFs of your reports and dashboards

## Alerts

See how to create and dispatch alerts in the *Alerting Manual*.

Task:	Look here:
Learn about alerts	Getting started with alerts
Set up email notifications, RSS notifications, or alert scripts	Set up alert actions
See alerting examples	Alert examples
See recently triggered alerts	Triggered alerts
Set up alerts using the configuration files	Configure alerts in savedsearches.conf

## Create dashboards and visualizations

See the *Dashboards and Visualizations* manual for more information on the visualization and dashboard workflow and using the **Splunk Web Framework**.

Task:	Look here:
Learn about creating and editing dashboards	Dashboard overview
Learn about types of visualizations	Visualization reference
Learn about the default activity and summary dashboards	Splunk Enterprise summary dashboard
Learn about the Splunk Web Framework	Splunk Web Framework Overview

## Manage Splunk Enterprise knowledge

This topic lists common tasks in Splunk software knowledge management and directs you to the relevant topics for understanding and managing knowledge objects, such as events, fields, lookups, and data models.

### Splunk Enterprise knowledge

See the *Knowledge Manager Manual* for more information on using and maintaining knowledge objects.

Task:	Look here:
Understand Splunk Enterprise knowledge objects	What is Splunk knowledge? Understand and use the Common Information Model Add-on
Manage knowledge objects	Monitor and organize knowledge objects Disable or delete knowledge objects

## Events and event processing

See the *Knowledge Manager Manual* for more information on events. See the *Getting Data In* manual for more information on configuring event processing.

Task:	Look here:
Understand events and event types	About event types Define event types in Splunk Web
Configure event processing	Overview of event processing
Manage event segmentation	About event segmentation

## Fields and field extractions

See the *Knowledge Manager Manual* for more information on fields and field extractions.

Task:	Look here:
-------	------------

Understand fields	<a href="#">About fields</a> <a href="#">Use default fields</a> <a href="#">Configure extractions of multivalue fields with fields.conf</a> <a href="#">About calculated fields</a>
Understand and manage field extractions	<a href="#">About fields</a> <a href="#">When Splunk software extracts fields</a> <a href="#">About Splunk regular expressions</a>

## Build data models

See the *Knowledge Manager Manual* for more information on data models and using the Data Model Editor.

Task:	Look here:
Learn about data models and datasets	<a href="#">About data models</a>
Manage data models and datasets	<a href="#">Manage data models</a>
Use the Data Model Editor	<a href="#">Design data models</a>

## Customize and extend Splunk Enterprise

Developers can build Splunk apps to integrate Splunk Enterprise with other tools and applications. Follow these links to help you get started.

### Develop Splunk apps

Develop Splunk apps to build customized solutions for your specific data needs. For more information, see the *Developer Guide for Splunk Cloud Platform and Splunk Enterprise on the Splunk Developer Portal*.

Task:	Look here:
Learn about the Splunk app lifecycle	<a href="#">Lifecycle of a Splunk app for Splunk Cloud Platform or Splunk Enterprise</a>
See an overview of Splunk app development	<a href="#">Develop Splunk apps for Splunk Cloud Platform or Splunk Enterprise</a>
Learn how to build a Splunk app using React	<a href="#">Splunk React components for developers</a>
See an overview of how to release a Splunk app	<a href="#">Release and maintain Splunk apps for Splunk Cloud Platform or Splunk Enterprise</a>
See tutorials about Splunk app development	<a href="#">Tutorials for the Splunk platform</a>

### Use the Splunk REST API

Use the Splunk REST API to programmatically index, search, and visualize data using Splunk Enterprise in an external app.

Task:	Look here:
Learn how to use the Splunk REST API	<a href="#">Basic concepts about the Splunk platform REST API</a>

Task:	Look here:
See Splunk REST API tutorials	Splunk Enterprise Rest API Tutorials
Improve your logs to work with Splunk software	Logging in an app for Splunk Cloud Platform or Splunk Enterprise Logging best practices in an app or add-on for Splunk Enterprise
See the REST API Reference	Using the REST API Reference

## Download and install Splunk SDKs

Find information about integrating with the Splunk platform using the Splunk SDKs.

Task:	Look here:
Learn more about the Splunk SDKs	Developer tools for Splunk Cloud Platform or Splunk Enterprise
See the code library and examples for a Splunk SDK	Downloads on the Splunk Developer Portal
See the documentation site for Splunk SDKs	Splunk SDKs

## Extend Splunk platform functionality

Expand the Splunk platform to meet your specific data analysis needs.

Task:	Look here:
Extend the Splunk Search Processing Language	Create custom search commands for apps in Splunk Cloud Platform or Splunk Enterprise Define search macros in Settings Configure a script for an alert action
Create custom data inputs	Create custom data inputs for Splunk Cloud Platform or Splunk Enterprise
Create custom REST API endpoints	Extend the Splunk platform REST API with custom endpoints
Create custom workflow actions	Custom workflow actions for Splunk Cloud Platform or Splunk Enterprise
Create external lookups	Create external lookups for apps in Splunk Cloud Platform or Splunk Enterprise

## Troubleshooting

The *Troubleshooting Manual* describes how to analyze activity and diagnose problems with Splunk Enterprise. For specific troubleshooting information, see the associated manual for that topic. For example, you can find topics on how to improve search performance in the *Search Manual*.

Task:	Look here:
Learn about new features, known issues, and fixed problems	Welcome to Splunk Enterprise Known issues Fixed issues



Task:	Look here:
Learn about Splunk Enterprise troubleshooting tools	Introduction to troubleshooting Splunk Enterprise Use btool to troubleshoot configurations About the Monitoring Console
Use the platform instrumentation framework	About Splunk Enterprise platform instrumentation
Understand Splunk Enterprise log files	What Splunk software logs about itself About metrics.log
Troubleshoot search performance	Write better searches View search job properties
Troubleshoot license violations	About license violations Troubleshoot the license usage report view