



Splunk® Enterprise Installation Manual 9.0.4

Generated: 5/08/2023 6:18 pm

Table of Contents

Welcome to the Splunk Enterprise Installation Manual	1
What's in this manual.....	1
Plan your Splunk Enterprise installation	2
Installation overview.....	2
System requirements for use of Splunk Enterprise on-premises.....	3
Splunk Enterprise architecture and processes.....	9
Information on Windows third-party binaries that come with Splunk Enterprise.....	12
Installation instructions.....	14
Secure your Splunk Enterprise installation	15
About securing Splunk Enterprise.....	15
Secure your system before you install Splunk Enterprise.....	15
Install Splunk Enterprise securely.....	15
More ways to secure Splunk Enterprise.....	16
Install Splunk Enterprise on Windows	18
Choose the Windows user Splunk Enterprise should run as.....	18
Prepare your Windows network to run Splunk Enterprise as a network or domain user.....	20
Install on Windows.....	27
Install on Windows using the command line.....	31
Change the user selected during Windows installation.....	38
Install Splunk Enterprise on Linux or macOS	40
Install on Linux.....	40
Install on MacOS.....	42
Run Splunk Enterprise as a different or non-root user.....	44
Install Splunk Enterprise in virtual and containerized environments	46
Deploy and run Splunk Enterprise inside a Docker container.....	46
Start using Splunk Enterprise	48
Start Splunk Enterprise for the first time.....	48
What happens next?.....	52
Learn about accessibility to Splunk Enterprise.....	52
Install a Splunk Enterprise license	54
About Splunk Enterprise licenses.....	54
Install a license.....	54
Upgrade or migrate Splunk Enterprise	56
How to upgrade Splunk Enterprise.....	56
About upgrading to 9.0 READ THIS FIRST.....	59
How to upgrade a distributed Splunk Enterprise environment.....	75
Changes for Splunk App developers.....	76
Upgrade to version 9.0 on UNIX.....	77
Upgrade to version 9.0 on Windows.....	78

Table of Contents

Upgrade or migrate Splunk Enterprise	
Migrate a Splunk Enterprise instance from one physical machine to another.....	79
Plan your Splunk Enterprise upgrade to work with the Python 3 migration.....	83
Upgrade using the Python 3 runtime and dual-compatible Python syntax in custom scripts.....	83
Uninstall Splunk Enterprise.....	86
Uninstall Splunk Enterprise.....	86
Reference.....	88
PGP Public Key.....	88

Welcome to the Splunk Enterprise Installation Manual

What's in this manual

The *Installation Manual* provides the information that you need to install Splunk Enterprise.

Install Splunk Enterprise

- [System requirements](#)
- [Licensing information](#)
- [Procedures for installing](#)
- [Procedures for upgrading from a previous version](#)

Install the universal forwarder

To install the Splunk **universal forwarder**, see Install the universal forwarder software in the *Universal Forwarder* manual. The universal forwarder is a separate executable with its own set of installation procedures. For an introduction to forwarders, see About forwarding and receiving in the *Forwarding Data* manual.

Install a heavy forwarder

If you want to configure forwarding on a Splunk Enterprise instance, install the software and start the instance first, then enable forwarding on it.

1. Follow the instructions in this manual to install the right version of Splunk software for your computer operating system.
2. See Deploy a heavy forwarder for instructions on enabling heavy forwarding.

Plan your Splunk Enterprise installation

Installation overview

Installing Splunk Enterprise on a host is the first step in realizing value from your data. Read this topic and the contents of this chapter before you begin an installation.

There are two ways you can install Splunk Enterprise:

- Download and install a Splunk Enterprise installation package
- Download the Splunk Enterprise Docker image and run Splunk Enterprise inside a Docker container

Containerized Splunk Enterprise provides a simplified and consistent way for you to quickly get started with Splunk Enterprise and gain hands-on experience with the software. While Splunk Enterprise Docker containers are portable across different environments and allow for complex and scalable deployments, in this release, Splunk only supports the standalone and single-server Splunk topology for container-based deployments. For information about Docker, see the Docker documentation.

Install Splunk Enterprise by using an installation package

1. See the [system requirements](#) for installation. Additional requirements for installation might apply based on the operating system on which you install Splunk Enterprise and how you use Splunk Enterprise.
2. (Optional) See Components of a Splunk Enterprise deployment to learn about the Splunk Enterprise ecosystem, and [Splunk architecture and processes](#) to learn what the installer puts on your machine.
3. See [Secure your Splunk Enterprise installation](#) and, where appropriate, secure the machine on which you will install Splunk Enterprise.
4. Download the installation package for your system from the Splunk Enterprise download page.
5. (Optional) Migrate your KV store storage engine from the Memory Mapped (MMAP) storage engine to the WiredTiger storage engine to significantly reduce the amount of storage you need and to improve performance. See Migrate the KV store storage engine in the *Admin* manual to plan your migration.
6. Perform the installation by using the installation instructions for your operating system. See [Installation instructions](#).
7. (Optional) If this is the first time you have installed Splunk Enterprise, see the *Search Tutorial* to learn how to index data into Splunk software and search that data using the Splunk Enterprise search language.
8. (Optional) After you install Splunk Enterprise, calculate the amount of space your data takes up. See Estimate your storage requirements in the *Capacity Planning Manual*.
9. To run Splunk Enterprise in a production environment and to understand how much hardware such an environment requires, see the *Capacity Planning Manual*.

Deploy and run Splunk Enterprise inside Docker containers

1. Confirm that your system meets the following requirements for container-based installation:
 1. See the Containerized computing platforms section in [Supported Operating Systems](#) for supported operating systems.
 2. Confirm that your system meets or exceeds the recommended hardware requirements. See [Recommended hardware](#).
 3. Confirm that any disk volumes that you use to store Splunk Enterprise data inside a Docker container use one of the supported file systems. See [Supported file systems](#).

2. See [Secure your Splunk Enterprise installation](#) and, where appropriate, secure the machine on which you want to install Splunk Enterprise.
3. Download and install Docker Enterprise or Community Edition Engine 17.06.2 or higher for your operating system.
4. Perform the installation. See [Deploy and run Splunk Enterprise Docker containers](#) for step-by-step installation instructions.
5. (Optional) Estimate the amount of space your Splunk Enterprise data will take up. See Estimate your storage requirements in the *Capacity Planning Manual*.
6. Create and mount volumes to the containers for storing data that Splunk Enterprise uses and generates, such as indexed data and configuration files.
For instructions on configuring storage for data persistence, see Data Storage on Splunk Github.
7. To run Splunk Enterprise in a production environment and to understand how much hardware such an environment requires, see the *Capacity Planning Manual*.

Upgrade or migrate a Splunk Enterprise instance

In many cases, you can upgrade Splunk Enterprise over an existing version.

- To upgrade from an earlier version of Splunk Enterprise, see [How to upgrade Splunk Enterprise](#) in this manual for information and specific instructions.
- For information on migrating from one version to another, see the [About upgrading - READ THIS FIRST](#) topic for the version that you want to upgrade to.
- To move a Splunk Enterprise instance from one host to another, see [Migrate a Splunk instance](#).

System requirements for use of Splunk Enterprise on-premises

Splunk supports using Splunk Enterprise on several computing environments. Learn about the supported environments before you download the software.

The universal forwarder has its own set of hardware requirements. See Universal forwarder system requirements in the *Universal Forwarder* manual.

If you have ideas or requests for new features, use the Splunk Ideas portal to search for, vote on, and request new enhancements (called an idea) for any of the Splunk solutions. See Splunk Ideas in the *Get Started with Splunk Community* manual.

Supported Operating Systems

The following tables list the computing platforms for which Splunk Enterprise has support. The first table lists availability for *nix operating systems and the second lists availability for Windows operating systems.

Each table shows available computing platforms (operating system and architecture) and types of Splunk software. A bold **X** in a box that intersects the computing platform and Splunk software type you want means that Splunk software is available for that platform and type.

An empty box means that Splunk software is not available for that platform and type.

If you do not see the operating system or architecture that you are looking for in the list, the software is not available for that platform or architecture. This might mean that Splunk has ended support for that platform. See the list of deprecated and removed computing platforms in *Deprecated Features* in the *Release Notes*.

Some boxes contain characters other than a bold **X**. See the bottom of each table to learn what the characters mean and how that could affect your installation.

Confirm support for your computing platform

1. Find the operating system on which you want to install Splunk Enterprise in the **Operating system** column.
2. Find the computing architecture in the **Architecture** column that matches your environment.
3. Find the type of Splunk software that you want to use: Splunk Enterprise, Splunk Free, Splunk Trial, or Splunk Universal Forwarder.
4. If Splunk software is available for the computing platform and software type that you want, proceed to the download page to get it.

Unix operating systems

Operating system	Architecture	Enterprise License	Free License	Trial License	Universal Forwarder package
Linux, kernel version 5.4.x and higher	x86 (64-bit)	X	X	X	X
Linux, all 3.x and 4.x kernel versions	x86 (64-bit)	X	X	X	X
AIX 7.1 and 7.2	PowerPC				X
ARM Linux	ARM64 Graviton				X
	ARM64				X
	ARMv8				X
FreeBSD 11	x86 (64-bit)				X
macOS 12	Universal (Intel, M1)				X
macOS 11	Universal (Intel, M1)				X
macOS 10.15	Intel		D	D	X
macOS 10.14	Intel		D	D	
PowerLinux, Little Endian kernel version 3.0 and higher	PowerPC				X
Solaris 11	x86 (64-bit)				X
	SPARC				X
z/Linux, kernel version 3.0 and higher	s390x				X

X: Splunk software is available for the platform.

D: Splunk supports this platform and architecture, but might remove support in a future release. See *Deprecated Features* in the *Release Notes* for information on deprecation.

An empty box indicates software is not supported for this platform.

Windows operating systems

The table lists the Windows computing platforms that Splunk Enterprise supports.

Operating system	Architecture	Enterprise License	Free License	Trial License	Universal Forwarder package
Windows Server 2022 (all installation options)	x86 (64-bit)	X	X	X	X
Windows Server 2019 (all installation options)	x86 (64-bit)	X	X	X	X
Windows Server 2016 (all installation options)	x86 (64-bit)	D	D	D	X
Windows Server 2012 and Server 2012 R2	x86 (64-bit)				X
Windows 11	x86 (64-bit)				X
Windows 10	x86 (64-bit)		X	X	X
	x86 (32-bit)				X

X: Splunk software is available for the platform.

An empty box indicates software is not supported for this platform.

Containerized computing platforms

The official repository containing Dockerfiles for building Splunk Enterprise and Universal Forwarder images can be found on Splunk-Docker on GitHub. The list of requirements for Docker and Splunk software is available in the Support Guidelines on the Splunk-Docker GitHub.

For container orchestration, the Splunk Operator for Kubernetes on GitHub enables you to quickly and easily deploy Splunk Enterprise on your choice of private or public cloud provider. The operator simplifies scaling and management of Splunk Enterprise by automating workflows while implementing Kubernetes best practices.

Splunk Enterprise architecture support	Product
A single instance Splunk Enterprise deployment.	Splunk-Docker on GitHub
A distributed or single instance Splunk Enterprise deployment.	Splunk Operator for Kubernetes on GitHub

Operating system notes

Windows

Some parts of Splunk Enterprise on Windows require elevated user permissions to function properly. See the following topics for information on the components that require elevated permissions and how to configure Splunk Enterprise on Windows:

- [Splunk Enterprise architecture and processes](#)
- [Choose the Windows user Splunk Enterprise should run as](#)
- Considerations for deciding how to monitor remote Windows data in *Getting Data In*

Operating systems that support the Monitoring Console

The Splunk Enterprise Monitoring Console works only on some versions of Linux and Windows. For information on supported platform architectures for the Monitoring Console, see Supported platforms in the *Troubleshooting Manual*. To learn about the other prerequisites for the Monitoring Console, see Monitoring Console setup prerequisites in *Monitoring Splunk Enterprise*.

Deprecated operating systems and features

As we update Splunk software, we sometimes deprecate and remove support of older operating systems. See Deprecated features in the Release Notes for information on which platforms and features have been deprecated or removed entirely.

Creating and editing configuration files on OSES that do not use UTF-8 character set encoding

Splunk software expects configuration files to be in ASCII or Universal Character Set Transformation Format-8-bit (UTF-8) format. If you edit or create a configuration file on an OS that does not use UTF-8 character set encoding, then ensure that the editor you use can save in ASCII or UTF-8.

IPv6 platform support

All Splunk-supported OS platforms can use IPv6 network configurations.

See Configure Splunk Enterprise for IPv6 in the *Admin Manual* for details on IPv6 support in Splunk Enterprise.

Supported browsers

Splunk Enterprise supports the following browsers:

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

Recommended hardware

To evaluate Splunk Enterprise for a production deployment, use hardware that is typical of your production environment. This hardware should meet or exceed the recommended hardware capacity specifications. See Reference hardware in the *Capacity Planning Manual*.

For a discussion of hardware planning for production deployment, see Introduction to capacity planning for Splunk Enterprise in the *Capacity Planning Manual*.

Splunk Enterprise and virtual machines

If you run Splunk Enterprise in a virtual machine (VM) on any platform, performance decreases. This is because virtualization works by providing hardware abstraction on a machine into pools of resources. VMs that you define on the system draw from these resource pools. Splunk Enterprise needs sustained access to a number of resources, particularly disk I/O, for indexing operations. If you run Splunk Enterprise in a VM or alongside other VMs, indexing and search performance can degrade.

Splunk Enterprise and containerized infrastructures

A containerized deployment must provide hardware resources that meet or exceed the recommended hardware capacity for Splunk Enterprise deployments. See [Containerized computing platforms](#).

Recommended hardware capacity

For information on hardware requirements for production deployments, see Reference hardware in the *Capacity Planning Manual*.

Hardware requirements for universal forwarders

The universal forwarder has its own set of hardware requirements. See Universal forwarder prerequisites in the *Universal Forwarder* manual.

Supported file systems

If you run Splunk Enterprise on a file system that does not appear in this table, the software might run a startup utility named `locktest` to test the viability of the file system. If `locktest` fails, then the file system is not suitable for using with Splunk Enterprise.

Platform	File systems
Linux	ext3, ext4, btrfs, XFS, NFS 3/4
Solaris (universal forwarder only)	UFS, ZFS, VXFS, NFS 3/4
FreeBSD (universal forwarder only)	FFS, UFS, NFS 3/4, ZFS
Mac OS X	HFS, APFS, NFS 3/4
AIX (universal forwarder only)	JFS, JFS2, NFS 3/4
Windows	NTFS, FAT32

Considerations regarding Network File System (NFS)

When you use Network File System (NFS) as a storage medium for Splunk indexing, consider all of the ramifications of file level storage.

Use block level storage rather than file level storage for indexing your data.

In environments with reliable, high-bandwidth, low-latency links, or with vendors that provide high-availability, clustered network storage, NFS can be an appropriate choice. However, customers who choose this strategy should work with their hardware vendor to confirm that their storage platform operates to the vendor specification in terms of both performance and data integrity.

If you use NFS, note the following:

- Do not use NFS to host hot or warm index **buckets**. Splunk Enterprise on NFS is supported only with cold or frozen buckets.
- Do not use NFS to share cold or frozen index buckets amongst an indexer cluster, as this potentially creates a single point of failure.
- Splunk Enterprise does not support "soft" NFS mounts. These are mounts that cause a program attempting a file operation on the mount to report an error and continue in case of a failure.

- Only "hard" NFS mounts, where the client continues to attempt to contact the server in case of a failure, are reliable with Splunk Enterprise.
- Do not disable attribute caching. If you have other applications that require disabling or reducing attribute caching, then you must provide Splunk Enterprise with a separate mount with attribute caching enabled.
- Do not use NFS mounts over a wide area network (WAN). Doing so causes performance issues and can lead to data loss.

Considerations regarding system-wide resource limits on *nix systems

Splunk Enterprise allocates system-wide resources like file descriptors and user processes on *nix systems for monitoring, forwarding, deploying, and searching. The `ulimit` command controls access to these resources which must be tuned to acceptable levels for Splunk Enterprise to perform adequately on *nix systems.

The more tasks your Splunk Enterprise instance performs, the more resources it needs. You should increase the `ulimit` values if you start to see your instance run into problems with low resource limits. See I get errors about ulimit in `splunkd.log` in the *Troubleshooting Manual*.

The following table shows the system-wide resources that Splunk Enterprise uses. It provides the minimum recommended settings for these resources for instances that are not forwarders, such as indexers, search heads, cluster manager, license manager, deployment servers, and Monitoring Consoles (MC).

System-wide Resource	ulimit invocation	Minimum recommended value
Open files	<code>ulimit -n</code>	64000
User processes	<code>ulimit -u</code>	16000
Data segment size	<code>ulimit -d</code>	The maximum RAM you want Splunk Enterprise to allocate in kilobytes. For example, 8GB is 8000000.
File size	<code>ulimit -f</code>	-1 A setting of -1 sets the file size to unlimited.

On machines that run Linux where Splunk Enterprise services are managed by `systemd`, you can update the `/etc/systemd/system/Splunkd.service` unit file to set the values shown in the table below. Review the values and adjust them depending on the machine resources available.

System-wide Resource	systemd unit file parameter	Minimum recommended value
Open files	<code>LimitNOFILE=</code>	64000
User processes	<code>LimitNPROC=</code>	16000
Data segment size	<code>LimitDATA=</code>	The maximum RAM you want Splunk Enterprise to allocate in bytes. For example, 8GB is 8000000000.
File size	<code>LimitFSIZE=</code>	<code>infinity</code> A setting of "infinity" sets the file size to unlimited.
Total threads	<code>TasksMax=</code>	The maximum number of tasks that a service can create. This setting aligns with the user process limit <code>LimitNPROC</code> and the value can be set to match. For example, 16000.

On machines that run FreeBSD, you might need to increase the kernel parameters for default and maximum process stack size. The following table shows the parameters that must be present in `/boot/loader.conf` on the host.

System-wide Resource	Kernel parameter	Recommended value
Default process data size (soft limit)	<code>dfldsiz</code>	2147483648

System-wide Resource	Kernel parameter	Recommended value
Maximum process data size (hard limit)	maxdsiz	2147483648

On machines that run AIX, you might need to increase the systemwide resource limits for maximum file size (fsize) and resident memory size (rss). The following table shows the parameters that must be present in `/etc/security/limits` for the user that runs Splunk software.

System-wide Resource	ulimit invocation	Recommended value
Data segment size	ulimit -d	1073741824
Resident memory size	ulimit -m	536870912
Number of open files	ulimit -n	8192
File size limit	ulimit -f	-1 (unlimited)

This consideration is not applicable to Windows-based systems.

Considerations regarding Common Internet File System (CIFS)/Server Message Block (SMB)

Splunk Enterprise supports the use of the CIFS/SMB protocol for the following purposes, on shares hosted by Windows hosts only:

- Storage of cold or frozen **Index buckets**.

When you use a CIFS resource for storage, confirm that the resource has write permissions for the user that connects to the resource at both the file and share levels. If you use a third-party storage device, confirm that its implementation of CIFS is compatible with the implementation that your Splunk Enterprise instance runs as a client.

Do not index data to a mapped network drive on Windows (for example "Y:\\" mapped to an external share.) Splunk Enterprise disables any index it encounters with a non-physical drive letter.

Considerations regarding environments that use the transparent huge pages memory management scheme

If you run Splunk Enterprise on a Unix machine that makes use of transparent huge memory pages, see Transparent huge memory pages and Splunk performance in the *Release Notes* before you attempt to install Splunk Enterprise.

This consideration is not applicable to Windows operating systems.

Further reading

See the Download Splunk Enterprise page to get the latest available version.

See the release notes for details on known and resolved issues in this release.

See Introduction to Capacity Planning for Splunk Enterprise in the *Capacity Planning Manual* for information on estimating capacity .

Splunk Enterprise architecture and processes

This topic discusses the internal architecture and processes of Splunk Enterprise at a high level. If you're looking for information about third-party components used in Splunk Enterprise, see the credits section in the Release notes.

Splunk Enterprise Processes

A Splunk Enterprise server installs a process on your host, `splunkd`.

`splunkd` is a distributed C/C++ server that accesses, processes and indexes streaming IT data. It also handles search requests. `splunkd` processes and indexes your data by streaming it through a series of pipelines, each made up of a series of processors.

- **Pipelines** are single threads inside the `splunkd` process, each configured with a single snippet of XML.
- **Processors** are individual, reusable C or C++ functions that act on the stream of IT data that passes through a pipeline. Pipelines can pass data to one another through **queues**.
- New for version 6.2, `splunkd` also provides the Splunk Web user interface. It lets users search and navigate data and manage Splunk Enterprise deployment through a Web interface. It communicates with your Web browser through REpresentational State Transfer (REST).
- `splunkd` runs a Web server on port 8089 with SSL/HTTPS turned on by default.
- It also runs a Web server on port 8000 with SSL/HTTPS turned off by default.

Splunk Enterprise processes require network connectivity. For a table and diagrams showing the network ports used, see Components and their relationship with the network in the *Inherit a Splunk Enterprise Deployment* manual.

`splunkweb` installs as a legacy service on Windows only. Prior to version 6.2, it provided the Web interface for Splunk Enterprise. Now, it installs and runs, but quits immediately. You can configure it to run in "legacy mode" by changing a configuration parameter.

On Windows systems, `splunkweb.exe` is a third-party, open-source executable that Splunk renames from `python-service.exe`. Because it is a renamed file, it does not contain the same file version information as other Splunk Enterprise for Windows binaries.

[Read information on other Windows third-party binaries that come with Splunk Enterprise.](#)

Splunk Enterprise and Windows in Safe Mode

If Windows is in Safe Mode, Splunk services do not start. If you attempt to start Splunk Enterprise from the Start Menu while in Safe Mode, Splunk Enterprise does not alert you to the fact that its services are not running.

Additional processes for Splunk Enterprise on Windows

On Windows instances of Splunk Enterprise, in addition to the two services described, Splunk Enterprise uses additional processes when you create specific data inputs on a Splunk Enterprise instance. These inputs run when configured by certain types of Windows-specific data input.

splunk.exe

`splunk.exe` is the control application for the Windows version of Splunk Enterprise. It provides the command-line interface (CLI) for the program. It lets you start, stop, and configure Splunk Enterprise, similar to the *nix `splunk` program.

The `splunk.exe` binary requires an elevated context to run because of how it controls the `splunkd` and `splunkweb` processes. Splunk Enterprise might not function correctly if this program does not have the appropriate permissions on

your Windows system. This is not an issue if you install Splunk Enterprise as the Local System user.

splunk-admon

`splunk-admon.exe` runs whenever you configure an Active Directory (AD) monitoring input. `splunkd` spawns `splunk-admon`, which attaches to the nearest available AD domain controller and gathers change events generated by AD. Splunk Enterprise stores these events in an index.

splunk-perfmon

`splunk-perfmon.exe` runs when you configure Splunk Enterprise to monitor performance data on the local Windows machine. This binary attaches to the Performance Data Helper libraries, which query the performance libraries on the system and extract performance metrics both instantaneously and over time.

splunk-netmon

`splunk-netmon` runs when you configure Splunk Enterprise to monitor Windows network information on the local machine.

splunk-regmon

`splunk-regmon.exe` runs when you configure a Registry monitoring input in Splunk. This input initially writes a baseline for the Registry in its current state (if requested), then monitors changes to the Registry over time.

splunk-winevtlog

You can use this utility to test defined event log collections, and it outputs events as they are collected for investigation. Splunk Enterprise has a Windows event log input processor built into the engine.

splunk-winhostmon

`splunk-winhostmon` runs when you configure a Windows host monitoring input in Splunk. This input gets detailed information about Windows hosts.

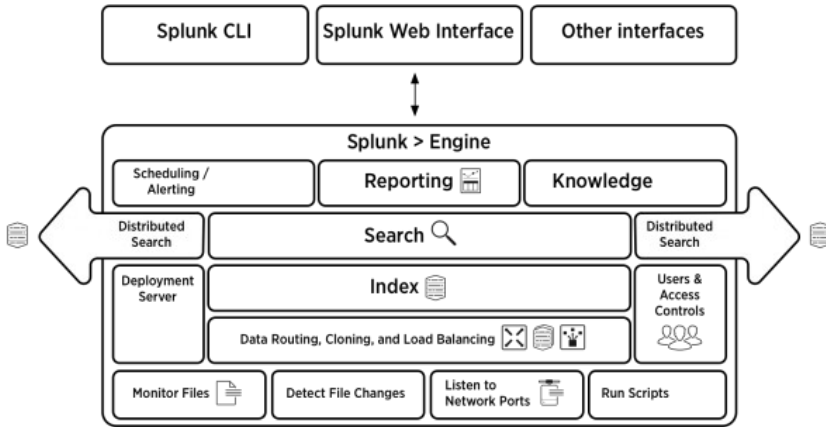
splunk-winprintmon

`splunk-winprintmon` runs when you configure a Windows print monitoring input in Splunk. This input gets detailed information about Windows printers and print jobs on the local system.

splunk-wmi

When you configure a performance monitoring, event log or other input against a remote computer, this program runs. Depending on how you configure the input, it either attempts to attach to and read Windows event logs as they come over the wire, or executes a Windows Query Language (WQL) query against the Windows Management Instrumentation (WMI) provider on the specified remote machine.

Architecture diagram



Information on Windows third-party binaries that come with Splunk Enterprise

Learn about the third-party Windows binaries that come with the Splunk Enterprise and the Splunk universal forwarder packages.

For more information about the universal forwarder, see *About forwarding and receiving data* in the *Forwarding Data Manual*.

Third-party Windows binaries that ship with Splunk Enterprise

The following third-party Windows binaries ship with Splunk Enterprise. The Splunk Enterprise product includes these binaries, except where indicated.

The binaries provide functionality to Splunk Enterprise as shown in their individual descriptions. The binaries do not contain file version information or authenticode signatures (certificates that prove the binary file's authenticity). Additionally, Splunk Enterprise does not provide support for debug symbols related to third-party modules.

Binaries, apps, and scripts that do not ship with Splunk Enterprise have not been tested for Certified for Windows Server 2008 R2 (CFW2008R2) Windows Logo compliance.

Archive.dll

Libarchive.dll is a multi-format archive and compression library.

Both Splunk Enterprise and the Splunk universal forwarder include this binary.

Bzip2.exe

Bzip2 is a patent-free, high-quality data compressor. It typically compresses files to within 10% to 15% of the best available techniques (the prediction by partial matching (PPM) family of statistical compressors), while being about twice as fast at compression and six times faster at decompression.

Jsmine.exe

Jsmine.exe is an executable that removes white space and comments from JavaScript files, reducing their size.

Libexslt.dll

Libexslt.dll is the Extensions to Extensible Stylesheet Language Transformation (EXSLT) dynamic link C library developed for libxslt (a part of the GNU is Not Unix Network Object Model Environment (GNOME) project).

Both Splunk Enterprise and the Splunk universal forwarder include this binary.

Libxml2.dll

Libxml2.dll is the Extensible Markup Language (XML) C parser and toolkit library. This library was developed for the GNOME project but can be used outside of the GNOME platform.

Both Splunk Enterprise and the Splunk universal forwarder include this binary.

Libxslt.dll

Libxslt.dll is the XML Stylesheet Language for Transformations (XSLT) dynamic link C library developed for the GNOME project. XSLT itself is an XML language to define transformation for XML. Libxslt is based on libxml2, the XML C library developed for the GNOME project. It also implements most of the EXSLT set of processor-portable extensions functions and some of Saxon's evaluate and expressions extensions.

Both Splunk Enterprise and the Splunk universal forwarder include this binary.

Minigzip.exe

Minigzip.exe is the minimal implementation of the 'gzip' compression tool.

Openssl.exe

The OpenSSL Project is a collaborative effort to develop a robust, commercial-grade, full-featured, and open source toolkit implementing the Secure Sockets Layer (SSL v2/v3) and Transport Layer Security (TLS v1) protocols as well as a full-strength general purpose cryptography library.

Both Splunk Enterprise and the Splunk universal forwarder include this binary.

Python.exe

Python.exe is the Python programming language binary for Windows.

Pythoncom.dll

Pythoncom.dll is a module that encapsulates the Object Linking and Embedding (OLE) automation API for Python.

Pywintypes27.dll

Pywintypes27.dll is a module that encapsulates Windows types for Python version 2.7.

Installation instructions

Use a link below for instructions to install Splunk Enterprise on your operating system:

- [Windows](#)
- [Windows \(from the command line\)](#)
- [Linux](#)

To use a containerized instance of Splunk Enterprise, see:

- [Deploy and run Splunk Enterprise inside a Docker container](#)

Splunk Enterprise for macOS 10.14, and 10.15 is available when using a free or trial license:

- [macOS](#)

Splunk Enterprise availability has been removed for some operating systems

Splunk Enterprise support was removed for these platforms, but a universal forwarder package is available:

- Solaris
- FreeBSD
- AIX

For a list of supported operating systems and versions for this release version, see [Supported Operating Systems](#).

Secure your Splunk Enterprise installation

About securing Splunk Enterprise

When you set up and begin using your Splunk Enterprise installation or upgrade, perform some additional steps to ensure that Splunk Enterprise and your data are secure. Taking the proper steps to secure Splunk Enterprise reduces its attack surface and mitigates the risk and impact of most vulnerabilities.

This section highlights some of the ways that you can secure Splunk Enterprise before, during, and after installation. The *Securing Splunk Enterprise* manual provides more information about the ways you can secure Splunk Enterprise.

Secure your system before you install Splunk Enterprise

Before you install Splunk Enterprise, make your operating system secure. Harden all Splunk Enterprise server operating systems.

- If your organization does not have internal hardening standards, use the CIS hardening benchmarks.
- At a minimum, limit shell and command-line access to your Splunk Enterprise servers.
- Secure physical access to all Splunk Enterprise servers.
- Ensure that Splunk Enterprise end users practice physical and endpoint security.

Install Splunk Enterprise securely

Verify integrity and signatures for your Splunk Installation when you download and install Splunk Enterprise.

Verify Integrity

Verify your Splunk Enterprise installation download using hash functions such as Message Digest 5 (MD5) and Secure Hash Algorithm-512 (SHA-512) to compare the hash fingerprints. Use a trusted version of OpenSSL.

MD5 validation

This procedure helps you compare the MD5 hash of the installation file you download from the Splunk website against the expected hash of the file. The tools you use to compare the files might be different based on the operating system that you run. You might need to download these tools before verifying the MD5 hash.

1. Download the installation package for the platform and version of Splunk software that you want.
2. On the "Thank you for downloading" page, click the link to the MD5 hash file for this package.
3. Open a shell prompt or Terminal window.
4. Print the contents of the MD5 hash file.

```
cat splunk-x.x.x-xxxxxxxxxxxx-Linux-x86-64.tgz.md5
MD5 (splunk-x.x.x-xxxxxxxxxxxx-Linux-x86_64.tgz) = c63c869754d420bb62f04f4096877481
```

5. Run the `md5` tool against the installer package.

```
md5 splunk-x.x.x-xxxxxxxxxxxx-Linux-x86-64.tgz
MD5 (splunk-x.x.x-xxxxxxxxxxxx-Linux-x86_64.tgz) = c63c869754d420bb62f04f4096877481
```

6. Compare the output of both commands.
7. If the hashes match, then you have confirmed that the installation package that you downloaded is the same as what is on the splunk.com website.

SHA512 validation

To compare the SHA512 hash of the installation file you download from the Splunk website against the expected hash of the file:

1. Check if the SHA comparison tool is already installed on your operating system.
2. Download the installation package for the platform and version of Splunk software that you want.
3. On the "Thank you for downloading" page, select and copy the link to the installer package. For example:
`https://download.splunk.com/products/splunk/releases/8.0.0/windows/splunk-8.0.0-1357bef0a7f6-x86-release.msi`
4. Append SHA512 to the end of the file extension in the link. For example:
`https://download.splunk.com/products/splunk/releases/8.0.0/windows/splunk-8.0.0-1357bef0a7f6-x64-release.msi.sha512`
5. Paste the link into a web browser to download the SHA512 hash file.
6. Verify the Splunk installation package and hash file are in the same location.
7. Run your SHA comparison tool against the hash file.
8. If the tool confirms that the hash matches the installation package, then you have confirmed that the installation package you downloaded is the same as the splunk.com hosted package.

Verify Signatures

Verify the authenticity of the downloaded RPM package by using the Splunk GnuPG Public key. The signature only applies to the RPM package. For all other package types, use the checksum files.

1. Download the GnuPG Public key file. (This link implements Transport Layer Security (TLS)).
2. Install the key.

```
rpm --import <filename>
```

3. Verify the package signature.

```
rpm -K <filename>
```

Example:

```
$ rpm -K splunk-8.0.0-1357bef0a7f6-linux-2.6-x86_64.rpm
```

```
splunk-8.0.0-1357bef0a7f6-linux-2.6-x86_64.rpm: rsa sha1 (md5) pgp md5 OK
```

More ways to secure Splunk Enterprise

After you install Splunk Enterprise, you have more options to secure your configuration.

Configure user authentication and role-based access control

Set up your users and use roles to control access. Splunk Enterprise lets you configure users in several ways. See the following information in *Securing Splunk Enterprise*.

- The built-in authentication system. See Set up user authentication with Splunk Enterprise native authentication.
- LDAP. See Set up user authentication with LDAP.
- A scripted authentication API for use with an external authentication system, such as Pluggable Authentication

Modules (PAM) or Remote Access Dial-In User Server (RADIUS). See Set up user authentication with external systems.

After you configure users, you can assign roles in Splunk Enterprise that determine and control capabilities and access levels. See About role-based user access.

Use SSL certificates to configure encryption and authentication

Splunk Enterprise comes with a set of default certificates and keys that, when enabled, provide encryption and data compression. You can also use your own certificates and keys to secure communications between your browser and Splunk Web as well as data sent from forwarders to a **receiver**, such as an indexer.

See "About securing Splunk with SSL" in this manual.

Audit Splunk Enterprise

Splunk Enterprise includes audit features that let you track the reliability of your data.

- Monitor files and directories in *Getting Data In*
- Search for audit events in *Securing Splunk Enterprise*

Harden your Splunk Enterprise installation

See the following topics in *Securing Splunk Enterprise* to harden your installation.

- Deploy secure passwords across multiple servers
- Use Splunk Enterprise Access Control Lists
- Secure your service accounts
- Disable unnecessary Splunk Enterprise components
- Secure Splunk Enterprise on your network

Install Splunk Enterprise on Windows

Choose the Windows user Splunk Enterprise should run as

When you install Splunk Enterprise on Windows, the software lets you select the Windows user that it should run as.

The user you choose depends on what you want Splunk Enterprise to monitor

The user that Splunk Enterprise runs as determines what Splunk Enterprise can monitor. The Local System user has access to all data on the local machine by default, but nothing else. A user other than Local System has access to whatever data you want, but you must give the user that access before you install Splunk Enterprise.

About the Local System user and other user choices

The Windows Splunk Enterprise installer provides two ways to install it:

- As the Local System user
- As another existing user on your Windows computer or network, which you designate

To do any of the following actions with Splunk Enterprise, you must install it as a domain user:

- Read Event Logs remotely
- Collect performance counters remotely
- Read network shares for log files
- Access the Active Directory schema using Active Directory monitoring

The user that you specify must meet the following requirements. If the user does not satisfy these requirements, Splunk Enterprise installation might fail. Even if installation succeeds, Splunk Enterprise might not run correctly, or at all.

- Be a member of the Active Directory domain or forest that you want to monitor (when using AD)
- Be a member of the local Administrators group on the server on which you install Splunk Enterprise
- Be assigned specific user security rights

If you are not sure which user Splunk Enterprise should run as, then see *Considerations for deciding how to monitor remote Windows data* in the *Getting Data In* manual for information on how to configure the Splunk Enterprise user with the access it needs.

User accounts and password concerns

The user that you select to run Splunk Enterprise as also has unique password constraints.

If you have a password enforcement security policy on your Windows network, that policy controls the validity of any user passwords. If that policy enforces password changes, you must do one of the following to keep Splunk Enterprise services running:

- Before the password expires, change it, reconfigure Splunk Enterprise services on every machine to use the changed password, and then restart Splunk Enterprise on each machine.
- Configure the account that Splunk Enterprise uses so that its password never expires.
- Use a managed service account. See "Use managed service accounts" later in this topic.

Use managed service accounts

You can use a managed service account (MSA) to run Splunk Enterprise if you can meet all of the following conditions:

- You run Windows Server 2008 R2 or later, or Windows 8 or later in Active Directory
- At least one domain controller in your Active Directory runs Windows Server 2008 R2 or later

The benefits of using an MSA are:

- Increased security from the isolation of accounts for services.
- Administrators no longer need to manage the credentials or administer the accounts. Passwords automatically change after they expire. They do not have to manually set passwords or restart services associated with these accounts.
- Administrators can delegate the administration of these accounts to non-administrators.

Some important things to understand before you install Splunk Enterprise with an MSA are:

- The MSA requires the same permissions as a domain account on the machine that runs Splunk Enterprise.
- The MSA must be a local administrator on the machine that runs Splunk Enterprise.
- You cannot use the same account on different machines, as you would with a domain account.
- You must correctly configure and install the MSA on the machine that runs Splunk Enterprise before you install Splunk Enterprise on the machine. See [Service Accounts Step-by-Step Guide on MS Technet](#).

To install Splunk Enterprise using an MSA, see [Prepare your Windows network for a Splunk Enterprise installation as a network or domain user](#).

Security and remote access considerations

Minimum permissions requirements

If you install Splunk Enterprise as a domain user, the machine that runs the instance requires that some default permissions change.

The `splunkd` and `splunkforwarder` services require specific user rights when you install Splunk Enterprise using a domain user. Depending on the sources of data you want to monitor, the Splunk Enterprise user might need additional rights. Failure to set these rights might result in a failed Splunk Enterprise installation, or an installation that does not function correctly.

Required basic permissions for the `splunkd` or `splunkforwarder` services

- Full control over the Splunk Enterprise installation directory.
- Read access to any files that you want to index.

Required Local/Domain Security Policy user rights assignments for the `splunkd` or `splunkforwarder` services

- Permission to log on as a service.
- Permission to log on as a batch job.
- Permission to replace a process-level token.
- Permission to act as part of the operating system.
- Permission to bypass traverse checking.

How to assign these permissions

This section provides guidance on how to assign the appropriate user rights and permissions to the Splunk Enterprise service account before you install. For procedures, see [Prepare your Windows network for a Splunk Enterprise installation as a network or domain user](#).

Use Group Policy to assign rights to multiple machines

To assign the policy settings to a number of machines in your AD forest, you can define a Group Policy object (GPO) with these rights, and deploy the GPO across the forest.

After you create and enable the GPO, the machines in the forest pick up the changes, either during the next scheduled AD replication cycle (usually every 1.5 to 2 hours), or at the next boot time. Alternatively, you can force AD replication by using the `GPUPDATE` command-line utility on the machine that you want to update Group Policy.

When you set user rights with a GPO, those rights override identical Local Security Policy rights on a machine. You cannot change this setting. To retain the Local Security Policy rights, you must assign those rights within the GPO.

Troubleshoot permissions issues

The rights described are the rights that the `splunkd` and `splunkforwarder` services require to run. The data you want to access might require that you assign additional rights. Many user rights assignments and other Group Policy restrictions can prevent Splunk Enterprise from running. If you have problems, consider using a tool such as Process Monitor or the `GPRESULT` command line tool to troubleshoot GPO application in your environment.

Prepare your Windows network to run Splunk Enterprise as a network or domain user

You can prepare your Windows network to run Splunk Enterprise as a network or domain user other than the "Local System" user.

This can be different from the user that you use to install the software. Regardless of the user that you run Splunk Enterprise as, you must install the software with an account that has local administrator privileges on the installation machine.

These instructions have been tested for Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2, and might differ for other versions of Windows.

The rights you assign by using these instructions are the minimum rights that are necessary for a successful Splunk Enterprise installation. You might need to assign additional rights, either within the Local Security Policy or a Group Policy object (GPO), or to the user and group accounts that you create, for Splunk Enterprise to access the data you want.

Security requirements and ramifications of changing system defaults through Group Policy

This procedure requires full administrative access to the host or Active Directory domain you want to prepare for Splunk Enterprise operations. Do not attempt to perform this procedure without this access.

The low-level access requirements for Splunk Enterprise operations necessitate these changes if you want to run Splunk Enterprise as a user other than the Local System user. You must make changes to your Windows network to complete

this procedure. Making these changes can present a significant security risk.

To mitigate the risk, you can prevent the user that Splunk Enterprise runs as from logging in interactively, and limit the number of machines from where the user can log in. Alternatively, on Windows Server 2008 R2 and later, you can set up managed user accounts (MSAs) that further limit risk.

If you are not comfortable with or do not understand the security risks that come with this procedure, then do not perform it.

Configure Active Directory for running Splunk software as a domain user

The following procedures prepare your Active Directory for installations of Splunk Enterprise or the Splunk universal forwarder as a domain user.

To use PowerShell to configure your Active Directory for installation of Splunk Enterprise, see "Use PowerShell to configure your AD domain" later in this topic.

Prerequisites

You must meet the following requirements to perform this procedure:

- Your Windows environment runs Active Directory.
- You are a domain administrator for the AD domains that you want to configure.
- The installation hosts are members of this AD domain.

Create users

When you create users for running Splunk Enterprise, follow Microsoft best practices . See Microsoft Best Practices on MS TechNet.

1. Run the Active Directory Users and Computers tool by selecting **Start > Administrative Tools > Active Directory Users and Computers**.
2. Select the domain that you want to prepare for Splunk Enterprise operations.
3. Click **Action > New > User**
4. Enter the username for the new user and click **Next**.
5. Uncheck **User must change password at next logon**.
6. Click **Next**.
7. Click **Finish**.
8. (Optional) Repeat this procedure to create additional users.
9. (Optional) Quit Active Directory Users and Computers.

Create groups

This procedure creates the groups for users and machines that run Splunk Enterprise.

1. Run the Active Directory Users and Computers tool by selecting **Start > Administrative Tools > Active Directory Users and Computers**.
2. Select the domain that you want to prepare for Splunk Enterprise operations.
3. Double-click an existing container folder, or create an Organization Unit by selecting **New > Group** from the **Action** menu.
4. Select **Action > New > Group**.

5. Type a name that represents Splunk Enterprise user accounts, for example, Splunk Accounts.
6. Confirm that the **Group scope** is set to **Domain Local** and **Group type** is set to **Security**.
7. Click **OK** to create the group.
8. Create a second group and specify a name that represents Splunk Enterprise enabled computers, for example, Splunk Enabled Computers. This group contains computer accounts that receive permissions to run Splunk Enterprise as a domain user.
9. Confirm that the **Group scope** is **Domain Local** and the **Group type** is **Security**.

Assign users and computers to groups

This part of the procedure assigns users and computers that you created in the previous part.

1. Add the accounts to the **Splunk Accounts** group.
2. Add the computer accounts of the computers that will run Splunk Enterprise to the **Splunk Enabled Computers** group.
3. (Optional) Quit **Active Directory Users and Computers**.

Define a Group Policy object (GPO)

The Group Policy Object you create here will be distributed to all of the machines that run Splunk Enterprise. It assigns rights to the machines that make running Splunk Enterprise easier.

1. Run the **Group Policy Management Console (GPMC)** tool by selecting **Start > Administrative Tools > Group Policy Management**
2. In the tree view pane on the left, select **Domains**.
3. Click the **Group Policy Objects** folder.
4. In the **Group Policy Objects in <your domain>** folder, right-click and select **New**.
5. Type a name that describes the fact that the GPO will assign user rights to the servers you apply it to. For example, "Splunk Access."
6. Leave the **Source Starter GPO** field set to "(none)".
7. Click **OK** to save the GPO.
8. Remain in the GPMC. You will perform additional work there in the next section.

Add rights to the GPO

1. While still in the GPMC, right-click on the newly-created group policy object and select **Edit**.
2. In the **Group Policy Management Editor**, in the left pane, browse to **Computer Configuration -> Policies -> Windows Settings -> Security Settings -> Local Policies -> User Rights Assignment**.
 1. In the right pane, double-click on the **Act as part of the operating system** entry.
 2. In the window that opens, check the **Define these policy settings** checkbox.
 3. Click **Add User or Group...**
 4. In the dialog that opens, click **Browse...**
 5. In the **Select Users, Computers, Service Accounts, or Groups** dialog that opens, type in the name of the "Splunk Accounts" group you created earlier and click **Check Names...** Windows underlines the name if it is valid. Otherwise it tells you that it cannot find the object and prompts you for an object name again.
 6. Click OK to close the "Select Users..." dialog.
 7. Click OK again to close the "Add User or Group" dialog.
 8. Click OK again to close the rights properties dialog.
3. Repeat Steps 2a-2h for the following additional rights:
 - ◆ **Bypass traverse checking**
 - ◆ **Log on as a batch job**

- ◆ **Log on as a service**
- ◆ **Replace a process-level token**

4. Remain in the Group Policy Management Editor. You will perform additional work there in the next section.

Change Administrators group membership on each host

This procedure restricts who is a member of the Administrators group on the hosts to which you apply this GPO.

Confirm that all accounts that need access to the Administrators group on each host have been added to the Restricted Groups policy setting. Failure to do so can result in losing administrative access to the hosts on which you apply this GPO!

1. While still in the Group Policy Management Editor window, in the left pane, browse to **Computer Configuration -> Policies -> Windows Settings -> Security Settings -> Restricted Groups**.
 1. In the right pane, right-click and select **Add Group...** in the pop-up menu that appears.
 2. In the dialog that appears, type in **Administrators** and click OK.
 3. In the properties dialog that appears, click the **Add** button next to **Members of this group:**.
 4. In the **Add Member** dialog that appears, click **Browse..."**
 5. In the **Select Users, Computers, Service Accounts, or Groups** dialog that opens, type in the name of the "Splunk Accounts" group you created earlier and click **Check Names...** Windows underlines the name if it is valid. Otherwise it tells you that it cannot find the object and prompts you for an object name again.
 6. Click OK to close the **Select Users...** dialog.
 7. Click OK again to close the "Add User or Group" dialog.
 8. Click OK again to close the group properties dialog.
2. Repeat Steps 1a-1h for the following additional users or groups:
 - ◆ Domain Admins
 - ◆ any additional users who need to be a member of the Administrators group on every host to which you apply the GPO.
3. Close the Group Policy Management Editor window to save the GPO.
4. Remain in the GPMC. You will perform additional work there in the next section.

Restrict GPO application to select computers

This procedure controls which machines will actually receive the new GPO, and thus have their user rights assignments changed so that they can run Splunk Enterprise.

1. While still in the GPMC, in the GPMC left pane, select the GPO you created and added rights to, if it is not already selected. The GPMC displays information about the GPO in the right pane.
2. In the right pane, under **Security Filtering**, click **Add...**
3. In the **Select User, Computer, or Group** dialog that appears, type in "Splunk Enabled Computers" (or the name of the group that represents Splunk-enabled computers that you created earlier.)
4. Click **Check Names**. If the group is valid, Windows underlines the name. Otherwise, it tells you it cannot find the object and prompts you for an object name again.
5. Click OK to return to the GPO information window.
6. Repeat Steps 2-5 to add the "Splunk Accounts" group (the group that represents Splunk user accounts that you created earlier.)
7. Under **Security Filtering**, click the **Authenticated Users** entry to highlight it.
8. Click **Remove**. GPMC removes the "Authenticated Users" entry from the "Security Filtering" field, leaving only "Splunk Accounts" and "Splunk Enabled Computers."
9. Remain in the GPMC. You will perform additional work there in the next section.

Apply the GPO

Active Directory controls when Group Policy updates occur and GPOs get applied to hosts in the domain. Under normal circumstances, replication happens every 90-120 minutes. You must either wait this amount of time before attempting to install Splunk as a domain user, or force a Group Policy update by running `GPUPDATE /FORCE` from a command prompt on the host whose Group Policy you want to update.

1. While still in the GPMC, in the GPMC left pane, select the domain that you want to apply the GPO you created.
2. Right click on the domain, and select **Link an Existing GPO...** in the menu that pops up.

If you only want the GPO to affect the OU that you created earlier, then select the OU instead and right-click to bring up the pop-up menu.

3. In the **Select GPO** dialog that appears, select the GPO you created and edited, and click **OK**. GPMC applies the GPO to the selected domain.
4. Close GPMC by selecting **File > Exit** from the GPMC menu.

Install Splunk with a managed system account

Alternatively, you can install Splunk Enterprise with a managed system account.

You can use the instructions in "Configure Active Directory for running Splunk software as a domain user" earlier in this topic to assign the MSA the appropriate security policy rights and group memberships.

When you grant file permissions to the MSA after installation, you might need to break NTFS permission inheritance from parent directories above the Splunk Enterprise installation directory and explicitly assign permissions from that directory and all subdirectories.

Windows grants the "Log on as a service" right to the MSA automatically if you use the Services control panel to make changes to Splunk services.

1. Create and configure the MSA that you plan to use to monitor Windows data.
2. [Install Splunk from the command line](#) and use the `LAUNCHSPLUNK=0` flag to keep Splunk Enterprise from starting after installation has completed.
3. After installation completes, use the Windows Explorer or the `ICACLS` command line utility to grant the MSA "Full Control" permissions to the Splunk Enterprise installation directory and all its sub-directories.
4. Change the default user for the `splunkd` and `splunkweb` service accounts, as described in the topic [Correct the user selected during Windows installation](#).

You must append a dollar sign (\$) to the end of the username when completing this step for the MSA to work. For example, if the MSA is `SPLUNKDOCS\splunk1`, then you must enter `SPLUNKDOCS\splunk1$` in the appropriate field in the properties dialog for the service. You must do this for both the `splunkd` and `splunkweb` services.

5. Confirm that the MSA has the **"Log on as a service"** right.
6. Start Splunk Enterprise. It runs as the MSA configured above, and has access to all data that the MSA has access to.

Use PowerShell to configure your AD domain

You can use PowerShell to configure your Active Directory environment for Splunk Enterprise services. This option is available when you do not want to use the GUI-based administrative applications.

Create the Splunk user account

1. Open a PowerShell window.
2. Import the ActiveDirectory PowerShell module, if needed:

```
> Import-Module ActiveDirectory
```

3. Create a new user:

```
> New-ADUser -Name <user> `
-SamAccountName <user> `
-Description "Splunk Service Account" `
-DisplayName "Service:Splunk" `
-Path "<organizational unit LDAP path>" `
-AccountPassword (Read-Host -AsSecureString "Account Password") `
-CannotChangePassword $true `
-ChangePasswordAtLogon $false `
-PasswordNeverExpires $true `
-PasswordNotRequired $false `
-SmartcardLogonRequired $false `
-Enabled $true `
-LogonWorkstations "<server>" `
```

In this example:

- ◆ The command creates an account whose password cannot be changed, is not forced to change after first logon, and does not expire.
- ◆ *<user>* is the name of the user you want to create.
- ◆ *<organizational unit LDAP path>* is the name of the OU in which to put the new user, specified in X.500 format, for example: `CN=Managed Service Accounts,DC=splk,DC=com`.
- ◆ *<server>* is a single host or comma-separated list which specifies the host(s) that the account can log in from.

The LogonWorkstations argument is not required, but lets you limit which workstations a managed service account can use to log into the domain.

Configure the Splunk Enterprise server

After you have configured a user account, use PowerShell to configure the server with the correct permissions for the account to run Splunk Enterprise.

This is an advanced procedure. Improper changes to your AD can render it unusable. Perform these steps only if you feel comfortable doing so and understand the ramifications of using them, including problems that can occur due to typos and improperly-formatted files.

In the following examples:

- *<user>* is the name of the user you created that will run Splunk Enterprise.
- *<domain>* is the domain in which the user resides.
- *<computer>* is the remote computer you want to connect to in order to make changes.

To configure local security policy from PowerShell:

1. Connect to the machine that you wish to configure.
 - ◆ If you use the local machine, log in and open a PowerShell prompt, if you have not already.
 - ◆ If you connect to a remote machine, create a new `PSSession` on the remote host, as shown in the following examples.
 - ◆ You might need to disable Windows Firewall before you can make the remote connection. To do so, see [Need to Disable Windows Firewall on MS TechNet](#) (for versions of Windows Server up to Server 2008 R2, and [Firewall with Advanced Security Administration with Windows PowerShell](#), also on MS TechNet.

```
> Enter-PSSession -Computersname <computer>
```

2. Add the service account to the local Administrators group.

```
> $group = [ADSI]"WinNT://<server>/Administrators,group"
> $group.Add("WinNT://<domain>/<user>")
```

3. Create a backup file that contains the current state of user rights settings on the local machine.

```
> secedit /export /areas USER_RIGHTS /cfg OldUserRights.inf
```

4. Use the backup to create a new user rights information file that assigns the Splunk Enterprise user elevated rights when you import it.

```
> Get-Content OldUserRights.inf `
| Select-String -Pattern `
"(SeTcbPrivilege|SeChangeNotify|SeBatchLogon|SeServiceLogon|SeAssignPrimaryToken|SeSystemProfile)" `
| %{ "$_,<domain>\<user>" }
| Out-File NewUserRights.inf
```

5. Create a header for the new policy information file and concatenate the header and the new information file together.

```
> ( "[Unicode]", "Unicode=yes" ) | Out-File Header.inf
> ( "[Version]", "signature=```$CHICAGO`$`", "Revision=1" ) | Out-File -Append Header.inf
> ( "[Privilege Rights]" ) | Out-File -Append Header.inf
> Get-Content NewUserRights.inf | Out-File -Append Header.inf
```

6. Review the policy information file to ensure that the header was properly written, and that the file has no syntax errors in it.
7. Import the file into the local security policy database on the host.

```
> secedit /import /cfg Header.inf /db C:\splunk-lsp.sdb
> secedit /configure /db C:\splunk-lsp.sdb
```

Prepare a local machine or non-AD network for Splunk Enterprise installation

If you do not use Active Directory, follow these instructions to give administrative access to the user you want Splunk Enterprise to run as on the hosts on which you want to install Splunk Enterprise.

1. Give the user Splunk Enterprise should run as administrator rights by adding the user to the local Administrators group.
2. Start Local Security Policy by selecting **Start > Administrative Tools > Local Security Policy**.
3. In the left pane, expand **Local Policies** and then click **User Rights Assignment**.
 1. In the right pane, double-click on the **Act as part of the operating system** entry.
 2. Click **Add User or Group...**
 3. Click **Browse...**
 4. Type in the name of the "Splunk Computers" group you created earlier, and click **Check Names...** Windows underlines the name if it is valid. Otherwise it tells you that it cannot find the object and prompts you for an object name again.
 5. Click **OK**.
 6. Click **OK**.

7. Click **OK**.
4. Repeat Steps 3a-3g for the following additional rights:
 - ◆ **Bypass traverse checking**
 - ◆ **Log on as a batch job**
 - ◆ **Log on as a service**
 - ◆ **Replace a process-level token**

After you have completed these steps, you can then [install Splunk Enterprise](#) as the desired user.

Install on Windows

You can install Splunk Enterprise on Windows with the Graphical User Interface (GUI)-based installer or from the command line. More options, such as silent installation, are available if you install from the command line. See [Install on Windows from the command line](#) for the command line installation procedure.

You cannot install or run the 32-bit version of Splunk Enterprise for Windows on a 64-bit Windows machine. You also cannot install Splunk Enterprise on a machine that runs an unsupported OS. For example, you cannot install Splunk Enterprise on a machine that runs Windows Server 2003. See [System requirements](#). If you attempt to run the installer in such a way, it warns you and prevents the installation.

Note: If, rather than installing Splunk Enterprise, you want to install the Splunk **universal forwarder**, see Install a Windows universal forwarder in the *Universal Forwarder* manual. The universal forwarder is a separate executable from Splunk Enterprise and uses a different installer.

Upgrading?

If you plan to upgrade Splunk Enterprise, see [How to upgrade Splunk Enterprise](#) for instructions and migration considerations before proceeding.

Before you install

Choose the Windows user Splunk should run as

Before installing, see [Choose the Windows user Splunk should run as](#) to determine which user account Splunk should run as to address your specific needs. The user you choose has ramifications on what you must do prior to installing the software, and more details can be found there.

Disable or limit antivirus software if able

The Splunk Enterprise indexing subsystem requires high disk throughput. Any software with a device driver that intermediates between Splunk Enterprise and the operating system can restrict processing power available to Splunk Enterprise, causing slowness and even an unresponsive system. This includes anti-virus software.

You must configure such software to avoid on-access scanning of Splunk Enterprise installation directories and processes before you start a Splunk installation.

Consider installing Splunk software into a directory with a short path name

By default, the Splunk MSI file installs the software to `\Program Files\Splunk` on the system drive (the drive that booted your Windows machine.) While this directory is fine for many Splunk software installations, it might be problematic for

installations that run in distributed deployments or that employ advanced Splunk features such as search-head or indexer clustering.

The Windows API has a path limitation of `MAX_PATH` which Microsoft defines as 260 characters including the drive letter, colon, backslash, 256-characters for the path, and a null terminating character. Windows cannot address a file path that is longer than this, and if Splunk software creates a file with a path length that is longer than `MAX_PATH`, it cannot retrieve the file later. There is no way to change this configuration.

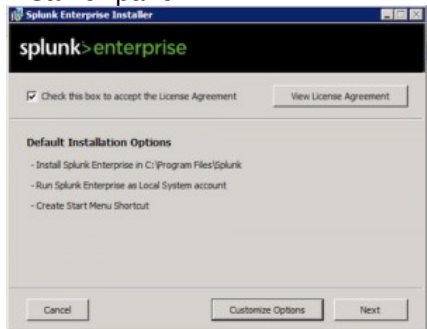
To work around this problem, if you know that the instance will be a member of a search head or indexer cluster, consider installing the software into a directory with a short path length, for example `C:\Splunk` or `D:\SPL`.

Install Splunk Enterprise via the GUI installer

The Windows installer is an MSI file.

Begin the installation

1. Download the Splunk installer from the Splunk download page.
2. To start the installer, double-click the `splunk.msi` file. The installer runs and displays the **Splunk Enterprise Installer** panel.



3. To continue the installation, check the "Check this box to accept the License Agreement" checkbox. This activates the "Customize Installation" and "Next" buttons.
4. (Optional) If you want to view the license agreement, click **View License Agreement**.

Installation Options

The Windows installer gives you two choices: Install with the default installation settings, or configure all settings prior to installing.

When you choose to install with the default settings, the installer does the following:


- Installs Splunk Enterprise in `\Program Files\Splunk` on the drive that booted your Windows machine.
- Installs Splunk Enterprise with the default management and Web network ports.
- Configures Splunk Enterprise to run as the Local System user.
- Prompts you to create a Splunk administrator password. You must do this before installation can continue.
- Creates a Start Menu shortcut for the software.

If you want to change any of these default installation settings, click **Customize Options** and proceed with the instructions in "Customize Options" in this topic.

Otherwise, click **Next**. You will be prompted for a password for the Splunk admin user. After you supply a password, installation begins and you can continue with the "Complete the install" instructions later in this topic.

Customize options during the installation

You can customize several options during the installation. When you choose to customize options, the installer displays the "Install Splunk Enterprise to" panel.



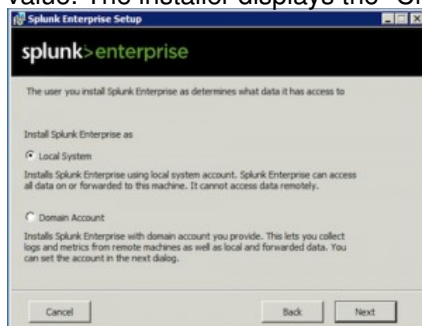
By default, the installer puts Splunk Enterprise into `\Program Files\Splunk` on the system drive. This documentation set refers to the Splunk Enterprise installation directory as `$SPLUNK_HOME` or `%SPLUNK_HOME%`.

Splunk Enterprise installs and runs two Windows services, `splunkd` and `splunkweb`. The `splunkd` service handles all Splunk Enterprise operations, and the `splunkweb` service installs to run only in legacy mode.

These services install and run as the user you specify on the "Choose the user Splunk Enterprise should run as" panel. You can choose to run Splunk Enterprise as the Local System user, or another user.

When the installer asks you the user that you want to install Splunk Enterprise as, you must specify the user name in `domain\username` format. The user must be a valid user in your security context, and must be an active member of an Active Directory domain. Splunk Enterprise must run under either the Local System account or a valid user account with a valid password and local administrator privileges. Failure to include the domain name with the user will cause the installation to fail.

1. Click **Change...** to specify a different location to install Splunk Enterprise, or click **Next** to accept the default value. The installer displays the "Choose the user Splunk Enterprise should run as" panel.



2. Select a user type and click **Next**.
3. If you selected the Local System user, proceed to Step 5. Otherwise, the installer displays the **Logon Information: specify a username and password** panel.



4. Enter the Windows credentials that Splunk Enterprise uses to run on the machine and click **Next**.

These credentials are different from the Splunk administrator credentials that you create in the next step.



5. Create credentials for the Splunk administrator user by entering a username and password that meets the minimum eligibility requirements as shown in the panel and click **Next**.

You must perform this action as the installation cannot proceed without your completing it. If you do not enter a username, the installer creates the admin user during the installation process.

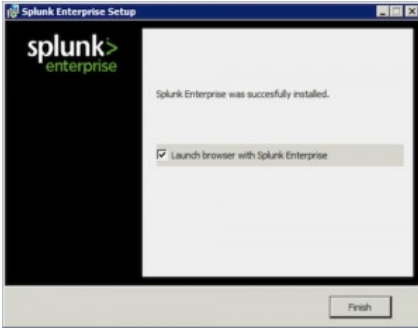
6. The installer displays the installation summary panel.



7. Click "Install" to proceed with the installation.

Complete the installation

The installer runs, installs the software, and displays the **Installation Complete** panel.



If you specified the wrong user during the installation procedure, you will see two pop-up error windows explaining this. If this occurs, Splunk Enterprise installs itself as the Local System user by default. Splunk Enterprise does not start automatically in this situation. You can proceed through the final panel of the installation, but uncheck the "Launch browser with Splunk" checkbox to prevent your browser from launching. Then, use [these instructions](#) to switch to the correct user before starting Splunk.

1. (Optional) Check the boxes to **Launch browser with Splunk** and **Create Start Menu Shortcut**.
2. Click **Finish**. The installation completes, Splunk Enterprise starts and launches in a supported browser if you checked the appropriate box.

Install or upgrade license

If this is a new installation of Splunk Enterprise or switching from one license type to another, you must install or update your license. See [Install a license](#).

Next steps

Now that you have installed Splunk Enterprise, you can find out how to start using Splunk Enterprise. See [What happens next?](#)

Alternatively, you can see the following topics in *Getting Data In* for help on adding Windows data:

- Monitor Windows Event Log data
- Monitor Windows Registry data
- Monitor WMI-based data
- Considerations for deciding how to monitor remote Windows data.

Install on Windows using the command line

You can install Splunk Enterprise on Windows from the command line.

Do not run the 32-bit installer on a 64-bit system. If you attempt this, the installer warns you and prevents installation.

If you want to install the Splunk **universal forwarder** from the command line, see "Install a Windows universal forwarder" in the *Universal Forwarder* manual.

When to install from the command line

You can manually install Splunk Enterprise on individual machines from a command prompt or PowerShell window. Here are some scenarios where installing from the command line is useful:

- You want to install Splunk Enterprise, but do not want it to start right away
- You want to automate installation of Splunk Enterprise with a script
- You want to install Splunk Enterprise on a system that you will clone later
- You want to use a deployment tool such as Group Policy or System Center Configuration Manager
- You want to install Splunk Enterprise on a system that runs a version of Windows Server Core

Install using PowerShell

You can install Splunk Enterprise from a PowerShell window. The steps to do so are identical to those that you use to install from a command prompt.

Upgrading?

To upgrade Splunk Enterprise, see [How to upgrade Splunk](#) for instructions and migration considerations.

Splunk Enterprise does not support changing the management or Splunk Web ports during an upgrade.

Prerequisites to installing Splunk Enterprise on Windows

Choose the Windows user Splunk Enterprise should run as

Before you install, see [Choose the Windows user Splunk Enterprise should run as](#) to determine which user account Splunk Enterprise should run as to address your data collection needs. The user you choose has specific ramifications on what you need to do before you install the software.

Prepare your domain for a Splunk Enterprise installation as a domain user

The Windows network should be configured to support a Splunk Enterprise installation.

Before you install, see [Prepare your Windows network for a Splunk Enterprise installation as a network or domain user](#) for instructions about how to configure your domain to run Splunk Enterprise.

Disable or limit antivirus software if able

The Splunk Enterprise indexing subsystem requires high disk throughput. Any software with a device driver that intermediates between Splunk Enterprise and the operating system can restrict the processing power that is available to Splunk Enterprise. This can cause slowness and even an unresponsive system. This includes anti-virus software.

You must configure such software to avoid on-access scanning of Splunk Enterprise installation directories and processes before you start a Splunk installation

Have credentials for the Splunk administrator user ready

When you install Splunk Enterprise, you must create a username and password for the Splunk administrator user. The installer does not create credentials for the user by default. Think of a username and password combination and be ready to supply it when you perform the installation. If you do not supply at least a password during a silent installation, Splunk Enterprise can install without any users defined, which prevents login. You must then create a user-seed.conf file to fix the problem and restart the software.

Consider installing Splunk software into a directory with a short path name

By default, the Splunk MSI file installs the software to `\Program Files\Splunk` on the system drive (the drive that booted your Windows machine.) While this directory is fine for many Splunk software installations, it might be problematic for installations that run in distributed deployments or that employ advanced Splunk features such as accelerated data models, search-head or indexer clustering.

The Windows API has a path limitation of `MAX_PATH` which Microsoft defines as 260 characters including the drive letter, colon, backslash, 256-characters for the path, and a null terminating character. Windows cannot address a file path that is longer than this, and if Splunk software creates a file with a path length that is longer than `MAX_PATH`, it cannot retrieve the file later. There is no way to change this configuration.

To work around this problem consider installing the software into a directory with a short path length, for example

`C:\Splunk` OR `D:\SPL`.

Install Splunk Enterprise from the command line

Invoke `msiexec.exe` to install Splunk Enterprise from the command line or a PowerShell prompt.

For 32-bit platforms, use `splunk-<...>-x86-release.msi`:

```
msiexec.exe /i splunk-<...>-x86-release.msi [<flag>]... [/quiet]
```

For 64-bit platforms, use `splunk-<...>-x64-release.msi`:

```
msiexec.exe /i splunk-<...>-x64-release.msi [<flag>]... [/quiet]
```

The value of `<...>` varies according to the particular release; for example, `splunk-6.3.2-aaff59bb082c-x64-release.msi`.

Command-line flags let you configure Splunk Enterprise at installation. Using command-line flags, you can specify a number of settings, including but not limited to:

- Which Windows event logs to index.
- Which Windows Registry hives to monitor.
- Which Windows Management Instrumentation (WMI) data to collect.
- The user Splunk Enterprise runs as. See [Choose the Windows user Splunk Enterprise should run as](#) for information about what type of user you should install your Splunk instance with.
- An included application configuration for Splunk to enable (such as the light forwarder.)
- Whether Splunk Enterprise should start automatically when the installation is finished.

Supported flags

The following is a list of the flags you can use when installing Splunk Enterprise for Windows from the command line.

The Splunk universal forwarder is a separate executable, with its own installation flags. See "supported command line flags" in Install a Windows universal forwarder from the command line in the *Universal Forwarder* manual.

Flag	Purpose	Default
AGREETOLICENSE=Yes No	Use this flag to agree to the EULA. You must set this flag to Yes to perform a silent installation. The flag does not work when you click the MSI to start installation.	No
INSTALLDIR="<directory_path>"	Use this flag to specify directory to install. The Splunk Enterprise installation directory is referred to as \$SPLUNK_HOME or %SPLUNK_HOME% throughout this documentation set.	C:\Program Files\Splunk
SPLUNKD_PORT=<port number>	Use this flag to specify alternate ports for splunkd and splunkweb to use. If you specify a port and that port is not available, Splunk Enterprise automatically selects the next available port.	8089
WEB_PORT=<port number>	Use this flag to specify alternate ports for splunkd and splunkweb to use. If you specify a port and that port is not available, Splunk Enterprise automatically selects the next available port.	8000
WINEVENTLOG_APP_ENABLE=1/0 WINEVENTLOG_SEC_ENABLE=1/0 WINEVENTLOG_SYS_ENABLE=1/0 WINEVENTLOG_FWD_ENABLE=1/0 WINEVENTLOG_SET_ENABLE=1/0	Use these flags to specify whether or not Splunk Enterprise should index a particular Windows event log. You can specify multiple flags: Application log Security log System log Forwarder log Setup log	0 (off)
REGISTRYCHECK_U=1/0 REGISTRYCHECK_BASELINE_U=1/0	Use these flags to specify whether or not Splunk Enterprise should index events from capture a baseline snapshot of the Windows Registry user hive (HKEY_CURRENT_USER). Note: You can set both of these at the same time.	0 (off)
REGISTRYCHECK_LM=1/0 REGISTRYCHECK_BASELINE_LM=1/0	Use these flags to specify whether or not Splunk Enterprise should index events from	0 (off)

Flag	Purpose	Default
	<p>capture a baseline snapshot of the Windows Registry machine hive (HKEY_LOCAL_MACHINE).</p> <p>Note: You can set both of these at the same time.</p>	
<p>WMICHECK_CPUTIME=1/0</p> <p>WMICHECK_LOCALDISK=1/0</p> <p>WMICHECK_FREEDISK=1/0</p> <p>WMICHECK_MEMORY=1/0</p>	<p>Use these flags to specify which popular WMI-based performance metrics Splunk should index:</p> <p>CPU usage</p> <p>Local disk usage</p> <p>Free disk space</p> <p>Memory statistics</p> <p>Note: If you need this instance of Splunk Enterprise to monitor remote Windows data, then you must also specify the LOGON_USERNAME and LOGON_PASSWORD installation flags. Splunk Enterprise cannot collect any remote data that it does not have explicit access to. Additionally, the user you specify requires specific rights, administrative privileges, and additional permissions, which you must configure before installation. Read "Choose the Windows user Splunk Enterprise should run as" in this manual for additional information about the required credentials.</p> <p>There are many more WMI-based metrics that Splunk can index. Review "Monitor WMI Data" in the Getting Data In Manual for specific information.</p>	0 (off)
<p>LOGON_USERNAME="<domain\username>"</p> <p>LOGON_PASSWORD="<pass>"</p>	<p>Use these flags to provide domain\username and password information for the Windows user that Splunk Enterprise will run as. The splunkd and splunkweb services are configured with these credentials. For the LOGON_USERNAME flag, you must specify the domain with the username in the format "domain\username." Do not use this flag to set the Splunk administrator password.</p> <p>These flags are mandatory if you want this Splunk Enterprise installation to monitor any remote data. Review "Choose the Windows user Splunk Enterprise should run as" in this manual for additional information about which credentials to use.</p>	none
<p>SPLUNK_APP="<SplunkApp>"</p>	<p>Use this flag to specify an included Splunk application configuration to enable for this installation of Splunk Enterprise. Currently supported options for <SplunkApp> are: SplunkLightForwarder and SplunkForwarder. These specify that this instance of Splunk will function as a light</p>	none

Flag	Purpose	Default
	<p>forwarder or heavy forwarder, respectively. Refer to the "About forwarding and receiving" topic in the <i>Forwarding Data</i> manual for more information.</p> <p>If you specify either the Splunk forwarder or light forwarder here, you must also specify FORWARD_SERVER="<server:port>".</p> <p>To install Splunk Enterprise with no applications at all, omit this flag.</p> <p>Note: The full version of Splunk Enterprise does not enable the universal forwarder. The universal forwarder is a separate downloadable executable, with its own installation flags.</p>	
FORWARD_SERVER="<server:port>"	Use this flag only when you also use the SPLUNK_APP flag to enable either the Splunk heavy or light forwarder. Specify the server and port of the Splunk server to which this forwarder will send data.	none
DEPLOYMENT_SERVER="<host:port>"	Use this flag to specify a deployment server for pushing configuration updates. Enter the deployment server name (hostname or IP address) and port.	none
LAUNCHSPLUNK=0/1	<p>Use this flag to specify whether or not Splunk software should start up after the installation completes, and automatically when the machine boots.</p> <p>Note: If you enable the Splunk Forwarder by using the SPLUNK_APP flag, the installer configures Splunk to start automatically, and ignores this flag.</p>	1 (on)
INSTALL_SHORTCUT=0/1	Use this flag to specify whether or not the installer should create a shortcut to Splunk on the desktop and in the Start Menu.	1 (on)
SPLUNKUSERNAME=<username>	Create a username for the Splunk administrator user. If you specify a quiet installation with the /quiet flag and do not specify this setting, then the software uses the default value of admin, but you must still specify a password with the SPLUNKPASSWORD or GENRANDOMPASSWORD flags for the installation to add the credentials successfully.	admin
SPLUNKPASSWORD=<password>	Create a password for the Splunk administrator user. The password must meet eligibility requirements. Each operating system can use a unique escape character syntax. When choosing special characters for your password, test the escaped password string before using it on a production installation. If you specify a /quiet installation, and do not define this field or the SPLUNKUSERNAME field, then the software installs without an admin user and you must create one by editing the user-seed.conf configuration file.	N/A
MINPASSWORDLEN=<positive integer>	When using the SPLUNKPASSWORD flag to set a password, you can also set password eligibility requirements for password creation and modification. The MINPASSWORDLEN flag specifies the minimum length that a password must be to meet these eligibility requirements going forward. It cannot be set to 0 or a	> 1

Flag	Purpose	Default
	negative integer. Any new password you create and any existing password you change must meet the new requirements after you set this flag.	
MINPASSWORDDIGITLEN=<integer>	When using the <code>SPLUNKPASSWORD</code> flag to set a password, you can also set password eligibility requirements for password creation and modification. The <code>MINPASSWORDDIGITLEN</code> flag specifies the minimum number of numeral (0 through 9) characters that a password must contain to meet these eligibility requirements going forward. It cannot be set to a negative integer. Any new password you create and any existing password you change must meet the new requirements after you set this flag.	0
MINPASSWORDLOWERCASELEN=<integer>	When using the <code>SPLUNKPASSWORD</code> flag to set a password, you can also set password eligibility requirements for password creation and modification. The <code>MINPASSWORDLOWERCASELEN</code> flag specifies the minimum number of lowercase ('a' through 'z') characters that a password must contain to meet these eligibility requirements going forward. It cannot be set to a negative integer. Any new password you create and any existing password you change must meet the new requirements after you set this flag.	0
MINPASSWORDUPPERCASELEN=<integer>	When using the <code>SPLUNKPASSWORD</code> flag to set a password, you can also set password eligibility requirements for password creation and modification. The <code>MINPASSWORDUPPERCASELEN</code> flag specifies the minimum number of uppercase ('A' through 'Z') characters that a password must contain to meet these eligibility requirements going forward. It cannot be set to a negative integer. Any new password you create and any existing password you change must meet the new requirements after you set this flag.	0
MINPASSWORDSPECIALCHARLEN=<integer>	When using the <code>SPLUNKPASSWORD</code> flag to set a password, you can also set password eligibility requirements for password creation and modification. The <code>MINPASSWORDSPECIALCHARLEN</code> flag specifies the minimum number of special characters that a password must contain to meet these eligibility requirements going forward. It cannot be set to a negative integer. The ':' (colon) character cannot be used as a special character. Any new password you create and any existing password you change must meet the new requirements after you set this flag.	0
GENRANDOMPASSWORD=1/0	Generate a random password for the <code>admin</code> user and write the password to the installation log file. The installer writes the credentials to <code>%TEMP%\splunk.log</code> . After the installation completes, you can use the <code>findstr</code> utility to search that file for the word "PASSWORD". After you get the credentials, delete the installation log file, as retaining the file represents a significant security risk.	0

Silent installation

To run the installation silently, add `/quiet` to the end of your installation command string. If your system has User Access Control enabled (the default on some systems), you must run the installation as Administrator. To do this:

- When opening a command prompt or PowerShell window, right click on the app icon and select "Run As Administrator".
- Use this command window to run the silent install command.

Examples

The following are some examples of using different flags.

Silently install Splunk Enterprise to run as the Local System Windows user and set the Splunk administrator credentials to "SplunkAdmin/MyNewPassword"

```
msiexec.exe /I Splunk.msi SPLUNKUSERNAME=SplunkAdmin SPLUNKPASSWORD=MyNewPassword /quiet
```

Enable the Splunk heavy forwarder and specify credentials for the user Splunk Enterprise should run as

```
msiexec.exe /i Splunk.msi SPLUNK_APP="SplunkForwarder" SPLUNKPASSWORD=MyNewPassword  
FORWARD_SERVER="<server:port>" LOGON_USERNAME="AD\splunk" LOGON_PASSWORD="splunk123"
```

Enable the Splunk heavy forwarder, generate a random password for the default Splunk administrator user, enable indexing of the Windows System event log, and run the installer in silent mode

```
msiexec.exe /i Splunk.msi SPLUNK_APP="SplunkForwarder" GENRANDOMPASSWORD=1 FORWARD_SERVER="<server:port>"  
WINEVENTLOG_SYS_ENABLE=1 /quiet
```

Where "<server:port>" are the server and port of the Splunk server to which this machine should send data.

Install Splunk Enterprise with verbose logging to C:\TEMP\SplunkInstall.log

```
msiexec.exe /I Splunk.msi /l*v C:\TEMP\SplunkInstall.log
```

See Command Line Options on Windows Dev Center for additional logging and command line options for `msiexec.exe`.

Next steps

Now that you have installed Splunk Enterprise, [learn what happens next](#).

You can also review this topic about considerations for deciding how to monitor Windows data in the Getting Data In manual.

Change the user selected during Windows installation

You can change the Windows user that Splunk Enterprise or a universal forwarder has been installed as prior to starting the software for the first time.

The user that you change to must be a user that has administrative privileges on the local machine where you installed the software. See [Choose the Windows user Splunk Enterprise should run as](#) for additional information on choosing the correct Windows user for Splunk Enterprise operations.

There are several scenarios where performing this task is helpful:

- If you selected "Domain user" during the Splunk Enterprise installation, and that user does not exist or you mistyped the information
- If you need to install a Splunk Enterprise instance as a managed system account (MSA)
- If you installed the software from a ZIP file and want to change the Windows user for the Splunk Enterprise services from the default SYSTEM user

You must perform this procedure before you start Splunk Enterprise. If Splunk Enterprise has started, then stop it, uninstall it, and reinstall it.

1. Run the Services tool. From the **Start** menu, click **Control Panel > Administrative Tools > Services**.
2. Find the `splunkd` and `splunkweb` (or `splunkforwarder` for the universal forwarder) services. These services must not be started. The Local System user owns them by default.
3. Right-click a service, and select **Properties**.
4. Click the **Log On** tab.
5. Click the **This account** button.
6. Fill in the correct domain\user name and password.
7. Click **Apply**.
8. Click **OK**.
9. (Optional) If you run Splunk Enterprise in legacy mode, repeat steps 2 through 6 for the second service.
10. Start the Splunk Enterprise services from the Service Manager or from the command-line interface.

Install Splunk Enterprise on Linux or macOS

Install on Linux

You can install Splunk Enterprise on Linux using RPM or DEB packages or a tar file, depending on the version of Linux your host runs.

To install the Splunk **universal forwarder**, see Install a *nix universal forwarder in the *Universal Forwarder* manual. The universal forwarder is a separate executable, with a different installation package and its own set of installation procedures.

Upgrading Splunk Enterprise

If you are upgrading, see How to upgrade Splunk Enterprise for instructions and migration considerations before you upgrade.

Tar file installation

What to know before installing with a tar file

Knowing the following items helps ensure a successful installation with a tar file:

- Some non-GNU versions of `tar` might not have the `-c` argument available. In this case, to install in `/opt/splunk`, either `cd` to `/opt` or place the tar file in `/opt` before you run the `tar` command. This method works for any accessible directory on your host file system.
- Splunk Enterprise does not create the `splunk` user. If you want Splunk Enterprise to run as a specific user, you must create the user manually before you install.
- Confirm that the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.

Installation procedure

1. Expand the tar file into an appropriate directory using the `tar` command:

```
tar xvzf splunk_package_name.tgz
```

The default installation directory is `splunk` in the current working directory. To install into `/opt/splunk`, use the following command:

```
tar xvzf splunk_package_name.tgz -C /opt
```

RedHat RPM installation

RPM packages are available for Red Hat, CentOS, and similar versions of Linux.

The `rpm` package does not provide any safeguards when you use it to upgrade. While you can use the `--prefix` flag to install it into a different directory, upgrade problems can occur if the directory that you specified with the flag does not match the directory where you initially installed the software.

After installation, software package validation commands (such as `rpm -Vp <rpm_file>`) might fail because of intermediate files that get deleted during the installation process. To verify your Splunk installation package, use the `splunk validate files` CLI command instead.

1. Confirm that the RPM package you want is available locally on the target machine.
2. Verify that the Splunk Enterprise user account that will run the Splunk services can read and access the file.
3. If needed, change permissions on the file.
`chmod 644 splunk_package_name.rpm`
4. Invoke the following command to install the Splunk Enterprise RPM in the default directory `/opt/splunk`.

```
rpm -i splunk_package_name.rpm
```

5. (Optional) To install Splunk in a different directory, use the `--prefix` argument.

```
rpm -i --prefix=<new_directory_prefix> splunk_package_name.rpm
```

For example, if you want to install the files into `/new_directory/splunk` use the following command:

```
rpm -i --prefix=/new_directory splunk_package_name.rpm
```

Replace an existing Splunk Enterprise installation with an RPM package

- Run `rpm` with the `--prefix` flag and reference the existing Splunk Enterprise directory.

```
rpm -i --replacepkgs --prefix=/splunkdirectory/ splunk_package_name.rpm
```

Automate RPM installation with Red Hat Linux Kickstart

- If you want to automate an RPM install with Kickstart, edit the kickstart file and add the following.

```
./splunk start --accept-license  
./splunk enable boot-start
```

The enable boot-start line is optional.

Debian .DEB installation

Prerequisites to installation

- You can install the Splunk Enterprise Debian package only into the default location, `/opt/splunk`.
- This location must be a regular directory, and cannot be a symbolic link.
- You must have access to the root user or have sudo permissions to install the package.
- The package does not create environment variables to access the Splunk Enterprise installation directory. You must set those variables on your own.

If you need to install Splunk Enterprise somewhere else, or if you use a symbolic link for `/opt/splunk`, then use a tar file to install the software.

Installation procedure

- Run the `dpkg` installer with the Splunk Enterprise Debian package name as an argument.

```
dpkg -i splunk_package_name.deb
```

Debian commands for showing installation status

Splunk package status:

```
dpkg --status splunk
```

List all packages:

```
dpkg --list
```

Information on expected default shell and caveats for Debian shells

On later versions of Debian Linux (for example, Debian Squeeze), the default non-interactive shell is the `dash` shell. Splunk Enterprise expects to run commands using the `bash` shell, and `bash` to be available from `/bin/sh`. Using the `dash` shell can result in zombie processes - processes that have completed execution, yet remain in the process table and cannot be killed or removed. If you run Debian Linux, consider changing your default shell to be `bash`.

To view an example on how to change the default shell to `bash`, see <https://unix.stackexchange.com/questions/442510/how-to-use-bash-for-sh-in-ubuntu> at StackExchange.

Next steps

Now that you have installed Splunk Enterprise:

- Start it and create administrator credentials. See [Start Splunk Enterprise for the first time](#).
- Configure it to start at boot time. See [Configure Splunk software to start at boot time](#).
- Learn what comes next. See [what happens next?](#)

Uninstall Splunk Enterprise

To learn how to uninstall Splunk Enterprise, see [Uninstall Splunk Enterprise](#).

Install on MacOS

You can install Splunk Enterprise on macOS 10.15 and 10.14 with a DMG package or a .tgz file.

Splunk Enterprise is not supported on macOS 11. A forwarder installation package is available.

Installation options

The macOS installation package comes in two forms: a DMG package and a .tgz file:

- If you require two installations in different locations on the same host, use the .tgz file. The DMG can only install Splunk Enterprise into the `/Applications/Splunk` path.

Graphical installation

1. Navigate to the folder or directory where the installer is located.

2. Double-click the DMG file.
A Finder window that contains the `splunk.pkg` opens.
3. Double-click the `Install Splunk` icon to start the installer.
4. The **Introduction** panel lists version and copyright information. Click **Continue**.
5. The **License** panel lists shows the software license agreement. Click **Continue**.
6. You will be asked to agree to the terms of the software license agreement. Click **Agree**.
7. In the **Installation Type** panel, click **Install**. This installs Splunk Enterprise in the default directory `/Applications/Splunk`.
8. You are prompted to type the password that you use to login to your computer.
9. When the installation finishes, a popup informs you that an initialization must be performed. Click **OK**.
10. A terminal window appears and you are prompted to specify a userid and password to use with Splunk Enterprise.

The password must be at least 8 characters in length. The cursor will not advance as you type. Make note of the userid and password. You will use these credentials to login Splunk Enterprise.

11. A popup appears asking what you would like to do. Click **Start and Show Splunk**. The login page for Splunk Enterprise opens in your browser window.
12. Close the **Install Splunk** window.

The installer places a shortcut on the Desktop so that you can launch Splunk Enterprise from your Desktop any time.

tar file install

Use the `.tgz` file to perform a manual installation of Splunk Enterprise. When you install Splunk Enterprise with the `.tgz` file:

- The service account is not created. If you want it to run Splunk Enterprise services with a specific user, you must create the user before starting the services.
- The default installation directory is the current working directory when you untar the `.tgz` file. The tar extraction will place all files in a `<working_directory>/Splunk` folder.

To install Splunk Enterprise on macOS:

1. Place the `<splunk_package_name.tgz>` file into a folder.
2. From the terminal, expand the tar file into the local directory using the `tar` command:

```
tar xvzf splunk_package_name.tgz
```

3. Change directory to `Splunk/bin` and start the services.

Next steps

Now that you have installed Splunk Enterprise:

- To start Splunk Enterprise services, see [Start Splunk Enterprise for the first time](#).
- To configure Splunk Enterprise services to start at boot time, see [Configure Splunk software to start at boot time in the Admin Manual](#).
- For more guidance on what to do, see [What happens next?](#).

Are you looking for the universal forwarder installation?

The universal forwarder is a separate installation package, with its own installation procedures. To install a Splunk **universal forwarder**, see [Install a *nix universal forwarder in the Universal Forwarder manual](#).

Upgrading?

If you are upgrading a Splunk Enterprise instance, see [How to upgrade Splunk Enterprise](#).

Uninstall Splunk Enterprise

If you want to remove Splunk Enterprise, see [Uninstall Splunk Enterprise](#).

Run Splunk Enterprise as a different or non-root user

On *nix based systems, you can run Splunk Enterprise as a user other than root. This is a Splunk best practice and you should configure your systems to run the software as a non-root user where possible.

If you run Splunk software as a non-root user, confirm that the software can perform the following:

- Read the files and directories that you configure it to monitor. Some log files and directories might require root or superuser access to be indexed.
- Write to the Splunk Enterprise directory and execute any scripts configured to work with your alerts or scripted input. See [Configure a script for an alert action](#) in the *Alerting Manual* or [Get data from APIs and other remote data interfaces through scripted inputs](#) in *Getting data in*.
- Bind to the network ports it is listening on. Network ports below 1024 are reserved ports that only the root user can bind to.

Because network ports below 1024 are reserved for root access only, Splunk software can only listen on port 514 (the default listening port for syslog) if it runs as root. You can, however, install another utility (such as syslog-ng) to write your syslog data to a file and have Splunk monitor that file instead.

Set up Splunk software to run as a non-root user

1. Install Splunk software as the root user, if you have root access. Otherwise, install the software into a directory that has write access for the user that you want Splunk software to run as.
2. Change the ownership of the `$(SPLUNK_HOME)` directory to the user that you want Splunk software to run as.
3. Start the Splunk software.

Example instructions on how to install Splunk software as a non-root user

In this example, `$(SPLUNK_HOME)` represents the path to the Splunk Enterprise installation directory.

1. Log into the machine that you want to install Splunk software as root.
2. Create the `splunk` user and group.

On Linux:

```
useradd splunk
groupadd splunk
```

On Mac OS: You can use the **System Preferences > Accounts** System Preferences panel to add users and groups.

3. Install the Splunk software, as described in the installation instructions for your platform. See [Installation instructions](#).

Do not start Splunk Enterprise yet.

4. Run the `chown` command to change the ownership of the `splunk` directory and everything under it to the user that you want to run the software.

```
chown -R splunk:splunk $SPLUNK_HOME
```

If the `chown` binary on your system does not support changing group ownership of files, you can use the `chgrp` command instead. See the man pages on your system for additional information on changing group ownership.

5. Become the non-root user.

```
su - <user>
```

You can also log out of the root account and log in as that user.

6. Start the Splunk software.

```
$SPLUNK_HOME/bin/splunk start
```

Use sudo to start or stop Splunk software as a different user

If you want to start Splunk Enterprise as the `splunk` user while you are logged in as a different user, you can use the `sudo` command.

```
sudo -H -u splunk $SPLUNK_HOME/bin/splunk start
sudo -H -u splunk $SPLUNK_HOME/bin/splunk stop
```

This example command assumes the following:

- That Splunk Enterprise has been installed in the default installation directory. If Splunk Enterprise is in an alternate location, update the path in the command accordingly.
- That your system has the `sudo` command available. If this is not the case, use `su` or `get` and install `sudo`.
- That you have already created the user that you want Splunk software to run as.
- That the `splunk` user has access to the `/dev/urandom` device to generate the certificates for the product.

Further reading

- To configure Splunk software to run at boot time as a non-root user, see [Enable boot-start as a non-root user in the *Admin* Manual](#).
- To learn how to install Splunk Enterprise on Windows using a user that is not an administrator, see [Choose the user Splunk Enterprise should run as](#).
- To learn how to change the Windows user that Splunk Enterprise services use, see [Change the user selected during Windows installation](#).

Install Splunk Enterprise in virtual and containerized environments

Deploy and run Splunk Enterprise inside a Docker container

Run Splunk Enterprise inside a Docker container to quickly deploy an instance and gain hands-on experience with Splunk software. The official repository containing Dockerfiles for building Splunk Enterprise and Universal Forwarder images can be found on GitHub for Splunk-Docker.

Container orchestration for Splunk Enterprise

For container orchestration, the Splunk Operator for Kubernetes on GitHub enables you to quickly and easily deploy Splunk Enterprise on your choice of private or public cloud provider. The operator simplifies scaling and management of Splunk Enterprise by automating workflows while implementing Kubernetes best practices.

See the splunk-operator documentation on GitHub for more information.

Containerized Splunk software prerequisites

The list of requirements for Docker and Splunk software is available in the Support Guidelines on the Splunk-Docker GitHub. The requirements include OS architecture, Docker version, and supported Splunk architectures.

Deploy Splunk Enterprise Docker containers

You can deploy Splunk Enterprise inside a Docker container by downloading and launching the required Splunk Enterprise image in Docker. The image is an executable package that includes everything you need to run Splunk Enterprise. For universal forwarder instructions, see

1. From a shell prompt, run the following command to download the required Splunk Enterprise image to your local Docker image library.

```
docker pull splunk/splunk:latest
```

2. Run the downloaded Docker image.

```
docker run -d -p 8000:8000 -e SPLUNK_START_ARGS='--accept-license' -e SPLUNK_PASSWORD='<password>' splunk/splunk:latest
```

- ◆ The `SPLUNK_PASSWORD='<password>'` parameter sets the login password for the admin user. There are minimum requirements when setting passwords, which can change with different versions of Splunk Enterprise. To review the minimum password requirements, see *Configure a Splunk password policy* in *Authentication.conf* in the *Securing the Splunk Platform* manual.
 - ◆ The port definition `-p <host_port>:<container_port>` will expose a port used by the containerized application to the outside network by mapping it to port on the local host. In the example above, the SplunkWeb port 8000 is mapped to the host port 8000. If a host port is already occupied by another service, you can use the `-p` parameter to re-map a port to another open port on the host, example: `-p 9000:8000`. You can later verify the ports in use by running `docker port <container_id>`
3. The output of the `docker run` command is a hash of numbers and letters that represents the container ID of your new Splunk Enterprise instance. Run the following command with the container ID to display the status of the

container.

```
docker ps -a -f id=<container_id>
```

- To verify the container ID, run `docker ps` to review the container ID, status, and port mappings of all running containers.
- Open a web browser on the host and access SplunkWeb inside the container using the address:

```
localhost:8000
```

- Log in to Splunk Enterprise inside the container using the username `admin` and the password you set when you ran the Docker image.

Administer Splunk Enterprise Docker containers

You can use the following Docker commands to manage containers.

- To see a list of example commands and environment variables for running Splunk Enterprise in a container, run:

```
docker run -it splunk/splunk help
```

- To see a list of your running containers, run:

```
docker ps
```

- To stop your Splunk Enterprise container, run:

```
docker container stop <container_id>
```

- To restart a stopped container, run:

```
docker container start <container_id>
```

- To access a running Splunk Enterprise container to perform administrative tasks, such as modifying configuration files, run:

```
docker exec -it <container_id> bash
```

To learn more about Splunk Enterprise and Docker commands, see the documentation on GitHub for Splunk-Docker.

Next steps

Now that you have Splunk Enterprise installed:

- For ideas on what to do, see [What happens next?](#).

Start using Splunk Enterprise

Start Splunk Enterprise for the first time

Before you begin using your new Splunk Enterprise upgrade or installation, take a few moments to make sure that the software and your data are secure.

As part of the initial startup process, Splunk Enterprise prompts you to create credentials for the administrator user. You can choose a username or use the default of `admin`. You can also enter a password. You must complete both steps for Splunk Enterprise to start and operate normally.

See the following topics in the *Securing Splunk* manual for more information:

- Hardening Standards
- Create secure administrator credentials

If you start Splunk Enterprise for the first time with the `--no-prompt` CLI argument, then the software does not prompt you to create the administrator credentials. If you do not create the credentials then Splunk Enterprise displays a message on login that there is no user. You must then manually create the credentials and restart Splunk Enterprise before you can log in. See "Create admin credentials manually" later in this topic for instruction on creating the credentials.

On Windows

You can start Splunk Enterprise on Windows using either the command line or the Services control panel. Using the command line offers more options.

From a command prompt or PowerShell window, run the following commands:

```
cd <Splunk Enterprise installation directory>\bin
splunk start
```

(For Windows users: in subsequent examples and information, replace `$(SPLUNK_HOME)` with `C:\Program Files\Splunk` if you have installed Splunk in the default location. You can also add `%SPLUNK_HOME%` as a system-wide environment variable by using the Advanced tab in the System Properties dialog box.)

On UNIX

1. Use the Splunk Enterprise command-line interface (CLI):

```
cd <Splunk Enterprise installation directory>/bin
./splunk start
```

2. Create the Splunk Enterprise admin username. This is the user that you log into Splunk Enterprise with, not the user that you use to log into your machine or onto `splunk.com`. You can press Enter to use the default username of `admin`.

This appears to be your first time running this version of Splunk.

Splunk software must create an administrator account during startup. Otherwise, you cannot log in. Create credentials for the administrator account.

Characters do not appear on the screen when you type in credentials.

Please enter an administrator username:

3. Create the password for the user that you just created. You use these credentials to log into Splunk Enterprise.

Password must contain at least:

* 8 total printable ASCII character(s).

Please enter a new password:

4. If the default management and Splunk Web ports are already in use (or are otherwise not available), Splunk Enterprise offers to use the next available ports. You can either accept this option or specify a port to use.
5. You can optionally set the `SPLUNK_HOME` environment variable to the Splunk Enterprise installation directory. Setting the environment variable lets you refer to the installation directory later without having to remember its exact location:

```
export SPLUNK_HOME=<Splunk Enterprise installation directory>
cd $SPLUNK_HOME/bin
./splunk start
```

6. Splunk Enterprise displays the license agreement and prompts you to accept before the startup sequence continues.

On Mac OS X

Start Splunk Enterprise from the Finder

1. Double-click the **Splunk** icon on the Desktop to launch the helper application, entitled "Splunk's Little Helper".
2. Click **OK** to allow Splunk to initialize and set up the trial license.
3. (Optional) Click **Start and Show Splunk** to start Splunk Enterprise and direct your web browser to open a page to Splunk Web.
4. (Optional) Click **Only Start Splunk** to start Splunk Enterprise, but not open Splunk Web in a browser.
5. (Optional) Click **Cancel** to quit the helper application. This does not affect the Splunk Enterprise instance itself, only the helper application.

After you make your choice, the helper application performs the requested application and terminates. You can run the helper application again to either show Splunk Web or stop Splunk Enterprise.

The helper application can also be used to stop Splunk Enterprise if it is already running.

Start Splunk Enterprise from the command line

1. On macOS, the default Splunk Enterprise installation directory is `/Applications/splunk`.

```
cd <Splunk Enterprise installation directory>/bin
./splunk start
```

2. Create the Splunk Enterprise admin username. This is the user that you log into Splunk Enterprise with, not the user that you use to log into your machine or onto splunk.com. You can press Enter to use the default username of `admin`.

This appears to be your first time running this version of Splunk.

Splunk software must create an administrator account during startup. Otherwise, you cannot log in. Create credentials for the administrator account. Characters do not appear on the screen when you type in credentials.

Please enter an administrator username:

3. Create the password for the user that you just created. You use these credentials to log into Splunk Enterprise.

Password must contain at least:

* 8 total printable ASCII character(s).

Please enter a new password:

Other start options

Accept the Splunk license automatically when starting for the first time

1. Add the `--accept-license` option to the `start` command:

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

2. Create the Splunk Enterprise admin username. This is the user that you log into Splunk Enterprise with, not the user that you use to log into your machine or onto `splunk.com`. You can press Enter to use the default username of `admin`.

This appears to be your first time running this version of Splunk.

Splunk software must create an administrator account during startup. Otherwise, you cannot log in. Create credentials for the administrator account.

Characters do not appear on the screen when you type in credentials.

Please enter an administrator username:

3. Create the password for the user that you just created. You use these credentials to log into Splunk Enterprise.

Password must contain at least:

* 8 total printable ASCII character(s).

Please enter a new password:

4. The startup sequence displays:

```
Splunk>
```

```
Checking prerequisites...
```

```
Checking http port [8000]: open
```

```
Checking mgmt port [8089]: open
```

```
Checking appserver port [127.0.0.1:8065]: open
```

```
Checking kvstore port [8191]: open
```

```
Checking configuration... Done.
```

```
Checking critical directories... Done
```

```
Checking indexes...
```

```
Validated: _audit _blocksignature _internal _introspection _thefishbucket history  
main msad msexchange perfmon sf_food_health sos sos_summary_daily summary windows wineventlog  
winevents
```

```
Done
```

```
Checking filesystem compatibility... Done
```

```
Checking conf files for problems...
```

```
Done
```

```
All preliminary checks passed.
```

```
Starting splunk server daemon (splunkd)...
```

```
Done
```

```
[ OK ]
```

```
Waiting for web server at http://127.0.0.1:8000 to be available... Done
```

If you get stuck, we're here to help.

Look for answers here: <http://docs.splunk.com>

The Splunk web interface is at <http://localhost:8000>

Start Splunk Enterprise without prompting, or by answering "yes" to any prompts

There are two other `start` options: `no-prompt` and `answer=yes`.

- If you run `$SPLUNK_HOME/bin/splunk start --no-prompt`, Splunk Enterprise proceeds with startup until it has to ask a question. Then, it displays the question and why it has to quit, and quits. In this scenario, it does not prompt for administrator credentials. You must manually create the credentials and restart before you can log in. See "Create administrator credentials manually" later in this topic for the procedure.
- If you run `$SPLUNK_HOME/bin/splunk start --answer-yes`, Splunk Enterprise proceeds with startup and automatically answers "yes" to all yes/no questions that it encounters during startup. It displays each question and answer as it continues.

If you run start Splunk Enterprise with all three options in one line, the following happens:

- The software accepts the license automatically and does not ask you to accept it.
- The software answers "yes" to any "yes/no" question.
- The software quits if it encounters a question that cannot be answered "yes" or "no".

Change where and how Splunk Enterprise starts

To learn how to change system environment variables that control how Splunk Enterprise starts and operates, see "Set or change environment variables" in the Admin manual.

Create administrator credentials manually

If you start Splunk Enterprise for the first time and use the `--no-prompt` CLI argument, Splunk Enterprise can start without an administrator user, which prevents login. To fix this problem, you must create the credentials and then restart Splunk Enterprise.

1. Stop Splunk Enterprise:


```
./splunk stop
```
2. With a text editor, create `$SPLUNK_HOME/etc/system/local/user-seed.conf`, substituting `$SPLUNK_HOME` for where you installed the software.
3. Within the file, add the following lines, substituting a password for your new password:


```
[user_info]
USERNAME = admin
PASSWORD = <your new password>
```
4. Save the file and close it.
5. Restart Splunk Enterprise by following the instructions shown earlier in this topic.

For more information on administrator credential creation, including password management for automated installations, see Create a secure administrator password in *Securing Splunk Enterprise*.

Troubleshoot Splunk Enterprise not starting the first time

If you encounter a situation where Splunk Enterprise does not start, especially after an upgrade, confirm that you have not passed any illegal arguments to the Splunk CLI as part of the start process. If you have passed illegal arguments, rerun the `splunk start` command without the arguments.

Launch Splunk Web

With a supported web browser, navigate to:

```
http://<host name or ip address>:8000
```

Use whatever host and port you chose during installation.

What happens next?

Now that you have Splunk Enterprise installed on one server, here are some links to get started:

- Learn what Splunk Enterprise is, what it does, and how it's different.
- Learn how to add your data to Splunk Enterprise. See What Splunk software can monitor in *Getting Data In*.
- Learn how to add Splunk users and roles. See About users and roles in *Securing Splunk Enterprise*.
- Learn how to estimate storage requirements for your data. See Estimate your storage requirements in *Capacity Planning*.
- Learn how to plan your Splunk Enterprise deployment, from gigabytes to terabytes per day. See the Capacity Planning Manual.
- Learn how to search, monitor, report, and more. See the Search Tutorial.
- One big way that Splunk Enterprise differs from traditional technologies is that it **classifies and interprets data at search-time**. See What is Splunk Knowledge?.

If you downloaded Splunk Enterprise packaged with an **app** (for example, Splunk + WebSphere), go to Splunk Web and select the app in Launcher to go directly to the app setup page. To see more information about the setup and deployment for a packaged app, search for the app name on Splunkbase.

Learn about accessibility to Splunk Enterprise

Splunk is dedicated to maintaining and enhancing its accessibility and usability for users of assistive technology (AT), both in accordance with Section 508 of the United States Rehabilitation Act of 1973, and in terms of best usability practices. This topic discusses how Splunk addresses accessibility within the product for users of AT.

Accessibility of Splunk Web and the CLI

The Splunk Enterprise command line interface (CLI) is fully accessible, and includes a superset of the functions available in Splunk Web. The CLI is designed for usability for all users, regardless of accessibility needs, and Splunk therefore recommends the CLI for users of AT (specifically users with low or no vision, or mobility restrictions).

Splunk also understands that use of a GUI is occasionally preferred, even for non-sighted users. As a result, Splunk Web is designed with the following accessibility features:

- Form fields and dialog boxes have on-screen indication of focus, as supported by the Web browser.
- No additional on-screen focus is implemented for links, buttons or other elements that do not have browser-implemented visual focus.
- Form fields are consistently and appropriately labeled, and ALT text describes functional elements and images.
- Splunk Web does not override user-defined style sheets.
- Data visualizations in Splunk Web have underlying data available via mouse-over or output as a data table, such that information conveyed with color is available without color.
- Most data tables implemented with HTML use headers and markup to identify data as needed.
- Data tables presented using Flash visually display headers. Underlying data output in comma separated value (CSV) format have appropriate headers to identify data.

Accessibility and real-time search

Splunk Web does not include any blinking or flashing components. However, using real-time search causes the page to update. Real-time search is easily disabled, either at the deployment or user/role level. For greatest ease and usability,

Splunk recommends the use of the CLI with real-time functionality disabled for users of AT (specifically screen readers). See *How to restrict usage of real-time search* in the *Search Manual* for details on disabling real-time search.

Keyboard navigation using Firefox and Mac OS X

To enable Tab key navigation in Firefox on Mac OS X, use system preferences instead of browser preferences. To enable keyboard navigation:

1. In the menu bar, click **[Apple icon] > System Preferences > Keyboard** to open the Keyboard preferences dialog.
2. In the Keyboard preferences dialog, click the **Keyboard Shortcuts** button at the top.
3. Near the bottom of the dialog, where it says **Full Keyboard Access**, click the **All controls** radio button.
4. Close the Keyboard preferences dialog.
5. If Firefox is already running, exit and restart the browser.

Install a Splunk Enterprise license

About Splunk Enterprise licenses

Splunk Enterprise takes in data from sources you designate and processes it so that you can analyze it. This process is called **indexing**. For information about the indexing process, see What Splunk Enterprise does with your data in *Getting Data In*.

Splunk Enterprise licenses specify how much data you can index per day.

For more information about Splunk licenses, see the following:

- How Splunk licensing works in the *Admin manual*.
- Types of Splunk Enterprise licenses in the *Admin manual*.
- More about Splunk Free in the *Admin manual*.

Install a license

After you install Splunk Enterprise, you'll have an Enterprise Trial license for 60 days. To continue using all of the features of the product after the trial, you will need to install a Splunk Enterprise license, or change to the Free license. See Switching to Free from an Enterprise Trial license in the *Admin Manual*.

Add a new license

1. Log into Splunk Web.
2. Navigate to **Settings > Licensing**.
3. Click **Add license**.
4. Either click **Choose file** and navigate to your license file and select it, or click **copy & paste the license XML directly...** and paste the text of your license file into the provided field.
5. Click **Install**. Splunk Enterprise installs your license.
6. If this is the first Enterprise license that you are installing, you will be prompted to restart Splunk Enterprise services.

If you install a Dev/Test license with an Enterprise license, the Enterprise license file will be overwritten.

Use the command line to manage licenses

For examples on using the command line to manage Splunk Enterprise licensing, see Manage licenses from the CLI in the *Admin Manual*.

Learn more about licensing

You can learn about licensing here:

- For an introduction to Splunk licensing, see How Splunk licensing works in the *Admin Manual*
- For information about allocating license volume across Splunk Enterprise instances, see Allocate license volume in the *Admin Manual*.

- To compare license types and learn which licenses can be combined, see [Types of Splunk software licenses](#) in the *Admin Manual*.
- To learn about license warnings and violations, see [About license violations](#) in the *Admin Manual*.

Upgrade or migrate Splunk Enterprise

How to upgrade Splunk Enterprise

The process of upgrading a single Splunk Enterprise instance is straightforward. In many cases, you upgrade the software by installing the latest Splunk Enterprise package over your existing installation. When you upgrade on Windows systems, the installer package detects the version that you have previously installed and offers to upgrade it for you.

The process of upgrading a distributed or clustered Splunk Enterprise deployment differs based on the type of deployment, and whether or not the instance hosts various Splunk apps and add-ons.

If the Splunk Enterprise instance or deployment that you want to upgrade has one or more premium Splunk apps installed, such as Splunk IT Service Intelligence, Enterprise Security, or User Behavior Analytics, you need to plan your upgrade sequence and target version levels to maintain version compatibility with the premium apps. The Splunk products version compatibility matrix shows which specific versions of Splunk Enterprise are compatible and supported with premium Splunk apps.

Regardless of deployment type, you must upgrade Splunk Enterprise using an operating system account with sufficient privileges to satisfy the following requirements:

- The account has administrative privileges on the machine where you perform the upgrade
- The account can write to the instance directory and all of its subdirectories.

This topic provides specific information for upgrading to version 9.0 from a previous version. If you do not want to upgrade to version 9.0, use the Version drop-down list to choose the release version that you want. Always use the upgrade instructions for the version to which you want to upgrade. Earlier or later versions of upgrade instructions can present information that appears to conflict with information for your target version.

Upgrade information for version 9.0

Read on to learn the information you need to upgrade your deployment of Splunk Enterprise, including the available upgrade paths, information that might affect you when you upgrade, and links to information on features and release notes.

Upgrade paths to version 9.0

The following table describes the upgrade paths that are available from previous versions of Splunk Enterprise.

Find the version you currently use in the first column and read across to determine the upgrade path for that version. If your version does not appear in the first column, then there is no supported upgrade path to the latest version. You must first upgrade to a version that is in this list. After selecting a hyperlink in the following table, choose your specific product version from the **Version** drop-down list.

Your current version	First upgrade to latest available	Then upgrade to latest available	README link	Rel. Notes link
7.0.x, 7.1.x, 7.2.x, 7.3.x	8.0.x or 8.1.x	9.0.x	8.0 README, 8.1 README	8.0 Rel. Notes, 8.1 Rel. Notes

Your current version	First upgrade to latest available	Then upgrade to latest available	README link	Rel. Notes link
8.0.x	8.1.x or 8.2.x	9.0.x	8.1 README, 8.2 README	8.1 Rel. Notes, 8.2 Rel. Notes
8.1.x, 8.2.x	9.0.x	N/A	9.0 README	9.0 Rel. Notes

Splunk Enterprise upgrade process

The upgrade process for Splunk Enterprise consists of three phases:

- Phase 1: Identify, back up, and verify that components work as you expect
- Phase 2: Install updated Splunk Enterprise components
- Phase 3: Confirm everything works after the upgrade

This process applies to upgrades of all Splunk Enterprise deployments. Depending on the kind of deployment you have, some steps might differ from what this page shows.

Phase 1: Identify, back up, and verify that components work as you expect

Use the following steps to prepare a Splunk Enterprise upgrade. Specific steps might differ based on the size and kind of deployment and whether or not your deployment runs a premium Splunk app.

1. Identify all of the components in your deployment. This determines the upgrade procedures that you must follow during the upgrade phase:
 - ◆ Identify all single-instance components.
 - ◆ Identify all distributed components that are not in a cluster.
 - ◆ Identify all clustered components.
2. Back up your existing deployment, including configurations and data. For more information about backing up your Splunk Enterprise deployment, see *Back up configuration information* in the *Admin Manual* and *Back up indexed data* in the *Managing Indexers and Clusters of Indexers* manual.
3. Validate your backups and confirm that they can be restored.
4. Where applicable, use the Monitoring Console to take a snapshot of the health of your existing Splunk Enterprise deployment.
5. If you run a clustered Splunk Enterprise environment, use the Monitoring Console to confirm that the cluster is healthy.
6. If you run a Splunk Enterprise license manager instance, confirm that it is healthy, that all indexers successfully connect to it, and that all license keys either are available for entry or exist on backup media.
7. If you run a deployer on a search head cluster, confirm that it is healthy and can push configuration bundles to all SHC peers without problems.
8. If you run a deployment server machine, confirm that it is healthy, that configurations reload successfully, and that all forwarders can connect to it.
9. Review the forwarder-indexer compatibility matrix in *Compatibility between forwarders and indexers* in *Splunk Products Version Compatibility Matrix* to confirm that all forwarders in your deployment work with the version of indexer to which you plan to upgrade. Older versions of forwarder might not be compatible due to various security cipher changes.
10. For distributed deployments of any kind, confirm that all machines in the indexing tier satisfy the following conditions:
 - ◆ They have sufficient disk space available for installation of the updated software
 - ◆ They run basic searches without problems
 - ◆ They do not run their own saved searches

11. On distributed deployments of any kind, confirm that all machines in the search tier satisfy the following conditions:
 - ◆ The version of Splunk Enterprise that you want to upgrade can run your apps, add-ons, and dashboards
 - ◆ You have all security keys, configurations, and credentials available for possible reentry
 - ◆ Searches do not fail because of incorrect authentication credentials

Phase 2: Install updated Splunk Enterprise components

After you complete the pre-upgrade steps in Phase 1, you can begin upgrading individual Splunk Enterprise components. Depending on your deployment type, you might need to perform additional steps.

1. Read [About upgrading to 9.0: READ THIS FIRST](#) completely prior to starting an upgrade.
2. If you run premium Splunk apps, see the Splunk Products version compatibility matrix to determine the versions that your apps support.
3. Upgrade the Splunk Enterprise components in your deployment, based on the deployment architecture you identified in Phase 1:
 - ◆ For distributed environments that do not have clusters, follow the instructions in [How to upgrade a distributed Splunk Enterprise environment](#).
 - ◆ For clustered environments, see one of the following topics:
 - ◇ To upgrade an indexer cluster, see Upgrade an indexer cluster in the *Managing Indexers and Clusters of Indexers* manual.
 - ◇ To upgrade a search head cluster, see Upgrade a search head cluster in the *Distributed Search* manual.
 - ◆ For single instance deployments, follow the upgrade instructions for your operating system type:
 - ◇ [Upgrade to 9.0 on *nix](#)
 - ◇ [Upgrade to 9.0 on Windows](#)
4. During the upgrade, depending on the component that you upgrade, you might need to perform validation steps to ensure the upgrade is successful.
 - ◆ On a cluster manager node, you might need to run validation searches or use operating system tools to determine cluster manager health and readiness before you proceed to the next upgrade phase.
 - ◆ On forwarders, you can use Monitoring Console to determine that data ingestion levels remain at pre-upgrade rates as forwarders come back online.
 - ◆ On standalone indexers, you can run searches to determine that data ingestion and search participation occur normally.
 - ◆ On clustered indexers, you can use Monitoring Console to determine that indexers come back online and appear as normal in the Clustering Status page.

Phase 3: Verify everything works after the upgrade

After you complete the upgrade of Splunk Enterprise components, follow these high-level steps to confirm that your upgrade was successful. As with the other phases, specific steps might differ based on the number and kind of Splunk Enterprise components that you have in your deployment.

1. Confirm that your Splunk apps and add-ons work like they did before the upgrade.
2. If you have a distributed deployment, use Monitoring Console to verify all Splunk Enterprise components.
 - ◆ Review resource utilization for all components and compare to what you benchmarked prior to the upgrade.
 - ◆ Confirm all components are available.
3. If you have a distributed deployment, confirm that the license manager machine works properly and all indexers connect to it, like they did before the upgrade.
4. If you have a clustered deployment, confirm that the cluster manager operates normally and that cluster peers are connecting properly.

5. If you have a distributed deployment, confirm that the search tier operates normally and that search and indexers communicate without problems
6. If you have a search head cluster, use the Monitoring Console to verify search head cluster state and individual cluster peer nodes.
7. If you have an indexer cluster, confirm that all indexer cluster nodes reestablish communications with the cluster manager.

Optional upgrade activities

The following section describes optional steps that you can perform after an upgrade.

Review and configure the `tsidxWritingLevel`

Splunk Enterprise 7.2 introduced a new file format and optimizations for `tsidx` files that resulted in improved search performance through decreased I/O, lowered storage usage, and improved utilization of SmartStore caches. These optimizations are encapsulated in levels, with new levels added in higher releases of Splunk Enterprise. Changing the default `tsidxWritingLevel` changes the optimizations used by both the index `tsidx` files and data model accelerations.

To determine whether the `tsidx` level available has changed since your last upgrade, and what value to set the `tsidxWritingLevel` to, see The `tsidx` writing level in the *Managing Indexers and Clusters of Indexers* manual.

About upgrading to 9.0 READ THIS FIRST

Here are some changes in behavior to note for Splunk Enterprise and the Splunk Universal Forwarder when you upgrade to version 9.0.

Begin by selecting the version you want to upgrade Splunk Enterprise to, and review this page. Lower or higher versions of this topic can present information that appears to conflict with information for your target version.

Splunk App and Add-on Compatibility

Not all Splunk apps and add-ons are compatible with Splunk Enterprise version 9.0.

- See the Splunk products version compatibility matrix for information about which versions of Splunk IT Service Intelligence and Splunk Enterprise Security are compatible with this version of Splunk Enterprise.
- You can visit Splunkbase to confirm that your apps and add-ons are compatible with this version.

If your app or add-on is not compatible with version 9.0, consider delaying your upgrade until a compatible version is available.

Key points for upgrading to version 9.0

The following is a list of important elements that you must consider before upgrading Splunk Enterprise and its components. The sections that follow provide supporting details for these key points. Read through all sections in the topic before you begin your upgrade activities.

- Splunk Enterprise 9.0 has some security changes and enhancements that you can enable after you upgrade. In a later release, Splunk might enable these enhancements for you. See Security updates in the *Securing Splunk*

Enterprise Manual to learn about the changes and how they might affect you.

- Splunk Enterprise 9.0 introduces fixes for a critical vulnerability in the deployment server. Depending on the set up of your Splunk Enterprise deployment, you might need to isolate your deployment server from other Splunk Enterprise components and perform upgrades of deployment clients that are currently lower than version 7.0.0 of Splunk Enterprise or the universal forwarder to at least version 7.0.0. You do not need to upgrade your deployment clients to version 9.0.0 at this time. See [Changes that can potentially break Splunk Enterprise installations](#) later in this topic for specific details.
- Migrate your App Key Value Store storage engine from the Memory Mapped (MMAP) storage engine to the WiredTiger storage engine during or before upgrade to 9.0.0, and update your MongoDB version from 3.6 to 4.2 during or immediately after upgrade. These updates are required in Splunk Enterprise 9.0. See *Migrate the KV store storage engine* in the *Admin* manual to plan your migration.
- Adjust your scripts and templates to use Python 3-compatible syntax before you upgrade. Splunk Enterprise 9.0 removes support for Python 2. See *Python 3 migration* for more information.
- Use the Splunk Products version compatibility matrix to ensure that any premium Splunk apps and add-ons you run are compatible with version 9.0. If they are not, do not upgrade until compatible versions become available.
- Data Fabric Search (DFS) and all associated server-side components are removed. Upgrading to Splunk Enterprise 9.0 removes DFS functionality.
- Upgrading Splunk Enterprise directly to version 9.0 is only supported from versions 8.1.x and higher. Upgrading a universal forwarder directly to version 9.0 is supported from versions 8.1.x and higher.
- To upgrade search head or indexer clusters, see [Follow specific instructions to upgrade clusters](#) later in this topic.
- Back up your App Key Value Store (KV Store) databases prior to starting an upgrade. If you run version 7.1 and lower of Splunk Enterprise, you must stop Splunk Enterprise instances first.
- If you run Linux machines that use the second extended (ext2) file system, upgrade that file system to third extended (ext3) prior to starting an upgrade.

Changes that can potentially break Splunk Enterprise installations

There is no supported method for rolling back an installation to a prior release. Follow the guidance for these critical items to avoid breaking your existing installation during an upgrade.

Change to the `search_listener` parameter for the `search/jobs` endpoint

The `search_listener` request parameter for the Splunk REST API `search/jobs` endpoint is disabled. This change takes effect in Splunk Enterprise versions 8.1.13, 8.2.10, and 9.0.4.

Splunk Enterprise 9.0.2 changes the default value for `deployer_lookups_push_mode` in `app.conf`

In 9.0, a change was made to the behavior of the default `preserve_lookups` value for the `deployer_lookups_push_mode` setting in `app.conf`. This change fixed a behavior issue that caused `preserve_lookups` to not conform to its documented and intended behavior. Rather, prior to 9.0.0, the `preserve_lookups` value conformed instead to the documented behavior of the `always_preserve` value.

However, although the 9.0.0 change fixed the behavior of the `preserve_lookups` value, the change also led to performance degradation when using that value because of additional processing needed to attain the intended result. In addition, the fix changed the default behavior of the `deployer_lookups_push_mode` setting, which introduced an additional problem, since some users had come to expect and rely on the pre-9.0.0 behavior of the default `preserve_lookups` value, buggy though it was,

To counteract the resultant performance degradation and change to expected default behavior introduced by the 9.0.0 change to the default value, in 9.0.2, the default value for the `deployer_lookups_push_mode` setting was changed to `always_preserve`. This change to the default causes the default behavior of the setting to conform to the unfixed behavior

of the `preserve_lookups` value prior to the change in 9.0.0,

See Preserve lookup files across app upgrades in the Distributed Search manual.

Splunk Enterprise 9.0 fixes a critical vulnerability in deployment server but might introduce problems for older deployment clients

If you run a deployment server, upgrade that server to version 9.0 of Splunk Enterprise as soon as possible. Before the upgrade, carefully review your deployment server setup and the current versions of the deployment clients in your Splunk Enterprise network. Depending on the setup of your deployment server and whether that component shares a computer with other Splunk Enterprise components, you might need to do the following to ensure your deployment server and clients communicate without problems:

- Isolate deployment server from other components on a machine. Isolating your deployment server means you only have to upgrade that component. The sole exception for isolation is if you run a deployment server and a license manager on the same machine.
- Confirm that all deployment clients in your network run version 7.0.0 or higher of Splunk Enterprise or the universal forwarder. You don't have to upgrade deployment clients to version 9.0.0, but they must be at version 7.0.0 or higher to communicate with version 9.0.0 deployment servers.

See the following topics for additional information:

- SVD-2022-0608 for more about the vulnerability.
- Security updates in the *Securing Splunk Enterprise Manual* for more about the security updates that come with Splunk Enterprise 9.0.
- Client version compatibility in the *Updating Splunk Enterprise Instances Manual* for more about which versions of deployment client that the version 9.0.0 deployment server supports.

REST endpoint access authorization will be more secure in environments with distributed search

Applicable components: Splunk Enterprise

Applicable OSEs: all

Version introduced: 9.0

Beginning with Splunk Enterprise version 9.0, REST endpoints that distributed search nodes use to dispatch search jobs to indexers will use their own user-level tokens during that dispatching, rather than using a system-level token. You can override that default behavior by changing a setting in the `limits.conf` configuration file, but Splunk might remove that capability in a future release.

For more information about this change, see Security Updates in the *Securing the Splunk Platform Manual*.

Communication between deployment clients and deployment servers will be more secure after a configuration change

Applicable components: Splunk Enterprise

Applicable OSEs: all

Version introduced: 9.0 Beginning in Splunk Enterprise version 9.0, the communication between deployment clients and deployment servers will receive security improvements.

Channel subscription logic, which is the logic that deployment clients use to establish a subscription to deployment servers, will change. Clients must subscribe to deployment server channels using an exact text match for the channel name. They must authenticate using a `pass4symmKey` that is identical on both the client and deployment server to

download content from those servers. Clients will verify that they are actually communicating with a deployment server before attempting to download client bundles.

This change requires you to make a configuration update to your deployment clients and servers. For more information about the change, see Security Updates in the *Securing the Splunk Platform Manual*.

New capability requirements for some risky search commands can potentially break existing user- and app-related searches

Applicable components: Splunk Enterprise

Applicable OSeS: all

Version introduced: 9.0

Beginning with version 9.0 of Splunk Enterprise, Splunk has created new capabilities that a user must possess through a role to run various risky search commands.

The following table lists the commands and the new capabilities that users need to run the commands:

Search command	New required capability
sendalert	run_sendalert
dump	run_dump
Any custom command	run_custom_command

By default, the user and power roles receive these new capabilities. If a user that does not hold these roles runs these commands, that user will no longer be able to do so after an upgrade unless you either add the capabilities to at least one role that the user holds, or inherit either the user or power roles to a role that the user holds.

For more information about the Splunk platform approach towards managing risky commands see Safeguards for risky commands in the *Securing the Splunk Platform manual*.

TLS certificate validation is now available for all Splunk Enterprise instance types

Applicable components: Splunk Enterprise

Applicable OSeS: all

Version introduced: 9.0

If you use TLS to secure your Splunk Enterprise deployment, you can now enable TLS host name validation across the instances in that deployment. A new setting in the server.conf configuration file lets you control whether or not this validation happens on machines where you require certificates for secure communications. Some Splunk Enterprise instance types use certificates automatically for this type of connection. This version of Splunk Enterprise will provide notice when it cannot validate TLS certificates. Later versions of Splunk Enterprise might begin enforcing this type of certificate validation, which can cause connectivity problems between indexers, search heads, deployment servers, cluster nodes, and App Key Value Store.

For more information about this change, including how to turn on enforcement of certificate validation, see Security Updates in the *Securing the Splunk Platform Manual*.

Transport Layer Security (TLS) certificate validation is available for Python 3 modules

Applicable components: Splunk Enterprise

Applicable OSeS: all

Version introduced: 9.0

Splunk Inc. ships Python 3 as part of the Splunk Enterprise software package. Some Splunk-provided Python 3 modules, including but not limited to the httplib Python module, help developers and users alike establish secure connections to other Splunk Enterprise instances and APIs. The modules make these secure connections as part of their normal operations without validating the TLS certificates that the modules use to make those secure connections. Beginning with version 9.0 of Splunk Enterprise, you can use a setting in the server.conf configuration file to force TLS certificate validation checks. Later versions of Splunk Enterprise might enable these checks by default. Apps, APIs, and instances that use Python 3 modules which do not check TLS certificates might break when TLS certificate validation is on.

For more information about this change, see Security Updates in the *Securing the Splunk Platform Manual*.

Splunk Enterprise indexers and App Key Value Store nodes can now verify TLS certificates that they use to connect to each other

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

Splunk Enterprise indexers and App Key Value Store nodes use TLS to communicate securely with one another. Beginning with version 9.0 of Splunk Enterprise, you can use a setting in the server.conf configuration file to force the nodes to verify the certificates that they use to make secure connections. Later versions of Splunk Enterprise might enable these checks by default. When this happens, indexers and App Key Value Store nodes will not be able to communicate with one another any longer over TLS.

For more information about this change, see Security Updates in the *Securing the Splunk Platform Manual*.

Splunk Enterprise search peers will can now verify TLS certificates that they use to connect to each other

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

Splunk Enterprise search peers use TLS to communicate securely with one another. Beginning with version 9.0 of Splunk Enterprise, you can use a setting in the server.conf configuration file to force the nodes to verify the certificates that they use to make secure connections. Later versions of Splunk Enterprise might enable these checks by default. When this happens, search peers will not be able to communicate with one another any longer over TLS.

For more information about this change, see Security Updates in the *Securing the Splunk Platform Manual*.

The Splunk Command Line Interface can now verify TLS certificates that it uses to connect to external Splunk Enterprise nodes

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

The Splunk Enterprise CLI uses TLS to communicate securely with Splunk Enterprise instances on other machines. Beginning with version 9.0 of Splunk Enterprise, you can use an argument in the CLI to force it to verify the certificates that it use to make secure connections. Later versions of Splunk Enterprise will enable these checks in the CLI by default. When this happens, the CLI will not be able to connect to other Splunk Enterprise nodes using TLS if it cannot validate the certificates.

For more information about this change, visit our Security Updates page in the *Securing the Splunk Platform* Manual.

The indexer cluster slave-apps directory is renamed to peer-apps

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

During the upgrade, the indexer cluster `slave-apps` directory is automatically renamed to `peer-apps`. To accommodate this renaming, you must manually change any external hardcoded references to `slave-apps`, such as external scripts, SSL certificates, and so on, to `peer-apps`. See Update common peer configurations and apps in *Managing Indexers and Clusters of Indexers*.

Plan your upgrade to work with the Python 3 migration

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

Beginning with version 9.0, Splunk Enterprise uses the Python 3 interpreter only. See Python 3 migration for more information.

Splunk Enterprise and Universal Forwarders must be version 8.1 or higher to upgrade to version 9.0

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 9.0

Splunk supports upgrades of Splunk Enterprise and the universal forwarder to version 9.0 from version 8.1 and higher only. If you run a version of Splunk Enterprise or universal forwarder that is lower than 8.1, you must upgrade to version 8.1 or higher as an intermediate step before attempting the upgrade to version 9.0. See [Upgrade paths to version 9.0](#) for more information.

Follow specific instructions to upgrade clusters

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 7.3

To upgrade indexer or search head clusters, follow the upgrade procedure for the type of deployment you have.

- If your deployment has indexer clusters, follow the index cluster upgrade instructions.
- If your deployment has search head clusters, follow the search head cluster upgrade instructions.

On Splunk Enterprise indexer cluster manager nodes, enable maintenance mode before you upgrade

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.1

Because of changes to bucket data with the release of the latest version of Metrics Store, indexers that run versions of Splunk Enterprise lower than 8.1 cannot handle bucket replications from versions that run 8.1 and higher. Before you

upgrade a cluster manager node, ensure it is in maintenance mode first. This is a standard part of the indexer cluster update process.

See Use maintenance mode in *Managing Indexers and Clusters of Indexers*.

Back up App Key Value Store prior to starting an upgrade

Applicable components: Splunk Enterprise

Applicable OSEs: all

Version introduced: 7.3

Back up the app key value store (KV store) before any maintenance like an upgrade.

If you are upgrading from version 7.1 or lower, you must stop Splunk Enterprise before you can back up the KV Store databases. Confirm that you have accounted for this downtime in your upgrade planning. See Back up and restore KV store for more information.

The Splunk Enterprise credential creation scheme might affect scripted upgrades

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSEs: all

Version introduced: 7.2

Splunk Enterprise 7.1 introduced an updated authentication scheme for users that requires that you create a password for the admin account. In version 7.2, the scheme was extended to let you customize administrator credential creation for Splunk Enterprise instances.

This scheme includes additional settings and configuration options, which can affect how you upgrade if you use scripts to automate the upgrade process. You might need to change your upgrade scripts before performing scripted upgrades. Specifically, confirm that you do not pass any illegal arguments to the Splunk CLI for starting or restarting Splunk Enterprise during the upgrade, as this could result in a situation where Splunk Enterprise does not start after the upgrade has finished.

Occurrences that appear to be problems but are not

You might see things happen during or immediately after an upgrade that appear to indicate that the upgrade is not working. In nearly all cases, the occurrences that happen here can be expected.

If the following things occur during the upgrade, let the upgrade continue and do not interrupt it. If they occur immediately afterward, then perform benchmark tests on the deployment and compare to any benchmarks that you set up as part of the first phase of upgrading. Consider involving Splunk Support only if those benchmarks differ by a significant margin.

- On indexers, memory and CPU usage increases due to the following:
 - ◆ New data ingestion pipelines
- Permissions on the Splunk Enterprise introspection directory might change. Confirm that the user that runs Splunk Enterprise has write permission to the `$(SPLUNK_HOME)/var/log/introspection` directory.
- Deployment servers might push updates to all deployment clients due to app bundle hash recalculations.

CPU and memory usage on indexers increases

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 7.3

By default, new data ingestion pipelines consume more resources when they are enabled.

Permissions on introspection directory might change

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 7.3

Confirm that the user who runs Splunk Enterprise has write permission to the `$SPLUNK_HOME/var/log/introspection` directory.

Deployment servers might push updates to all deployment clients due to app bundle hash recalculations.

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 7.3

Considerations for changed or removed features

The following major features have been changed or removed from this version. If you use features that have been removed, and have not yet migrated off them, consider delaying your upgrade until you have.

Index buckets use zstd compression by default

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

The `indexes.conf` setting `journalCompression` defaults to using zstd compression, beginning with this release.

The default `tsidxWritingLevel` for indexes has changed

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

The default `tsidxWritingLevel` in `indexes.conf` has changed from 2 to 3, beginning with this release.

Splunk Enterprise 7.2 introduced a new file format and optimizations for `tsidx` files that resulted in improved search performance through decreased I/O, lowered storage usage, and improved utilization of SmartStore caches. These optimizations are encapsulated in levels, with new levels added in higher releases of Splunk Enterprise. Changing the default `tsidxWritingLevel` changes the optimizations used by both the index `tsidx` files and data model accelerations. See The `tsidx` writing level in the *Managing Indexers and Clusters of Indexers* manual.

The IOWait status check is delayed on startup

Applicable components: Splunk Enterprise

Applicable OSES: Linux

Version introduced: 9.0

The IOWait Health Report incorporates a 120 second delay after the Splunk Enterprise service starts to prevent false alerts.

The Bucket Health alert threshold is more aggressive

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

The threshold values for the Health Report "Bucket" feature are set to alert more aggressively when an indexer rolls a large number of small buckets.

The web.conf setting 'simplexml_dashboard_create_version' is deprecated

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 9.0

The `simplexml_dashboard_create_version` setting in the web.conf configuration file has been deprecated. Changing the default value might introduce security risks. Do not change the value without consulting Splunk Support.

The use of SSL compression is replaced with HTTP compression except when forwarding

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.2

In an effort to improve scalability and efficiency, the use of SSL compression is no longer enabled by default for SSL communications between Splunk Enterprise instances. The SSL compression is replaced with HTTP compression by default, except when forwarding data over SSL. This change does not prevent the use of SSL compression when negotiating communications over SSL with earlier Splunk Enterprise instances. For example, an earlier client negotiating with an 8.2 instance will still use SSL compression, and an 8.2 client negotiating with earlier servers will use HTTP compression.

This change effects the existing settings:

Setting	New default with 8.2	Prior default
server.conf useHTTPClientCompression	true	false
server.conf useClientSSLCompression	false	true
	true	Referenced the server.conf useClientSSLCompression setting (default: true)

Setting	New default with 8.2	Prior default
outputs.conf useClientSSLCompression		

If you have modified the `useHTTPClientCompression` or `useClientSSLCompression` settings in the past, check your environment to determine if the new defaults will impact your compression settings after an upgrade. Use `btool` to verify how compression is set in your environment based upon the settings above, and where those changes are made.

If you have explicitly enabled the `server.conf` setting `useClientSSLCompression` using a custom app or a local `.conf` file, after the upgrade both the SSL and HTTP compression will be enabled simultaneously. After the upgrade is complete, verify the `useClientSSLCompression` setting, and make sure the SSL compression is disabled.

The splunkd.log events now include thread names

Applicable components: Splunk Enterprise

Applicable OSes: all

Version introduced: 8.2

The default logging level for `splunkd.log` events now includes thread names to improve troubleshooting.

An additional option is available when performing a searchable rolling restart for multisite clusters

Applicable components: Splunk Enterprise

Applicable OSes: all

Version introduced: 8.2

The new `server.conf` setting `searchable_rolling_site_down_policy` improves searchable rolling restart performance when the prerequisites are met.

See How the manager determines the number of multisite peers to restart in each round in the *Managing Indexers and Clusters of Indexers* manual.

A Splunk Enterprise or universal forwarder installation will no longer create an inputs.conf with the hostname

Applicable components: Splunk Enterprise

Applicable OSes: all

Version introduced: 8.1

There's a change in the default installation behavior for Splunk Enterprise and universal forwarder instances. When installing earlier releases, the installation process would determine the hostname and set it in a newly created `$(SPLUNK_HOME)/etc/system/local/inputs.conf` file. On 8.1 and higher releases, `splunkd` no longer creates the `inputs.conf` file during installation. The service checks and sets the hostname as part of the `inputs.conf` setting `host = $decideOnStartup` when the service starts.

Some logging categories for the authentication system have changed

Applicable components: Splunk Enterprise

Applicable OSEs: all

Version introduced: 8.1

As part of overall improvements to support of the Security Assertion Markup Language (SAML) authentication scheme, various logging channels for the Splunk platform authentication system have changed. The logging category names that previously started with `AuthenticationManager` now begin with `AuthenticationProvider`. Following an upgrade, review the `log.cfg` in `$$SPLUNK_HOME/etc/` to confirm that the names have been changed.

Splunk Enterprise license enforcement change

Applicable components: Splunk Enterprise

Applicable OSEs: all

Version introduced: 8.1

For license stack volumes of less than 100GB of data per day, Splunk Enterprise will disable search when license limits are violated, after 45 warnings within a 60-day rolling window. For more information, see the License Enforcement FAQ on splunk.com.

The default logging level for audit logs has changed

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSEs: all

Version introduced: 8.1

The default logging level for auditing of certain administrative events, information that goes to the `_audit` index, has been lowered from the `INFO` level to the `DEBUG` level. Additionally, you now gain the ability to control the logging level that auditing events happens.

While this is a net positive for most customers due to the lower default verbosity of logging, it also means that some audit events, such as capability checks, are no longer logged by default due to the lower log level. To restore the old behavior, you can configure the Splunk platform logging utility by editing the `$$SPLUNK_HOME/etc/log-local.cfg` file and raising the `category.AuditLogger` entry from `DEBUG` back to `INFO`.

Splunk data collection practices have changed

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSEs: all

Version introduced: 8.0

In an effort to provide better support to its customers and improve its products and services, Splunk has changed its data collection practices. After you upgrade, Splunk automatically opts you in to sharing telemetry data. To learn what data Splunk collects, see *What data Splunk collects* in the *Admin Manual*.

The change happens when you first log into an instance with an administrator user after the upgrade. You receive a pop-up message that indicates the policy change, and the instance overwrites any existing data-sharing configurations after you acknowledge the pop-up.

You can still opt out of sharing telemetry data at any time. To learn how, see [How to opt out](#) in the "Share data in Splunk Enterprise" topic of the *Admin Manual*.

The "access controls" Splunk Web menu options have changed

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.0

The "Access controls" menu item under the "Users and Authentication" header in the Settings menu in Splunk Web has been replaced with individual links to user, role, password, and authentication scheme management.

The default memory resource allocation for workload categories has changed

Applicable components: Splunk Enterprise

Applicable OSES: Linux

Version introduced: 8.0

On upgrade to version 8.x, the default memory resource allocation for the ingest workload category changes from 20% to 100%. This might cause an increase in memory usage in the ingest category after upgrade. The default memory resource allocation for the search and misc. categories remains the same, at 70% and 10%, respectively.

For more information on workload categories, see [Configure workload categories](#) in the *Workload Management* manual.

The "failure to localize search" Splunk Web error message has been replaced with a "partial results" message

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.0

When you upgrade to version 8.x and run a search that encounters a localization error, the message that appears in Splunk Web has changed. Instead of displaying a "failure to localize" message, Splunk Web notifies users that the search process might have returned partial results.

Roles with the "install_apps" capability also need the "list_settings" capability to access certain REST endpoints

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.0

After you upgrade to version 8.x of Splunk Enterprise, any roles that hold the `install_apps` capability must also hold the `list_settings` capability for role users to be able to manage app installations through REST (using the `manager/appinstall/<app>` endpoint, for example).

Splunk Enterprise forwarders at version 8.x cannot forward metrics data to indexers that run lower than version 8.x

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0

There is no support for the forwarding of metrics data from forwarders that run version 8.x to indexers that run version 7.3 or lower. If you need to forward metrics data from a version 8.x forwarder, confirm that all indexers that receive this kind of data also run version 8.x.

Upgraded search head nodes will complain of missing metrics indexes in lower-version indexer nodes

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0

If you upgrade your distributed search environment to version 8.x of Splunk Enterprise but keep your indexer nodes at a version below 8.0, your search head nodes will generate alerts about a missing `_metrics` index when they try to send `metrics.log` events to those indexers. You cannot work around this problem by adding an `_metrics` index to the indexer cluster, as the format of the metrics data is different on indexers that run a lower version. Upgrade all Splunk Enterprise components to the latest version where possible.

After an upgrade, configure `limits.conf` settings on all search head cluster nodes that handle metrics data

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0

After you upgrade your search head and indexer clusters to version 8.x of Splunk Enterprise, edit `limits.conf` on each search head cluster and set the `always_use_single_value_output` setting under the `[mcollect]` stanza to `false`. This lets these nodes use the "multiple measures per metric data point" schema when you convert logs to metrics with the `mcollect` command or use metrics rollups. This new schema increases your data storage capacity and improves metrics search performance.

In metric indexes only, by default the Splunk software removes rawdata journal files from buckets when they roll from hot to warm

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.1

Rawdata journal files are not used by metric searches. However, after the optimizations introduced in 8.0, rawdata journal files take up a large amount of the volume of a metrics index bucket. In version 8.1 and later, the Splunk software removes these files from metric index buckets when they roll from hot to warm by default. The software removes the file from `/rawdata` directory in the bucket and replaces it with a dummy journal file that has no usable data other than the journal header. This reduces the amount of space that metric indexes take up on disk.

This does not apply to metric indexes that have replication enabled (`repFactor = auto` in `limits.conf`) for indexer clustering. For more information, or to learn how to disable rawdata journal removal for a specific index, see the description of the `metric.stubOutRawdataJournal` setting in `limits.conf.spec`.

Metric data points containing metric names that are empty or composed entirely of white spaces cannot be indexed or searched upon

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.1

In version 8.1 and later, Splunk Enterprise cannot index metric data points containing metric names that are empty or composed entirely of white spaces. This can happen with ingested metrics data as well as event data that is converted to metrics data through the log-to-metrics sourcetype or some other method.

After the upgrade, Splunk Enterprise searches cannot return metric data points containing empty or white-spaced metric names, if those metric data points were indexed in a previous version of Splunk.

StatsD metric data inputs now produce single-measurement metric data points by default

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0.3

For ease of use, by default Splunk Enterprise version 8.0.3 and later converts StatsD metric data into single-measurement metric data points that have one key-value pair for the metric name and one key-value pair for the measurement value. If, prior to 8.0.3, you used StatsD inputs that converted their data to metric data points that could carry multiple measurements, you need to add `STATSD_EMIT_SINGLE_MEASUREMENT_FORMAT=false` to a stanza for the metric source type in `props.conf`.

Streaming metric alerts are not available on indexer clusters that run version 7.2 and lower

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0

Splunk Enterprise version 8.0 introduces streaming metric alerts, which are evaluated on a continuous basis, and which can reduce the load on your system by enabling similar alerts to share the same search process. There is no support for streaming metric alerts on indexer clusters that run version 7.2 or lower. If you have an indexer cluster and a search head cluster, and you upgrade the search head cluster to version 8.x, you lose the ability to trigger streaming metric alerts for that search head cluster. If you need to retain this ability, upgrade the indexer cluster to version 8.x first.

You can no longer use wildcards in real-time searches that use the mstats command

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.0

Beginning with Splunk Enterprise version 8.0, you can no longer use wildcards in real-time searches that contain the `mstats` command. If you have these kinds of searches, change them so that they no longer use wildcards of any kind, either in aggregation fields or metric names, prior to upgrading.

Metrics indexing and search is now case-sensitive

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0

The indexing and search schema for metrics data is case-sensitive as of Splunk Enterprise 8.0. Instances of Splunk Enterprise that are lower than 8.0 might encounter problems retrieving results when searching instances that run version 8.0 or higher.

Older, deprecated configuration file settings for knowledge bundle replication have been removed

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.0

Many older, previously deprecated settings under the `[replicationSettings]` stanza in `distsearch.conf` have been removed. If your configuration uses these settings, consider changing them before you start an upgrade.

Configuration file settings for specifying the mounted bundle option for knowledge bundle replication have changed

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 8.0

The `distsearch.conf` setting for specifying the mounted bundle option for knowledge bundle replication has changed. Previously, the setting was `shareBundles=false`. The new setting is `replicationPolicy=mounted`. If the upgrade process finds the old setting in your configuration files, it converts it automatically to the new setting, so no manual intervention is required on your part.

The Django framework and all its associated components have been removed from the Splunk platform. Apps and dashboards that use this framework will not function anymore. Before you upgrade, confirm that apps and dashboards you regularly use no longer use Django, as an upgrade will render them inoperable.

Updated field alias behavior might result in null values

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 7.2

When you upgrade to a version of Splunk Enterprise that is 7.2 or higher, the behavior of certain **field alias** configurations changes. This behavior affects events where the alias field is already present before field alias processing takes place.

- If both the source field and the alias field are present in the event, the search head overwrites the value of the alias field to match the value of the source field.
- If the alias field is present in the event but the source field is missing or has a null value, the search head removes the alias field from the event.

This is how field aliasing should work in these situations. However, you might have searches that rely on the old field alias behavior. A new `ASNEW` syntax has been added to enable you to provide the pre-7.2 behavior to your field alias configurations.

The pre-7.2 field alias behavior enabled users to create sets of field alias configurations that matched multiple source fields with one alias field. You can use the `ASNEW` syntax to continue doing this. A better practice would be to use a **calculated field** that uses the `coalesce` function to create a new field that takes the value of one or more existing fields. This method lets you be explicit about ordering of input field values in the case of null fields. For example: `EVAL-ip = coalesce(clientip, ipaddress).`

See Field alias behavior change in the *Release Notes* for further details.

Splunk Enterprise meters metric data points under 150 bytes by volume on a scale that is similar to the scale used for event data

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 7.3

There is a change to the way that Splunk Enterprise meters metrics data in version 7.3. Prior to this release, all metric data points counted against the license as if they were a flat 150 bytes. As of 7.3, Splunk Enterprise meters each metric data point that measures less than 150 bytes by volume on a scale that is similar to the scale used for event data. This scale is capped at 150 bytes. Metric events with volumes over 150 bytes are metered as if they are only 150 bytes. As a result, license usage for such events might be lower than in prior release.

For additional information on licensing, see How Splunk Enterprise licensing works in the *Admin Manual*.

The tstats command now reports the actual number of matched events in an index bucket as the scan count of those events

Applicable components: Splunk Enterprise

Applicable OSES: all

Version introduced: 7.3

The `tstats` search command, which performs statistical queries on indexed fields in tsidx files, has updated behavior with regard to the number of events it finds in a bucket.

Previously, the command reported the scan count of events in a bucket. After an upgrade, the command now reports the actual number of matched events in a bucket as the scan count. If you use `tstats` for searches, you might see scan counts fall significantly. This does not indicate incorrect results. -->

Considerations for new features

Splunk has introduced the following new features in this version of Splunk Enterprise. You might need to perform some configuration after an upgrade to enable and take advantage of these features.

New configuration settings for importing structured data files

Applicable components: Splunk Enterprise and Splunk Universal Forwarder

Applicable OSES: all

Version introduced: 8.0

Splunk Enterprise 8.0.0 and higher has added some new settings for configuring the indexing of structured data. Previously, the structured data processor converted characters that it encountered in header field names that were neither alphanumeric nor a space into underscores. With the new `HEADER_FIELD_ACCEPTABLE_SPECIAL_CHARACTERS` setting in `props.conf`, you can control which characters the processor accepts as valid in header field names. Some characters might cause the processor to malfunction, so if you import lots of structured data and want to use the feature, use a test index to confirm that the processor imports the data in the way you want.

Learn about known upgrade issues

To learn about any additional upgrade issues for Splunk Enterprise, see the Known Issues - Upgrade Issues page in the *Release Notes*.

How to upgrade a distributed Splunk Enterprise environment

Distributed Splunk Enterprise environments vary widely. Some have multiple indexers or search heads, and others have indexer- and search-head clusters. These types of environments present challenges over upgrading single-instance installations.

Determine the upgrade procedure to follow for your type of environment

Depending on the kind of distributed environment you have, you might have to follow separate instructions to complete the upgrade. This topic provides guidance on how to upgrade distributed environments that do not have any clustered elements like index- or search-head clusters. Environments with clustered elements, such as indexer clusters and search head clusters, have different upgrade procedures in different topics. Search head pooling has been removed in version 8.0 of Splunk Enterprise, so there are no upgrade instructions for that type of distributed deployment.

- To upgrade a distributed environment that does not have any clustered elements, follow the procedures in this topic.
- To upgrade an environment with index clusters, see Upgrade an indexer cluster in *Managing Indexers and Clusters of Indexers*.
- To upgrade an environment with search head clusters, see Upgrade a search head cluster in *Distributed Search*.
- If you have additional questions about upgrading your distributed Splunk Enterprise environment, log a case at the Splunk Support Portal.

Cross-version compatibility between distributed components

While there is some range in compatibility between various Splunk software components, they work best when they are all at a specific version. If you have to upgrade one or more components of a distributed deployment, you should confirm that the components you upgrade remain compatible with the components that you don't.

- For information on compatibility between different versions of **search heads** and **search peers** (indexers), see System requirements and other deployment considerations for distributed search in *Distributed Search*.
- For information on compatibility between indexers and forwarders, see Compatibility between forwarders and indexers in *Forwarding Data*.

Test apps prior to the upgrade

Before you upgrade a distributed environment, confirm that Splunk apps work on the version of Splunk Enterprise that you want to upgrade to.

1. On a reference machine, install the full version of Splunk Enterprise that you currently run.
2. Install the apps on this instance.
3. Access the apps to confirm that they work as you expect.
4. Upgrade the instance.
5. Access the apps again to confirm that they still work.

If the apps work as you expect, move them to `$SPLUNK_HOME/etc/apps` on each search head during the search head upgrade process.

Upgrade a distributed environment with multiple indexers and non-pooled search heads

This procedure upgrades the search head tier, then the indexing tier, to maintain availability.

Prepare the upgrade

1. Confirm that any apps that the non-pooled search heads use will work on the upgraded version of Splunk, as described in "[Test your apps prior to the upgrade](#)" in this topic.
2. (Optional) If you use a **deployment server** in your environment, disable it temporarily. This prevents the server from distributing invalid configurations to your other components.
3. (Optional) Upgrade the deployment server, but do not restart it.

Upgrade the search heads

1. Stop Splunk Enterprise services on one of the search heads.
2. Upgrade the search head. Do not let it restart.
3. After you upgrade the search head, place the confirmed working apps into the `$(SPLUNK_HOME)/etc/apps` directory of the search head.
4. Re-enable and restart the search head.
5. Test apps on the search head for operation and functionality.
6. If there are no problems with the search head, then disable and upgrade the remaining search heads, one by one. Repeat this step until you have reached the last search head in your environment.
7. (Optional) Test each search head for operation and functionality after you bring it up.
8. After you upgrade the last search head, test all of the search heads for operation and functionality.

Upgrade the indexers

1. Stop Splunk Enterprise services, and upgrade the indexers, one by one. You can restart the indexers immediately after you upgrade them.
2. Test search heads to ensure that they find data across all indexers.
3. After you upgrade all indexers, restart your deployment server.

Upgrade forwarders

After your distributed environment upgrade, review the forwarder versions used in your environment and check for feature compatibility and support.

To upgrade universal forwarders, see the Forwarder Manual.

Changes for Splunk App developers

If you develop apps for the Splunk platform, read this topic to find out what changes we've made to how the software works with apps in version 7.3.x, and how to migrate any existing apps to work with the new version.

Changes

- Django Web Framework is removed from the product in Splunk Enterprise 7.3.0. Django Web Framework was first deprecated in Splunk Enterprise 6.3.0. Apps and pages that are built using Django Web Framework will not work in Splunk Enterprise 7.3.0 and later, and will return a 404 error. Customers should migrate Django-based apps and pages to Simple XML or HTML Dashboards framework.
- For other changes, or to learn more about Splunk app development, visit the Splunk Dev portal.

Visit the Splunk Dev portal

To learn more about Splunk app development, visit the Splunk Dev portal.

Upgrade to version 9.0 on UNIX

Before you upgrade

Before you upgrade, see [About upgrading to 9.0: READ THIS FIRST](#) for information on changes in the new version that can impact you if you upgrade from an existing version.

Splunk Enterprise does not provide a means of downgrading to previous versions. If you need to revert to an older Splunk release, 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**, including Splunk Enterprise configurations, indexed data, and binaries.

For information on backing up data, see Back up indexed data in *Managing Indexers and Clusters of Indexers*.

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

How upgrading works

To upgrade a Splunk Enterprise installation, you must install the new version directly on top of the old version (into the same installation directory.) When Splunk Enterprise 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 Enterprise does not change your configuration until after you restart it.

Upgrade Splunk Enterprise

1. Go to the machine with the Splunk Enterprise instance you want to upgrade, and open a shell prompt.
2. Verify the folder where Splunk Enterprise is installed, and change to the `$(SPLUNK_HOME)/bin` directory.
3. Stop the Splunk Enterprise services by running `systemctl stop Splunkd.service` OR `$(SPLUNK_HOME)/bin/splunk stop`
4. Confirm that no other processes will automatically start Splunk Enterprise, such as a configuration management or service management tool.
5. To upgrade and migrate the existing configurations, install the latest Splunk Enterprise package directly over your existing deployment.
 - ◆ If you are using a `.tar` file, expand it into the same directory with the same ownership as your existing Splunk Enterprise instance. This overwrites and replaces the default files, but does not remove unique files or file paths. Example: `tar xzf splunk-8.3.0-12345678-Linux-x86_64.tgz -C /opt`
 - ◆ If you use a package manager, such as RPM, type `rpm -U splunk_package_name.rpm`
 - ◆ If you use a `.dmg` file on MacOS, double-click it and follow the instructions. Specify the same installation directory as your existing installation.

6. Start the Splunk Enterprise services by running `$SPLUNK_HOME/bin/splunk start`. Splunk Enterprise 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]
```

7. (Optional) 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 now. If you choose to view the expected changes, the script provides a list but does not start any services. After you review the migration changes and are ready to proceed with migration and upgrade, start the Splunk Enterprise services again.

Upgrade and accept the license agreement simultaneously

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

- To accept the license and begin the upgrade without viewing the changes, use the following command:

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

Upgrade to version 9.0 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 Enterprise, if you need to downgrade, you must uninstall the upgraded version and then reinstall the previous version of Splunk Enterprise 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.

Before you upgrade

Before you upgrade, see [About upgrading to 9.0: READ THIS FIRST](#) for information on changes in the new version that can impact you if you upgrade from an existing version.

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

Changing Splunk Enterprise ports during an upgrade is not supported

Splunk Enterprise 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 Enterprise configurations, indexed data, and binaries.

- For information on backing up data, see Back up indexed data in *Managing Indexers and Clusters of Indexes*.
- 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 Enterprise.

Upgrade Splunk Enterprise using the GUI installer

1. Go to the Splunk.com Free Trials and Downloads page (Login required.)
2. Select "Splunk Enterprise".
3. Select "Download now" to get the latest release, or click the link to "Previous Releases" to find a specific version.
4. Download the MSI file to the host.
5. Double-click the MSI file. The installer runs and attempts to detect the existing version of Splunk Enterprise installed on the machine. When it locates the prior installation, it displays a pane that asks you to accept the licensing agreement.
6. Accept the license agreement. The installer then installs the updated Splunk Enterprise. This method of upgrade retains all parameters from the existing installation. The installer restarts Splunk Enterprise services when the upgrade is complete, and places a log of the changes made to configuration files during the upgrade in `%TEMP%`.

Upgrade using the command line

1. Go to the Splunk.com Free Trials and Downloads page (Login required.)
2. Select "Splunk Enterprise."
3. Select "Download now" to get the latest release, or click the link to "Previous Releases" to find a specific version.
4. Download the MSI file to the host.
5. 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 Enterprise 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.
6. Depending on your specification, Splunk Enterprise might start automatically when you complete the installation.

Migrate a Splunk Enterprise instance from one physical machine to another

Important: These migration instructions are for on-premises Splunk Enterprise instances only.

Important: These migration instructions are for on-premises Splunk Enterprise instances only.

If you are a Splunk Cloud Platform customer or want to migrate your data from Splunk Enterprise to Splunk Cloud Platform, do not use these instructions. Contact Professional Services for assistance.

You can migrate a Splunk Enterprise instance from one server, operating system, architecture, or filesystem to another, while maintaining the indexed data, configurations, and users. Migrating an instance of Splunk Enterprise different than upgrading one, which is merely installing a new version on top of an older one.

Do not attempt to migrate a Splunk Enterprise installation to Splunk Cloud Platform using these instructions. Doing so could result in data loss. Speak with Professional Services or your Splunk account representative for information and instructions.

When to migrate

There are a number of reasons to migrate a Splunk Enterprise install:

- Your Splunk Enterprise installation is on a host that you wish to retire or reuse for another purpose.
- Your Splunk Enterprise installation is on an operating system that either your organization or Splunk no longer supports, and you want to move it to an operating system that does have support.
- You want to switch operating systems (for example, from *nix to Windows or vice versa)
- You want to move your Splunk Enterprise installation to a different file system.
- Your Splunk Enterprise installation is on 32-bit architecture, and you want to move it to a 64-bit architecture for better performance.
- Your Splunk Enterprise installation is on a system architecture that you plan to stop supporting, and you want to move it to an architecture that you do support.

Considerations for migrating Splunk Enterprise

While migrating a Splunk Enterprise instance is simple in many cases, there are some important considerations to note when doing so. Depending on the type, version, and architecture of the systems involved in the migration, you might need to consider more than one of these items at a time.

When you migrate a Splunk Enterprise instance, note the following.

Differences in Windows and Unix path separators

The path separator (the character used to separate individual directory elements of a path) on *nix and Windows is different. When you move index files between these operating systems, you must confirm that the path separator you use is correct for the target operating system. You must also make sure that you update any Splunk configuration files (in particular, `indexes.conf`) to use the correct path separator.

For more information about how path separators can impact Splunk Enterprise installations, see Differences between *nix and Windows in Splunk operations in the *Admin* manual.

Windows permissions

When moving a Splunk Enterprise instance between Windows hosts, make sure that the destination host has the same rights assigned to it that the source host does. This includes but is not limited to the following:

- Ensure that the file system and share permissions on the target host are correct and allow access for the user that runs Splunk Enterprise.

- If Splunk Enterprise runs as an account other than the Local System user, that the user is a member of the local Administrators group and has the appropriate Local Security Policy or Domain Policy rights assigned to it by a Group Policy object

Architecture changes

If you downgrade the architecture that your Splunk Enterprise instance runs on (for example, 64-bit to 32-bit), you might experience degraded search performance on the new host due to the larger files that the 64-bit operating system and Splunk Enterprise instance created.

Distributed and clustered Splunk environments

When you want to migrate data on a distributed Splunk Enterprise instance (that is, an indexer that is part of a group of search peers, or a search head that has been configured to search indexers for data), you should remove the instance from the distributed environment before attempting to migrate it. Distributed environments must run on the same operating system, and have other requirements. See Summary of key requirements in the *Managing Indexers and Clusters of Indexers* manual.

Bucket IDs and potential bucket collision

If you migrate a Splunk Enterprise instance to another Splunk instance that already has existing indexes with identical names, you must make sure that the individual buckets within those indexes have bucket IDs that do not collide. Splunk Enterprise does not start if it encounters indexes with buckets that have colliding bucket IDs. When you copy index data, you might need to rename the copied bucket files to prevent this condition.

How to migrate

When you migrate on *nix systems, you can extract the tar file you downloaded directly over the copied files on the new system, or use your package manager to upgrade using the downloaded package. On Windows systems, the installer updates the Splunk files automatically.

1. Stop Splunk Enterprise services on the host from which you want to migrate.
2. Copy the entire contents of the \$SPLUNK_HOME directory from the old host to the new host. Copying this directory also copies the `mongo` subdirectory.
3. Install Splunk Enterprise on the new host.
4. Verify that the index configuration (`indexes.conf`) file's volume, sizing, and path settings are still valid on the new host.
5. Start Splunk Enterprise on the new instance.
6. Log into Splunk Enterprise with your existing credentials.
7. After you log in, confirm that your data is intact by searching it.

How to move index buckets from one host to another

If you want to retire a Splunk Enterprise instance and immediately move the data to another instance, you can move individual buckets of an index between hosts, as long as:

When you copy individual bucket files, you must make sure that no bucket IDs conflict on the new system. Otherwise, Splunk Enterprise does not start. You might need to rename individual bucket directories after you move them from the source system to the target system.

1. Roll any hot buckets on the source host from hot to warm.

2. Review indexes.conf on the old host to get a list of the indexes on that host.
3. On the target host, create indexes that are identical to the ones on the source system.
4. Copy the index buckets from the source host to the target host.
5. Restart Splunk Enterprise.

Plan your Splunk Enterprise upgrade to work with the Python 3 migration

Upgrade using the Python 3 runtime and dual-compatible Python syntax in custom scripts

Splunk Enterprise versions 7.x and lower don't support Python 3, and Splunk Enterprise versions 9.0 and higher don't support Python 2. If you choose to write your private apps to only be compatible with Python 3, then you might experience breaking changes if you install your app before upgrading to Splunk Enterprise version 8.0 or later, or if the previous version of the app has Python 2 syntax in Splunk Web-dependent scripts and templates.

Follow this upgrade path to transition to use the Python 3 runtime. This upgrade path has three parts:

1. [Pre-upgrade steps](#)
2. [Upgrade Splunkbase apps and validate](#)
3. [Upgrade Splunk Enterprise](#)

Pre-upgrade steps

Before you migrate to Python 3 during your upgrade, take the following steps. Using a test environment is optional. If you don't use a test environment, this process might impact performance.

Prepare for upgrade with the Upgrade Readiness app

1. Optional: In a test environment, install all your production apps and replicate any custom scripts you included in your production environment outside of an app context.
2. Install either the Upgrade Readiness App 3.0.1 or the Python Readiness App 2.0.0 and run it, preferably on your test instance. See *About the Upgrade Readiness App* in the *Upgrade Readiness* manual for more information.
3. Review the test results for your private apps and follow the instructions in the Upgrade Readiness App to resolve any issues. Rewrite Python scripts to be compatible with both Python 2 and Python 3. See *Writing scripts compatible with Python 2 and Python 3* in the *Python 3 Migration* manual for more information.
4. Review the test results for the apps you downloaded from Splunkbase and check for file paths that contain content you have customized or extended. Follow the instructions in the Upgrade Readiness App to resolve the issues.

Manually prepare for upgrade

If you choose not to use the Upgrade Readiness App, manually check your Splunk Enterprise deployment for issues. Confirm you've met these upgrade readiness requirements in your private apps and extensions before you continue:

1. Optional: In a test environment, install all your production apps and replicate any custom scripts you included in your production environment outside of an app context.
2. Remove all Advanced XML.
3. If you are running SplunkWeb Legacy Mode, disable it by editing `%SPLUNK_HOME%\etc\system\local\web.conf` to remove the `appServerPorts` and `httpport` attributes.
4. Rename any files named `test.py` to a non-reserved name.
5. Adjust custom Mako templates to be Python 2 and Python 3 compatible.
6. Adjust custom CherryPy endpoints in apps to be Python 2 and Python 3 compatible.

7. Rewrite all remaining Python code to be dual-compatible with Python 2 and 3 either by using the Python SDK or by leveraging the Six or Future libraries. Keep a list of the files you rewrite so that you can test them when you upgrade.

For more information about these upgrade requirements, see Python 3 migration with the Splunk platform in the *Python 3 Migration* manual.

Upgrade Splunkbase apps and validate

Only after you have completed all the preparation steps, upgrade Splunk Enterprise in a test environment, following the three-phase Splunk Enterprise upgrade process. For more information, see [Splunk Enterprise upgrade process](#).

To reduce performance impact, upgrade Splunk Enterprise in a test environment first to test for breaking changes. Don't upgrade in your production environment directly.

Follow these steps to validate the upgrade in your test environment:

1. Upgrade your installed Splunkbase apps to versions that are dual-compatible with Splunk Enterprise 7.x and 8.x, since these should be compatible for both Python 2 and Python 3. Test for any breaking changes. If the developer hasn't provided a version of the app that is compatible with your Splunk Enterprise version, update it yourself or remove the app. If the developer has provided separate versions of the app for Splunk Enterprise 7.x and 8.x, follow any upgrade instructions provided by the developer. If none are provided, assume the app is written for Python 3 only, and choose one of the following options:
 - ◆ Upgrade the app to the Splunk Enterprise 8.x compatible version only after upgrading your instance to Splunk Enterprise 8.x, and note any potential breaking changes in Splunk Web.
 - ◆ If there are many Splunkbase apps that are Python 3-only and might cause unplanned downtime, take your deployment offline for maintenance during the Splunk Enterprise upgrade. At this time, remove the apps, take your deployment offline, upgrade the apps to the versions compatible with Splunk Enterprise 8.x, and then take your deployment back online.
2. Test your Splunk Web-dependent Python scripts by accessing all views that rely on custom CherryPy endpoints or Mako templates. Ensure that they function as expected.
3. Test the functionality of your dual-compatible Python scripts. If you prefer, you can run each dual-compatible script with Python 3 instead. Consult the list of files you rewrote, and, one at a time, go to the configuration file responsible for invoking each script and add `python.version=python3` to the stanza that invokes the script. Test for errors using the following table:

Script type	File	Restart required?	How to test
Custom search commands	<code>commands.conf</code>	Yes	Run the command and check that it works as expected.
Modular inputs	<code>inputs.conf</code>	Yes	Enable the input and check that data arrives as expected.
Scripted inputs	<code>inputs.conf</code>	Yes	Enable the input and check that data arrives as expected.
Custom alert actions	<code>alert_actions.conf</code>	No	Run the custom alert action and verify that it works as expected.
Scripted lookups	<code>transforms.conf</code>	Yes	Verify the lookup is working as expected.
Custom REST endpoints	<code>restmap.conf</code>	No	Access the endpoint and check that it responds as expected.

Script type	File	Restart required?	How to test
Scripted authentication	authentication.conf	No. Go to Settings > Access controls > Authentication method and click Reload authentication configuration .	

4. If you use scripted authentication, check that the Python interpreter path specified in your `scriptPath` setting is the canonical path, `$(SPLUNK_HOME)/bin/python`. If you enter a custom path to another interpreter, scripted authentication will use the interpreter you specify rather than the interpreter you choose in the `python.version` setting.
5. If you encountered any errors, adjust the Python files and test them again.
6. Once you're satisfied that all your scripts are working as expected, specify the global Python 3 runtime in your test environment.
 1. Go to `$(SPLUNK_HOME)/etc/system/local/server.conf` and set `python.version=python3`.
 2. Restart Splunk Enterprise.
 3. Test any Python scripts that are on a critical path, such as scripted authentication, to ensure they continue to work.

Upgrade Splunk Enterprise

Upgrade your production environment on your search heads, indexers, and forwarders, following the three-phase Splunk Enterprise upgrade process. For more information, see [Splunk Enterprise upgrade process](#). As you upgrade the components of your deployment, follow this order of operations:

1. Upgrade Splunk Enterprise.
2. Note any potential breaking changes, such as apps that aren't compatible yet with Splunk Enterprise version 8.x at time of upgrade.
3. Install any apps that are only compatible with Python 3 and Splunk Enterprise 8.x.
4. Set your scripts to run Python 3 only. Go to the appropriate script `.conf` file and set `python.version=python3`, as covered when validating apps.
5. (Optional) Specify the global Python 3 runtime in your production environment. Go to `$(SPLUNK_HOME)/etc/system/local/server.conf` and set `python.version=python3`.
6. Start Splunk Enterprise.

Uninstall Splunk Enterprise

Uninstall Splunk Enterprise

Learn how to remove Splunk Enterprise from a host by following the procedures in this topic.

Prerequisites

1. If you configured Splunk Enterprise to start on boot, remove it from your boot scripts before you uninstall.

```
./splunk disable boot-start
```

2. Stop Splunk Enterprise. Navigate to `$SPLUNK_HOME/bin` and type `./splunk stop` (or just `splunk stop` on Windows).

Uninstall Splunk Enterprise with your package management utilities

If you used local package management tools to install Splunk Enterprise, use those same tools to uninstall Splunk Enterprise. In most cases, files that were not originally installed by the package are retained. These files include your configuration and index files which are located in the Splunk Enterprise installation directory.

In these instructions, `$SPLUNK_HOME` refers to the Splunk installation directory. On Windows, this is `C:\Program Files\Splunk` by default. For most Unix platforms, the default installation directory is `/opt/splunk`. On Mac OS X, it is `/Applications/splunk`.

RedHat Linux

```
rpm -e splunk_product_name
```

Debian Linux

```
dpkg -r splunk
```

Remove all Splunk files, including configuration files

```
dpkg -P splunk
```

Other things you might want to delete

- If you created any indexes and did not use the Splunk Enterprise default path, you must delete those directories as well.
- If you created a user or group for running Splunk Enterprise, you should also delete them.

Windows

- Use the **Add or Remove Programs** option in the Control Panel. In Windows 8.1 and 10, and Windows Server 2012 R2, 2016, and 2019, that option is available under **Programs and Features**.
- (Optional) You can also uninstall Splunk Enterprise from the command line by using the `msiexec` executable against the Splunk installer package.

```
msiexec /x splunk-<version>-x64.msi
```

Under some circumstances, the Microsoft installer might present a reboot prompt during the uninstall process. You can safely ignore this request without rebooting.

Uninstall Splunk Enterprise manually

If you can't use package management commands, use these instructions to uninstall Splunk Enterprise.

1. Stop Splunk Enterprise.

```
$(SPLUNK_HOME)/bin/splunk stop
```

2. Find and `kill` any lingering processes that contain "splunk" in their name.

For Linux

```
kill -9 `ps -ef | grep splunk | grep -v grep | awk '{print $2;}'`
```

For Mac OS

```
kill -9 `ps ax | grep splunk | grep -v grep | awk '{print $1;}'`
```

3. Remove the Splunk Enterprise installation directory, `$(SPLUNK_HOME)`.

For Linux

```
rm -rf /opt/splunk
```

For Mac OS

```
rm -rf /Applications/splunk
```

You can also remove the installation directory by dragging the folder into the Trash.

4. Remove any Splunk Enterprise datastore or indexes outside the top-level directory, if they exist.

```
rm -rf /opt/splunkdata
```

5. Delete the `splunk` user and group, if they exist.

For Linux

```
userdel splunk  
groupdel splunk
```

For Mac OS

Use the **System Preferences > Accounts** panel to manage users and groups.

For Windows

Open a command prompt and run the command `msiexec /x` against the msi package that you used to install Splunk Enterprise. If you don't have that package, get the correct version from the download page.

Reference

PGP Public Key

You can copy the Pretty Good Privacy (PGP) public key for Splunk software from this page or download the file using HTTPS.

- This PGP public key is used to sign packages of Splunk software that are released on or after August 15, 2018.
- The signature only applies to the RPM package. For all other package types, use the checksum files. See [Install Splunk Enterprise securely](#).

PGP public key block

-----BEGIN PGP PUBLIC KEY BLOCK-----

```
mQINBftbeBEADjLzD+QXyTqLwT2UW1D1e5MpBj+C5cbaCIpFEh1+KemcnUKHls
TlCxEpzJczZPiYtcp+wtKCaNG/zoEvYcQ0jKk6Wgoa2cLkDeHtNiuBCHrztgeDTe
FpPT+xmtLoJvulT0JV/iPG7p5FBGYKOKApnd/awRRC47p1CGfVA3VVdQP8jhpMZV
T9C86hWbNo/NRjNH69x1xAe/9POc8KmvxZQb+KGG5tulGIWa7j1TMw850HZwFcft
F13DiAVGcj516K8oZBb5bjgu2ZvpCtMRbCmrzx26ilcB7VJRSTaB6G8MqRzVgLuJ
11dTG2XMuBw+3UcjAlZ/y6Cut0Gc5FHIKqMVXf29y9uXddvIqQnke0AkOj6flm6
OElvmq7v+NVYLRb9XTy0oWTwOtyGTTso2xwZ8itDT4rIWeta0FxtQPt8Kq369ZGy
CbK11PU9IrkAeST0AkXyFQXqPc0IzHxz3AhOLzvwM/9/00Ws0ONbxdyTCxQjrhe6
1YBoVv2T5K27fTp7rMFEstyU0NFI3J5P/oxg5ts6y21CMUB7Q71yAOWVZPgucOAH
7iiNmvrytuGT0c8TfJkulcneaJw9jmNvKVD/r3qj6YTAL3mqC0yYx3PiLyUVm8OZ
q90hpFHAI7zV1u6zmqV4EkWg5tEknMwCjQnyIfn0Jx8LedDjbTM8Dt9VKQARAQAB
tCFTcGx1bmssIEluYy4gPHJlbGVhc2VAc3BsdW5rLmNvbT6JAK4EEwEiADgWIQRy
wzMQt6NUwSedtmle+gHts81EIAUCW1t5sQIbAwULCQgHAgYVCAKkCwIEFgIDAQIe
AQIXgAAKCRBe+gHts81EIEsUD/9urCsBW40ahPr1gBsu6T1FbVWFN6TK7NpByecr
KzhD1OGJbh7g1ulqR088ncUb/iPfbjppJ0RbskrZQKVVbmnhLeNPw4oqHq4kNmN
Kc8iV9tywnw55Ww5Y0cJoeWrx9Ireub3+1GhKzUomIK0TuQtMULmW7Tdmw46iEDgC
qox2hOutLMFjrT9XOFnluCeyi8HL9m6xUlvvsxYxqWIzWUvoWH3AwpGSPMwg/nzH
V11Wz9IJOlqjQFbiA1Vmb/UEkP60JatXWtNKJ70qTLaq29XBSaJ01NiQFZYb8uCU
GSqNOKYUwiO3ZivmVY1XB7fC2uHpU45g/d2PrRKgVvIOC9xKiG8+jh/WuW1T14i
vVjAIEFIW08Ni7uoR9xi+0ZxzKp00tGO2Cgv0Cf3TYQRsgrD7QDRBN2az4HtF
WvxJuOYjNL17mp+Lx0Aj9wtb1WkYNBV0NMXTThnZsDU6Uo6iJJa2uBwkT8M1jCHX
n7DjVFZYoz6m2cwUdR5XswfSq1a7LcSbef4CIC1H0mVxVzeB2B6xGxpVIMNGs4
B1RXW1amVeKmv9ZbTAQpGNVMyGJ8oOhksBFL2Ng0Z5kA9aCuwr10jyrxBdglfGd/
wmEGIX2cLNNvs+Elh4JzFuKsURWbJ8qF17cQvKQkS+UTwu7e3CCp8VztFRqPvgQi
A+2oI7kCDQRbW3mxARAAtOBTC9nNiY3301QKzTyPvud3XI03RZTXVsSHVP4yV0x
fobD2aRhmjxwRjrajZnMCEfKb7yYtSbyiRfznLoycFBse6p4y9gguWEIgaW6TTQP
zQTEgi6AKt38nqDN42L/WurNhAKq9R5X/85vr2t6b18Yp2kw62okbuTtVLjuNwzh
tnZE/HziWVbtBy0KfZ0c6QMuhN7j0U67+QJeIzLcQuBn4qnb177TrtnqN29aFTXX
mnUA7qTOAvL+wsoyOcuOboj4N45H5s/izPSiXkoUM1ITuuUI3QH46zw5cEvSLg+
WImmwZCN4tC275abjxw7XbirglV1E01CWoALIOAh1BwXDA/JJGwbGOp+ueE7askJ
TiAtP9EM1mJSWnbE9uKDUvEMiAavwtt0kWmQOrB4HFY0AstOnCxwQYCOB0CDImyq
Sb1c3tqvoZzbjPBHQFvxClzxfGdmvQwoxr2WRfsspLPuG1FzgmX29/WaOV747W
TwJP9xw10tJmAkq/+CH6J12PmXHy9sJRdk6d1PPEuHjJ588U3Kwc7B5uAtgnwQO8
aS4zPM45y6+J1D2SdM0ydwuqQ9z9wwa022EGTa89k5Vfigx+C/VadMa1Bu/NSKZ8
7S0NpQGbrWdP76gSKv1T/15hYVg2n0s1lhtVmM8hVZQO3k04zFj10rNNjwWor0A
EQEAAYkCNQYQAQgAIBYhBFjDMx3o1TBJ522aV76Ae2zzUQgBQJbW3mxAhSMAAJ
EF76Ae2zzUQg26YP/0dj631dEluB8L7+dFm9stebcMpgxAugmntdlprDKgi6Rhfd
ks7uff+mny731GZPcJW1YKi797qerG501AI4siaK9FRKzW4PLIGvh0oNg2wrSP/+
7qTFf+ZbT7H5VpIqvcnntR05pi1KiMIXW82h47daFYVnhQPbV4+USHwFG7r3Lku
XdiS4hrcoe+Y/a9zGVAdU9QwrT8CuNAw8SYNYx1rJECHiMxmMaEw42a5NARoFdbh
swrR6Mwy5sPhzOHjSI/ZPyM/W9TKAoXfmDQSGDrvnU6NAdpIbP1Ab1FtMjuARfRg
```

```
8ndqfm/n8MIvAxjzoBBZkdV5HLondX3fLVNewnvSWQx90lV4a7+dKXeQ8TueOMq+
XMA4RKsh3gEMJWbVRZwZnxy+3UKGJD3e10+C7m483ptR8Tj8qBq5KELO0vkcq8+a
eHIbzmQsSj9iAdNfGVLYhimzpz5NCT12sgmy4g33pdljMtUzdFZhvelVzMNlkLZ
AmAJX7yZLQwLsXDEpffg2S/U8vYAZNTdeZqKvmvCCO+fweRRC7NnnPJQ7nVhL7r
VDxHuk8oMqBQIUdE7Z+WDfyagMMhJWbeMNNnhTZdoPmpXEGkjUKwPDYl+GmF50c1
6vjXtbrCP42pu2IQxiqiaTSLei8LRwPck1eE+78sSUxjVuWRuThoYRhGYoXt
=ivRW
-----END PGP PUBLIC KEY BLOCK-----
```

GPG key ID

The GPG public key ID is `key ID b3cd4420`.

Install the PGP public key

See [Verify Signatures](#).