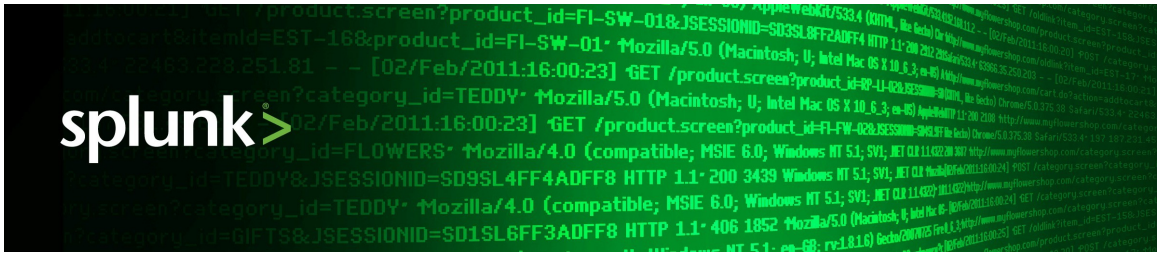

Splunk 4.3.3

Installation Manual

Generated: 7/25/2012 1:16 pm

Table of Contents

Table of Contents

Welcome to the Splunk Installation Manual

What's in the Installation Manual

Use this guide to find system requirements, licensing information, and procedures for installing or migrating Splunk.

Note: If you want to install the Splunk **universal forwarder**, see the Distributed Deployment manual: "Universal forwarder deployment overview". Unlike Splunk heavy and light **forwarders**, which are just full Splunk instances with some features disabled, the universal forwarder is an entirely separate executable, with its own set of installation procedures. For an introduction to forwarders, see "About forwarding and receiving".

Find what you need

You can use the table of contents to the left of this panel, or simply search for what you want in the search box in the upper right.

If you're interested in more specific scenarios and best practices, you can visit the Splunk Community Wiki to see how other users Splunk IT.

Make a PDF

If you'd like a PDF of any version of this manual, click the red **Download as PDF** link below the table of contents on the left side of this page. A PDF version of the manual is generated on the fly for you, and you can save it or print it out to read later.

System requirements

Before you download and install the Splunk software, read the following sections for the supported system requirements. If you have ideas or requests for new features to add to future releases, email Splunk Support. Also, you can follow our Product Roadmap.

Refer to the download page for the latest version to download. Check the release notes for details on known and resolved issues.

For a discussion of hardware planning for deployment, check out the topic on capacity planning in this manual.

Supported OSes

Splunk is supported on the following platforms.

- Solaris 9, 10 (x86, SPARC)
- Linux Kernel vers 2.6.x and above (x86: 32 and 64-bit)
- FreeBSD 6.1 (x86: 32-bit), 6.2, 7.x, 8.x (x86: 32 and 64-bit)
- Windows Server 2003/2003 R2 (64-bit, supported but not recommended on 32-bit)
- Windows Server 2008/2008 R2 (64-bit, supported but not recommended on 32-bit)
- Windows XP (32-bit)
- Windows Vista (32-bit, 64-bit)
- Windows 7 (32-bit, 64-bit)
- MacOSX 10.5 and 10.6 (32-bit and 64-bit in one download. **10.6 is only supported in 32-bit mode.**)
- AIX 5.2, 5.3, and 6.1
- HP-UX 11iv2 (11.22) and 11iv3 (11.31) (PA-RISC or Itanium, **gnu tar is required to unpack the tar.gz archive**)

Windows

Certain parts of Splunk on Windows require elevated permissions to function properly. For additional information about what is required, read the following topics:

- "Splunk architecture and processes" in this manual.
- "Choose the user Splunk should run as" in the "Install on Windows" topic in this manual.
- "Considerations for deciding how to monitor remote Windows data" in the Getting Data In Manual.

FreeBSD 7.x

To run Splunk 4.x on 32-bit FreeBSD 7.x, install the compat6x libraries. Splunk Support will supply "best effort" support for users running on FreeBSD 7.x. For more information, refer to this Community wiki topic.

Fedora Core 13

Users of Fedora Core 13 must be sure to update glibc to 2.12-2 or higher (released 2010-06-07) to resolve a glibc memory allocator bug - https://bugzilla.redhat.com/show_bug.cgi?id=594784 The symptom of the glibc-2.12-1 problem are program crashes with the message 'invalid fastbin entry (free)'. This is only expected to affect the 32 bit splunk build, but as it will likely cause crashes in system tools as well, the update is recommended for all Fedora Core 13 splunk users, 32-bit and 64-bit.

Creating and editing configuration files on non-UTF-8 OSes

Splunk expects configuration files to be in ASCII/UTF-8. If you are editing or creating a configuration file on an OS that is non-UTF-8, you must ensure that the editor you are using is configured to save in ASCII/UTF-8.

IPv6 platform support

All Splunk-supported OS platforms are supported for use with IPv6 configurations except for the following:

- HPUX PA-RISC
- Solaris 8 and 9
- AIX

Refer to "Configure Splunk for IPv6" in the Admin Manual for details on Splunk IPv6 support.

Supported browsers

- Firefox 3.6, 10.x, and latest
- Internet Explorer 6, 7, 8, and 9. Internet Explorer 8 is supported in IE7 compatibility mode only. Internet Explorer 9 is not supported in compatibility mode.
- Safari (latest)
- Chrome (latest)

You should also make sure you have the latest version of Flash installed to render any charts that use options not supported by the JSChart module. For more information about this subject, see "Advanced charting options" in the Developing Dashboards, Views, and Apps for Splunk Web manual.

Recommended hardware

Splunk is a high-performance application. If you are performing a comprehensive evaluation of Splunk for production deployment, we recommend that you use hardware typical of your production environment; this hardware should **meet or exceed** the recommended hardware capacity specifications below.

For a discussion of hardware planning for production deployment, check out the topic on capacity planning in this manual.

Splunk and virtual machines

Running Splunk in a virtual machine (VM) on any platform will degrade performance. This is because virtualization works by abstracting the hardware on a system into resource pools from which VMs defined on the system can draw from as needed. Splunk needs sustained access to a number of resources, particularly disk resources for indexing operations, which can cause problems when running it in a VM, or alongside other VMs.

Recommended and minimum hardware capacity

Platform	Recommended hardware capacity/configuration	Minimum supported hardware capacity
Non-Windows platforms	2x quad-core Xeon, 3GHz, 8GB RAM, RAID 1+0 or 0, with a 64 bit OS installed.	1x1.4 GHz CPU, 1 GB RAM
Windows platforms	2x quad-core Xeon, 3GHz, 8GB RAM, RAID 1+0 or 0, with a 64 bit OS installed.	Pentium 4 or equivalent at 2Ghz, 2GB RAM

Note: Be certain that your data reliability needs are met by a RAID 0 configuration before deploying a Splunk indexer on RAID 0.

- All configurations other than universal and light forwarder instances require at least the recommended hardware configuration.
- The minimum supported hardware guidelines are designed for personal use of Splunk.

Important: For all installations, including forwarders, a minimum of 2GB hard disk space for your Splunk installation is required **in addition to the space required for your indexes, if any**. Refer to this topic on estimating your index size requirements in this manual for some planning information.

Hardware requirements for universal and light forwarders

Recommended	Dual Core 1.5Ghz+ processor, 1GB+ RAM
Minimum	1.0 Ghz processor, 512MB RAM

Supported file systems

Platform	File systems
Linux	ext2/3/4, reiser3, XFS, NFS 3/4
Solaris	UFS, ZFS, VXFS, NFS 3/4
FreeBSD	FFS, UFS, NFS 3/4
Mac OS X	HFS, NFS 3/4
AIX	JFS, JFS2, NFS 3/4
HP-UX	VXFS, NFS 3/4
Windows	NTFS, FAT32

Note: If you run Splunk on a filesystem that is not listed above, Splunk may run a startup utility named `locktest` to test the viability of a filesystem for running Splunk. Locktest is a program that tests the start up process. If `locktest` runs and fails, the filesystem is not suitable for running Splunk.

Considerations regarding File Descriptors (FDs)

Splunk will allocate file descriptors for actively monitored files, forwarder connections, deployment clients, users running searches, and so on. Usually, the default ulimit on an OS is 1024. Your Splunk administrator should determine the correct level, but it should be at least 2056. Even if Splunk allocates just a single file descriptor for each of the activities above, it?s easy to see how a few hundred files being monitored, a few hundred forwarders sending data, a handful of very active users on top of reading/writing to/from the datastore can easily exhaust the default setting.

The more tasks your Splunk instance is doing, the more FDs it will need, so you should increase the ulimit value if you start to see your instance run into problems with low FD limits.

Considerations regarding NFS

NFS is usually a poor choice for Splunk indexing activity, for reasons of performance, resilience, and semantics. In environments with very high bandwidth, very low latency links, that are kept highly reliable, it can be an appropriate choice. Typically, this is a SAN (Storage Area Network) accessed via

the NFS protocol, an unusual choice for SANs but sometimes done.

"Soft" NFS mounts are not supported. Only "hard" NFS mounts can be reliable with Splunk.

Attribute caching should not be disabled. If you have other applications which require disabling or reducing attribute caching, a seperate mount with attribute caching enabled should be provided to Splunk.

Note: On FreeBSD, mounting as nullfs is not supported.

Considerations regarding solid state drives

Solid state drives (SSDs) gain most of their performance through read operations. Splunk relies on fast disk write performance in order to index data with low latency. SSDs do not provide a significant write-speed advantage in Splunk over fast conventional hard drives.

Supported server hardware architectures

32 and 64-bit architectures are supported for some platforms. See the download page page for details.

Unix/Linux file system permissions

The user running Splunk should have full permission to $SPLUNK_HOME and $SPLUNK_DB directories. Avoid changing the default umask for the user running Splunk. This can result in permission issues.

Components of a Splunk deployment

Splunk is simple to deploy by design. By using a single software component and easy to understand configurations, Splunk can coexist with existing infrastructure or be deployed as a universal platform for accessing machine data.

The simplest deployment is the one you get by default when you install Splunk: indexing and searching on the same server. Data comes in from the sources you've configured, and you log into Splunk Web or the **CLI** on this same server to search, monitor, alert, and report on your machine data.

Depending on your needs, you can also deploy components of Splunk on different servers to address your load and availability requirements. This section introduces the types of components. For a more thorough introduction, see the Distributed Deployment manual, particularly the topic, "Scale your deployment: Splunk components".

Indexer

Splunk **indexers**, or index servers, provide indexing capability for local and remote data and host the primary Splunk data store, as well as Splunk Web. Refer to "How indexing works" in the Admin manual for more information.

Search peer

A search peer is an indexer that services requests from **search heads** in a **distributed search** deployment. Search peers are also sometimes referred to as **indexer nodes**.

Search head

A **search head** is a Splunk instance configured to distribute searches to indexers, or **search peers**. Search heads can be either **dedicated** or not, depending on whether they also perform indexing. Dedicated search heads don't have any indexes of their own (other than the usual internal indexes). Instead, they consolidate results originating from remote search peers.

See "About distributed search" in the Distributed Deployment Manual to learn how to configure a search head to search across a pool of indexers.

Forwarder

Forwarders are Splunk instances that forward data to remote indexers for indexing and storage. In most cases, they do not index data themselves. Refer to the "About forwarding and receiving" topic in the Distributed Deployment manual for additional information on forwarders.

Deployment server

Both indexers and forwarders can also act as **deployment servers**. A deployment server distributes configuration information to running instances of Splunk via a push mechanism which is enabled through configuration. Refer to "About deployment server" for additional information.

Functions at a glance

Functions	Indexer	Search head	Forwarder	Deployment server
Indexing	x			
Web	x			
Direct search		x		
Forward to indexer			x	
Deploy configurations	x		x	x

Hardware capacity planning for your Splunk deployment

Splunk is a very flexible product that can be deployed to meet almost any scale and redundancy requirement. However, that doesn't remove the need for care and planning. This article discusses high level considerations for Splunk deployments, including sizing issues.

After you've worked through the general layout of your Splunk search topology, the other sections in this document can explain more thoroughly how to implement them, along with the formal Admin guide for Splunk.

Reference hardware

Let's consider a common, commodity hardware server as our standard:

- Intel x86-64-bit chip architecture
- Standard Linux or Windows 64-bit distribution
- 2 CPU, 4 core per CPU, 2.5-3Ghz per core
- 8GB RAM
- 4x300GB SAS hard disks at 10,000 rpm each in RAID 10
    - ♦ capable of 800 IO operations / second (IOPS)
- standard 1Gb Ethernet NIC, optional 2nd NIC for a management network

For the purposes of this discussion this will be our single server unit. Note that the only exceptional item here is the disk array. Splunk is often constrained by disk I/O first, so always consider that first when selecting your hardware.

Performance checklist

The first step to deciding on a reference architecture is sizing - can your Splunk handle the load? For the purposes of this guide we assume that managing forwarder connections and configurations (but not their data!) to be free. Therefore we need to look at index volume and search load.

Question 1: Do you need to index more than 2GB per day?

Question 2: Do you need more than 2 concurrent users?

If the answer to both questions is '**NO'** then your Splunk instance can safely share one of the above servers with other services, with the caveat that Splunk be allowed sufficient disk I/O on the shared box. If you answered yes, continue.

Question 3: Do you need to index more than 100GB per day?

Question 4: Do you need to have more than 4 concurrent users?

If the answer to both questions is '**NO'**, then a single dedicated Splunk server of our reference architecture should be able to handle your workload.

Question 5: Do you need more than 500GB of storage?

At a high level, total storage is calculated as follows:

```
daily average rate x retention policy x 1/2
```

You can generally safely use this simple calculation method. If you want to base your calculation on the specific type(s) of data that you'll be feeding into Splunk, you can use the method described in "Estimate your storage requirements" in this manual.

Splunk can generally, including indexes, store raw data at approximately half the original size thanks to compression. Given allowances for operating system and disk partitioning, that suggests about 500GB of usable space. In practical terms, that's ~6 months of fast storage at 5GB/day, or 10 days at 100GB/day.

If you need more storage, you can either opt for more local disks for fast access (required for frequent searching) or consider attached or network storage (acceptable for occasional searching). Low-latency connections over NFS or CIFS are acceptable for searches over long time periods where instant search returns can be compromised to lower cost per GB. Shares mounted over WAN

connections and standby storage such as tape are never acceptable.

Beyond 100GB per day

If you have requirements greater than 100GB/day or 4 concurrent users, you'll want to leverage Splunk's scale-out capabilities. That involves using **distributed search** to run searches in parallel across multiple indexers at once, and possibly load balancing the incoming data with load-balanced Splunk forwarders.

While Splunk does continue to scale linearly across multiple indexers, at this scale other considerations often become important. Larger user counts, more forwarders, and more scheduled searches begin to overshadow indexing throughput in your deployment design. Also, at this scale it is very likely that you will have high availability or redundancy requirements, which are covered in greater detail in "High availability reference architecture".

Dividing up indexing and searching

At daily volumes above 100GB/day, it makes sense to slightly modify our reference hardware to reflect the differing needs of indexers and search heads. Dedicated search heads do not need disk I/O, nor much local storage. However they are far more CPU bound than indexers. Therefore we can change our recommendations to:

Dedicated search head

- Intel 64-bit chip architecture
- Standard Linux or Windows 64-bit distribution
- 4 CPU, 4 core per CPU, 2.5-3Ghz per core
- 4GB RAM
- 2 300GB SAS hard disks at 10,000 rpm each in RAID 0
- standard 1Gb Ethernet NIC, optional 2nd NIC for a management network

Given that a search head will be CPU bound, if fewer, more performant servers are desired, adding more and faster CPU cores is best.

Note: The guideline of 1 core per active user still applies. Don't forget to account for scheduled searches in your CPU allowance as well.

Indexer

- Intel 64-bit chip architecture
- Standard Linux or Windows 64-bit distribution

- 2 CPU, 4 core per CPU, 2.5-3Ghz per core
- 8GB RAM
- 8 300GB SAS hard disks at 10,000 rpm each in RAID 10
    - ♦ capable of 1200 IO operations / second (IOPS)
- standard 1Gb Ethernet NIC, optional 2nd NIC for a management network

The indexers will be busy both writing new data and servicing the remote requests of search heads. Therefore disk I/O is the primary bottleneck.

At these daily volumes, likely local disk will not provide cost effective storage for the time frames that speedy search is desired, suggesting fast attached storage or networked storage. While there are too many types of storage to be prescriptive, here are guidelines to consider:

- indexers do many bulk reads
- indexers do many disk seeks

Therefore...

- more disks (specifically, more spindles) are better
- total throughput of the entire system is important, but...
- disk to controller ratio should be higher, similar to a database

Ratio of indexers to search heads

Technically, there is no practical Splunk limitation on the number of search heads an indexer can support, or the number of indexers a search head can search against. However systems limitations suggest a ratio of approximately 8 to 1 for most use cases. That is a rough guideline however; if you have many searchers compared to your total data volume, more search heads make sense, for example. In general, the best use of a separate search head is to populate summary indexes. This search head will then act like an indexer to the primary search head that users log into.

Accommodating many simultaneous searches

A common question for a large deployment is: how do I account for many concurrent users? Let's take as an example a system that may have at peak times 48 concurrent searches. The short answer is that we can accommodate 48 simultaneous searches on a cluster of indexers and search heads where each machine has enough RAM to prevent swapping. Assuming that each search takes 200MB of RAM per system, that is roughly 10GB additional RAM (beyond indexing requirements). This is because CPU will degrade gracefully with more

concurrent jobs but once the working set of memory for all processes exceeds the physical RAM, performance drops catastrophically with swapping.

The caveat here is that a search's run time will be longer in proportion to the number of free cores when no searches were running. For example, suppose the indexers were doing nothing before the searches arrived and have 8 cores each. Suppose the first (of identical searches) takes 10s to complete. Then the first 8 searches will each take 10s to complete since there is no contention. However, since there are only 8 cores, if there are 48 searches running, each search will take 48/8 = 6x longer than if only 1-8 searches were running. So now, every search takes ~1 minute to complete.

This leads to the observation that the most important thing to do here is add indexers. Indexers do the bulk of the work in search (reading data off disk, decompressing it, extracting knowledge and reporting). If we want to return to the world of 10s searches, we use 6 indexers (one search head is probably still fine, though it may be appropriate to set aside a search head for summary index creation) and searches 1-8 now take 10/6 = 1.6s and with 48 searches, each takes 10s.

Unfortunately, the system isn't typically idle before searches arrive. If we are indexing 150 GB/day, at peak times, we probably are using 4 of the 8 cores doing indexing. That means that the first 4 searches take 10s, and having 48 searches running takes 48/4 = 12x longer, or 2 min to complete each.

Now one might say: let me put sixteen cores per indexer rather than eight and avoid buying some machines. That makes a little bit of sense, but is not the best choice. The number of cores doesn't help searches 1-16 in this case; they still take 10s. With 48 searches, each search will take 48/16 = 3x longer, which is indeed better than 6x. However, it's usually not too much more expensive to buy two 8 core machines, which has advantages: the first few searches will now just take 5s (which is the most common case) and we now have more aggregate I/O capacity (doubling the number of cores does nothing for I/O, adding servers does).

The lesson here is to add indexers. Doing so reduces the load on any system from indexing, to free cores for search. Also, since the performance of almost all types of search scale with the number of indexers, searches will be faster, which mitigates the effect of slowness from resource sharing. Additionally making every search faster, we will often avoid the case of concurrent searches with concurrent users. In realistic situations, with hundreds of users, each user will run a search every few minutes, though not at the exact same time as other users. By reducing the search time by a factor of 6 (by adding more indexers), the

concurrency factor will be reduced (not necessarily by 6x, but by some meaningful factor). This in turn, lowers the concurrency related I/O and memory contention.

Summary of performance recommendations

Daily Volume	Number of Search Users	Recommended Indexers	Recommended Search Heads
< 2GB/day	< 2	1, shared	N/A
2GB/day to 100GB/day	up to 4	1	N/A
200GB/day	up to 8	2	1
300GB/day	up to 12	3	1
400GB/day	up to 8	4	1
500GB/day	up to 16	5	2
1TB/day	up to 24	10	2
20TB/day	up to 100	100	24
60TB/day	up to 100	300	32

Note that these are approximate guidelines only. You should feel free to modify based on the discussion here for your specific use case, and to contact Splunk for more guidance if needed.

Performance considerations

Splunk has three primary roles - indexer, searcher and forwarder. In many cases a single Splunk instance performs both indexing and searching. Although a Splunk indexer can also perform forwarding, in most cases, it makes more sense to use a separate Splunk instance, the **universal forwarder**, to handle forwarding. All roles have their own performance requirements and bottlenecks.

- Indexing, while relatively resource inexpensive, is often disk I/O bound.
- Searching can be both CPU and disk I/O bound.
- Forwarding uses few resources and is rarely a bottleneck.

As you can see, disk I/O is frequently the limiting factor in Splunk performance. It deserves extra consideration in your planning. That also makes Splunk a poor virtualization candidate, unless dedicated disk access can be arranged.

CPU

- Allow 1 CPU core for every 1MB/s of indexing volume
- Allow 1 CPU core for Splunk's optimization routines for every 2MB/s of indexing volume
- Allow 1 CPU per active searcher (be sure to account for scheduled searches)

Disk I/O

- Assume 50 Iopps per 1 MB/s of indexing volume
- Allow 50 Iopps for splunk's optimize routines
- Allow 100 Iopps per **search**, or an average of 200 Iopps per **search user**

Memory

- allow 200-300MB for indexing
- allow 500MB per concurrent search user
- allow 1GB for the operating system to accommodate OS caching

Total storage

- Allow 15% overhead for OS and disk partitioning
    - ♦ On this system there is ~500GB of usable storage
- Conservatively Splunk can, including indexes, compress original logs by ~50%
    - ♦ Compression rates vary based on the data

Based on these estimates, this machine **will be disk I/O bound** if there are too many active users or too many searches per user. That is the most likely limitation for this hardware, possibly followed by CPU if the searches are highly computational in nature, such as many uses of *stats* or *eval* commands in a single search.

Applied performance

With the information above, it is possible to estimate required hardware for most Splunk use cases by considering the following:

- The amount of daily indexed volume (disk I/O, CPU)
- The required retention period (total storage)
- The number of concurrent search users (disk I/O, CPU)

Although not all search users consume the same amount of resources, consider these very rough guidelines:

- Dashboard-heavy users trigger many searches at once
- Dashboards also suggest many scheduled searches
- Searching for rare events across large datasets (for example, **all time**) is disk I/O intensive
- Calculating summary information is CPU intensive
    - ♦ If done over long time intervals, can also be disk I/O intensive
- Alerts and scheduled searches run even if no one sees their results

What does that mean in real life?

- Executive users with many dashboards and summaries require both CPU and disk I/O
- Operations users searching over recent and small datasets require less resources
- Forensic and compliance users searching over long timeframes require disk I/O
- Alerting and scheduled searches over short timeframes are inexpensive; over long timeframes potentially very expensive.

Answers

Have questions? Visit Splunk Answers to see what questions and answers other Splunk users had about hardware and Splunk.

High availability reference architecture

Splunk provides the flexibility and capability to handle machine data for any type of computing environment, including the most stringent needs of medium and large enterprises. In some environments, maintaining data integrity and high availability can be of critical importance.

How you define high availability, and the approach you take to implement it, will vary greatly according to the needs of your particular business and the state of your existing system. This topic will help you make the right decisions about how best to deploy Splunk to promote a highly available, highly reliable system. It does not attempt to dictate any single approach to high availability. Rather, it offers a starting point for planning an approach that suits your enterprise.

As part of planning a highly available Splunk deployment, you must also take into account all aspects of your existing system - not only its components and topology, but also its overall reliability and availability. The specifics of your current system will determine how you integrate Splunk into it.

Before reading this topic, you should already be familiar with Splunk deployments and components, as described in "Distributed Splunk overview".

Note: This topic is intended for planning purposes only. It is not meant to serve as a detailed implementation guide. If you want to implement a high availability Splunk deployment, contact Splunk Professional Services for guidance.

The elements of a high availability architecture

Splunk collects data and it queries data. To implement end-to-end Splunk availability, you need to consider both functions.

If you are using Splunk forwarders in a load-balanced configuration, in which you send data alternately to multiple Splunk indexers in a group, then you already have high availability on the **data collection** side of Splunk. If one indexer goes down, the forwarder will automatically start sending the data to the other indexers in the load-balanced group.

To provide high availability for Splunk's **data querying** capability, you must maintain availability for:

- The indexer(s)
- The indexed data

The rest of this topic describes ways to maintain high availability for querying Splunk data.

Depending on your requirements, you might also need to consider availability of other components of your Splunk deployment, such as search heads and forwarders. You must also provide high availability for non-Splunk (but Splunk-dependent) aspects of your system, such as your data sources, hardware, and network.

There are two basic choices for implementing high availability for Splunk indexers and data:

- Use a highly reliable storage system
- Use a mirrored cluster of Splunk indexers

High reliability storage

There are a number of ways that you can use an underlying storage system to promote high availability for Splunk. The exact architecture you implement will depend on your existing environment and specific needs.

For example, in a typical SAN-based architecture, you could install your Splunk indexers directly on the SAN and then mount the Splunk volumes on server nodes. If a node goes down, you can remount its volume on another node. The new node takes on the identity of the failed node, with the same configurations and access to the same set of indexed data. You just need to point your search head at that node in place of the old one. You can further configure your SAN to attain the level of redundancy you require.

Data replication across indexer clusters

Another way to achieve high availability of both the indexed data and the indexing/searching capabilities is to create primary and secondary clusters of mirrored indexers. If an indexer in the primary cluster fails, you can reconfigure forwarders and search heads to point to its mirror on the secondary cluster.

Here's an example of this strategy. It starts by showing two forwarders using load balancing to distribute data to the indexers in the primary cluster:

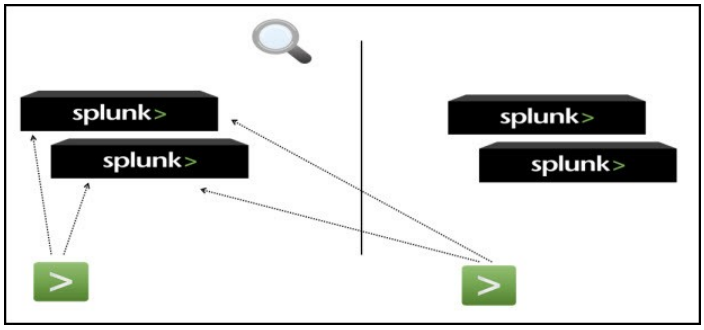

The primary indexers index the data locally and also forward the raw (unindexed) data onwards to secondary indexers, which then index the data a second time:

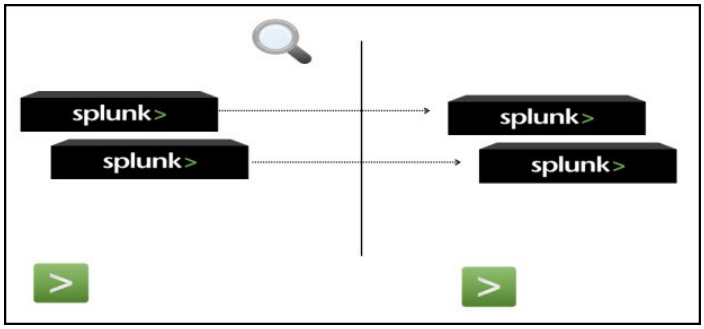

You now have copies of the indexed data in two places. Each indexer in the secondary cluster contains an exact copy of the data on its corresponding indexer in the primary cluster. You can search against either the primary or the secondary cluster:

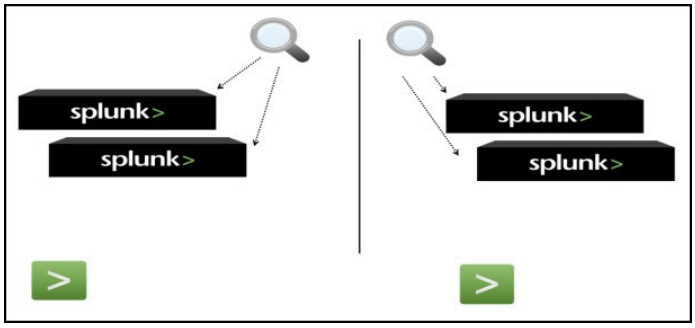

If one of the indexers in the primary cluster goes down, the forwarders' load-balancing capability means that they will automatically start sending all their data to the remaining indexer(s) in that cluster. Those indexers will continue to send copies of their data on to their mirrored instances in the secondary cluster:

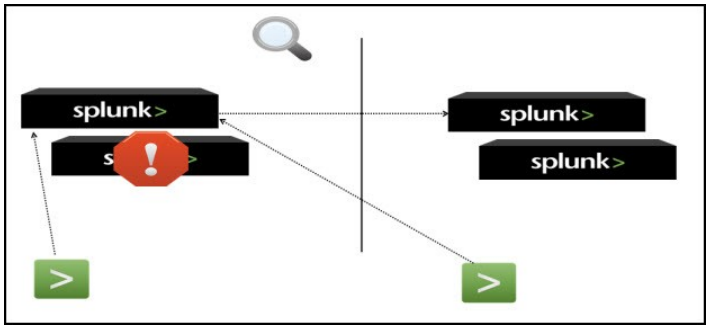

You can continue to search across the full set of data. You just redirect the search head(s) to point to the secondary instance of the downed indexer. At the same time, the search head can continue to point to the remaining indexer in the primary cluster. Alternatively, you can redirect the search head to point exclusively to the secondary indexers. In either case, searching continues across the entire set of data:

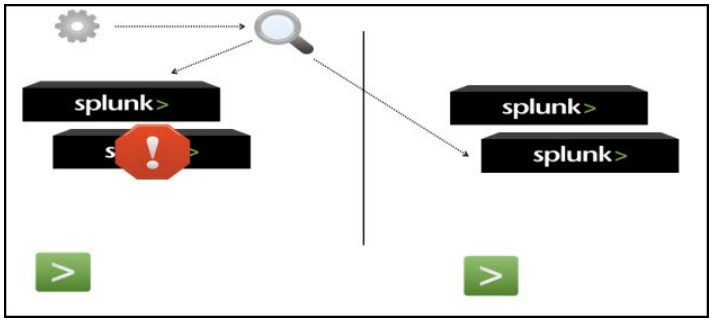

There are many ways you can implement specific aspects of this architecture. For guidance, talk with Splunk Professional Services.

This second solution has the advantage that it depends less on the capabilities of your underlying storage system. On the downside, it requires double the hardware (since you're doubling the indexers), as well as a license for twice the indexing volume (since you're indexing everything twice).

Estimate your storage requirements

This topic describes how to estimate the size of your Splunk index, so that you can plan your storage capacity requirements.

When Splunk indexes your data, it creates two main types of files: the rawdata file containing the original data in compressed form and the index files that point to this data. (It also creates a few metadata files, which don't consume much space.) With a little experimentation, you can estimate how much index disk space you will need for a given amount of incoming data.

Typically, the compressed rawdata file is approximately 10% the size of the

incoming, pre-indexed raw data. The associated index files range in size from approximately 10% to 110% of the rawdata file. This value is affected strongly by the number of unique terms in the data. Depending on the data's characteristics, you might want to tune your segmentation settings, as described in "About segmentation".

The best way to get an idea of your space needs is to experiment by indexing a representative sample of your data, and then checking the sizes of the resulting directories in `defaultdb`.

On *nix systems, follow these steps

Once you've indexed your sample:

1. Go to `$SPLUNK_HOME/var/lib/splunk/defaultdb/db`.

2. Run `du -shc hot_v*/rawdata` to determine how large the compressed persisted raw data is.

This is the persisted data to which the items in the index point. Typically, this file's size is about 10% of the size of the sample data set you indexed.

3. Run `du -ch hot_v*` and look at the last `total` line to see the size of the index.

4. Add the values you get together.

On Windows systems, follow these steps

1. Download the du utility from Microsoft TechNet.

2. Extract `du.exe` from the downloaded ZIP file and place it into your `%SYSTEMROOT%` folder.

Note: You can also place it anywhere in your `%PATH%`.

3. Open a command prompt.

4. Once there, go to `%SPLUNK_HOME%\var\lib\splunk\defaultdb\db`.

5. Run `del %TEMP%\du.txt & for /d %i in (hot_v*) do du -q -u %i\rawdata | findstr /b "Size:" >> %TEMP%\du.txt`.

6. Open the `%TEMP%\du.txt file`. You will see `Size: n`, which is the size of each `rawdata` directory found.

7. Add these numbers together to find out how large the compressed persisted raw data is.

8. Next, run `for /d %i in (hot_v*) do dir /s %i`, the summary of which is the size of the index.

9. Add this number to the total persistent raw data number.

This is the total size of the index and associated data for the sample you have indexed. You can now use this to extrapolate the size requirements of your Splunk index and `rawdata` directories over time.

Answers

Have questions? Visit Splunk Answers to see what questions and answers other Splunk users had about data sizing.

Splunk architecture and processes

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

Processes

A Splunk server runs two processes (installed as services on Windows systems) on your host, `splunkd` and `splunkweb`:

- **`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 passing through a pipeline. Pipelines can pass data to one another via **queues**. `splunkd` supports a command line interface for searching and viewing results.

- **splunkweb** is a Python-based application server based on CherryPy that provides the Splunk Web user interface. It allows users to search and navigate IT data stored by Splunk servers and to manage your Splunk deployment through a Web interface.

splunkweb and splunkd can both communicate with your Web browser via REST:

- splunkd also runs a Web server on port 8089 with SSL/HTTPS turned on by default.
- splunkweb runs a Web server on port 8000 without SSL/HTTPS by default.

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

Read information on other Windows third-party binaries distributed with Splunk.

Splunk and Windows in Safe Mode

Neither the splunkd nor the splunkweb services will start if Windows is in Safe Mode. Additionally, if you attempt to start Splunk from the Start Menu while in Safe Mode, Splunk does not alert you to the fact that its services are not running.

Additional processes for Splunk on Windows

On Windows instances of Splunk, in addition to the two services described above, there are additional processes that are used by the data inputs you create on a Splunk instance. These scripted 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. It provides the command line interface (CLI) for the program, and allows you to start, stop, and configure Splunk, similar to the *nix splunk program.

Important: splunk.exe requires an elevated context to run because of how it controls the splunkd and splunkweb processes. Splunk might not function correctly if this executable is not given the appropriate permissions on your Windows system. This is not an issue if you install Splunk as the Local System user.

splunk-admon

`splunk-admon.exe` is spawned by `splunkd` whenever you configure an Active Directory (AD) monitoring input. `splunk-admon`'s purpose is to attach to the nearest available AD domain controller and gather change events generated by AD. Those change events are then stored in Splunk.

splunk-perfmon

`splunk-perfmon.exe` (new for version 4.2) runs when Splunk has been set up to monitor performance data on the local machine. This service 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-regmon

`splunk-regmon.exe` runs when a Registry monitoring input is configured in Splunk. This scripted input initially writes a baseline for the Registry as it currently exists (if desired), then monitors changes to the Registry over time. Those changes come back into Splunk as searchable events.

splunk-winevtlog

This utility is used to test defined event log collections, and can output events as they are collected for investigation. Splunk has a Windows event log input processor built into the engine.

splunk-wmi

When you configure a performance monitoring, event log or other input against a remote computer, this program starts up. Depending on how the input is configured, either it attempts to attach to and read Windows event logs as they come over the wire, or it executes a Windows Query Language (WQL) query against the WMI provider on the specified remote machine(s). Those events are then stored in Splunk.

Architecture diagram

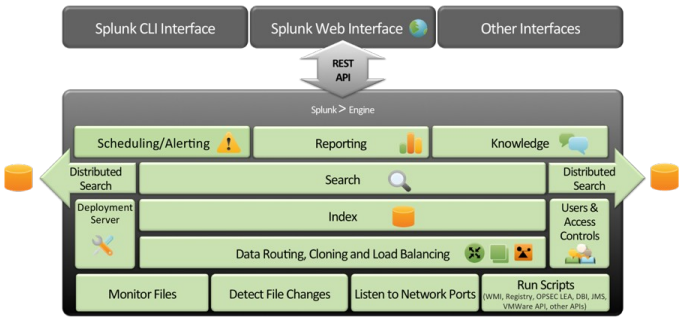

About Splunk licenses

Splunk takes in data from sources you designate and processes it so that you can analyze it in Splunk. We call this process "indexing". For information about the exact indexing process, refer to "What Splunk does with your data" in the Getting Data In Manual.

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

For more information about Splunk licenses, begin by reading:

- "How Splunk licensing works" in the Admin Manual.
- "Types of Splunk licenses" in the Admin Manual.
- "More about Splunk Free" in the Admin Manual.

Before you install

Before you install Splunk, be sure to review the system requirements and ensure that you've downloaded the right installation package for your system from the Splunk download page.

If you're installing Splunk on a Windows system, make sure that any anti-virus software is turned **off.** You can enable it again after the install is done, but be sure to configure it so that it doesn't attempt to scan either the Splunk installation directory, or any file accesses by Splunk processes (basically, any process that begins with "splunk".)

If you're upgrading from an earlier version of Splunk, review the information in "About upgrading to 4.3: READ THIS FIRST" before proceeding.

Accessibility options

Splunk is dedicated to maintaining and enhancing its accessibility and usability for users of assistive technology (AT), both in accordance with section 508, and in terms of best usability practices. The Splunk 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 has been 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 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.) Refer to "How to disable real-time search" in the User 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.

Step by step installation procedures

Choose your platform

Choose from the list below for detailed installation procedures:

- Windows
- Windows commandline instructions
- Linux
- Solaris
- MacOS
- FreeBSD
- AIX
- HP-UX

Install on Windows

This topic describes the procedure for installing Splunk on Windows with the Graphical User Interface (GUI)-based installer. More options (such as silent installation) are available if you install from the command line.

Important: Running the 32-bit version of Splunk for Windows on a 64-bit platform is not recommended. If you attempt to run the 32-bit installer on a 64-bit system, the installer will warn you of this.

If you can run 64-bit Splunk on 64-bit hardware, we strongly recommend it. The performance is greatly improved over the 32-bit version.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding.

Splunk for Windows and anti-virus software

Splunk's indexing subsystem requires lots of disk I/O bandwidth. Any software with a device driver that intermediates between Splunk and the operating system can rob Splunk of processing power, causing slowness and even an unresponsive system. This includes anti-virus software.

It's extremely important to configure such software to avoid on-access scanning of Splunk installation directories and processes, before starting a Splunk installation.

Choose the user Splunk should run as

When you run the Splunk Windows installer, you are given the option to select the user that Splunk will run as.

If you install as the Local System user, Splunk will have access to all of the important information on your local machine. However, the Local System user has no privileges on other Windows machines by design.

If you intend to do any of the following things, you must give Splunk a domain account:

- read Event Logs remotely
- collect performance counters remotely
- read network shares for log files
- enumerate the Active Directory schema using Active Directory monitoring

The domain account you use must be a member of the Active Directory domain you wish to monitor. It must also be a member of the local Administrators group.

Note: Splunk might not function properly if the Splunk user is not a local administrator on computers running versions of Windows prior to Windows Server 2008.

If you're not sure which account to run Splunk under, speak with your Windows domain administrator about the best way to proceed. If you are the domain administrator, then review "Considerations for deciding how to monitor remote Windows data" in the Getting Data In Manual for additional information on how to configure your Splunk user with the access it needs.

Important: If you decide to change the user Splunk runs as after you have installed, you must ensure that the new account:

- Has the necessary resource access rights.
- Is a member of the machine's local Administrators group.
- Has "Full Control" permissions to the `%SPLUNK_HOME%` directory and all its subdirectories.

Managed service accounts on Windows Server 2008 and Windows 7

If you run WIndows Server 2008, Windows Server 2008 R2 or Windows 7, and your domain is properly configured or has at least one Windows Server 2008 R2 domain controller present, you can use managed server accounts (MSA) on your Splunk instance.

The major benefits of using a MSA are:

- Increased security from the isolation of accounts for services.
- Administrators no longer need to manage the credentials or administer service principle names (SPNs).
- Administrators can delegate the administration of these accounts to non-administrators.

Some important things to understand before installing Splunk under a MSA are:

- The MSA requires the same permissions as a domain account on the machine that runs Splunk.
- The MSA must be a local administrator on the machine that runs Splunk.
- You cannot use the same account on different computers, as you would with a domain account.
- You must correctly configure and install the MSA on the machine that runs Splunk **before** you install Splunk on the machine. For information and instructions on how to do this, review "Service Accounts Step-by-Step Guide" (http://technet.microsoft.com/en-us/library/dd548356%28WS.10%29.aspx) on MS Technet.

To install Splunk using a managed service account:

1. Ensure that the MSA you plan to use is properly installed and configured.

Important: The MSA must have appropriate rights configured for the Windows resources you need to monitor, and must also be a local Administrator on the machine that runs Splunk.

2. Install Splunk from the command line as the "Local System" user.

Important: You **must** use the `LAUNCHSPLUNK=0` flag to keep Splunk from starting after installation is completed.

3. After installation is complete, use the Windows Explorer or the `ICACLS` command line utility to grant the MSA "Full Control" permissions to the Splunk installation directory and all its sub-directories.

4. Follow the instructions in the topic "Correct the user selected during Windows installation" in this manual. In this instance, the correct user is the MSA you configured prior to installing Splunk.

Important: You **must** append a dollar sign ($) to the end of the username when completing Step 4 in order for the MSA to work correctly. For example, if the MSA is `SPLUNKDOCS\splunk1`, then you must enter `SPLUNKDOCS\splunk1$` in the appropriate field in the properties dialog for the service. You must do this for both the `splunkd` and `splunkweb` services.

5. Make sure that the MSA has the "**Log on as a service**" right.

6. Restart Splunk. Splunk will run as the MSA configured above, and will have access to all data the MSA has access to.

Security and remote access considerations

In the interests of security, Splunk strongly recommends that you take the following steps when assigning rights for the Splunk user:

- Create a domain group that the Splunk user will be a member of.
- Place the Splunk user into this group.
- Then, place that group into local groups on member servers or workstations.

This helps maintain security integrity and makes it a lot easier to control access in the event of a breach or site-wide security change.

Minimum permissions required to run Splunk as a user other than Local System

The following is a list of the minimum **local** permissions required for the `splunkd` and `splunkweb` services, when Splunk is installed using a user. Depending on the sources of data you need to access, the Splunk user might need a significant amount of additional permissions.

Required basic permissions for the `splunkd` service:

- Full control over Splunk's installation directory

- Read access to any flat files you want to index

Required Local Security Policy user rights assignments for the `splunkd` service:

- 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

Required basic permissions for the `splunkweb` service:

- Full control over Splunk's installation directory

Required Local Security Policy user rights assignments for the `splunkweb` service:

- Permission to log on as a service

Note: These permissions are not required when Splunk runs as the Local System account.

Use Group Policy to assign user rights to multiple machines

If you want to assign the policy settings shown above to a number of workstations and servers in your AD domain or forest, you can define a Group Policy object (GPO) with these specific rights, and deploy that GPO across the domain or forest using the **Domain Security Policy** Microsoft Management Console (MMC) snap-in. For domain controllers, use the **Domain Controller Security Policy** snap-in.

Once you've created and enabled the GPO, the workstations and member servers in your domain will pick up the changes either during the next scheduled AD replication cycle (usually every 2-3 hours) or at the next boot time.

Remember that identical Local Security Policy user rights defined on a workstation or member server are overridden by the rights inherited from a GPO, and you can't change this setting. If you wish to retain previously existing rights that are explicitly defined through Local Security Policy on your member servers, they'll also need to be assigned within the GPO.

If you accidentally specify the wrong user the first time you install

If you specified the wrong user during the installation procedure, you'll see two popup error dialogs telling you this. Complete the installation and then use these instructions to switch to the correct user. You must not start Splunk before doing this.

Troubleshoot permissions issues

The rights described above are the rights that the `splunkd` and `splunkweb` services specifically invoke. Other rights might be required, depending on your usage and what data you want to access. Additionally, many user rights assignments and other Group Policy restrictions can prevent Splunk from running. If you have issues, consider using a tool such as Process Monitor to troubleshoot your environment.

You can use the `GPRESULT` command line tool or the Group Policy Management Console (GPMC) to troubleshoot issues related to GPO application in your enterprise. As a last resort, you can revert to running the `splunkd` service under a domain administrator or equivalent account.

Install Splunk via the GUI installer

The Windows installer is an MSI file.

1. To start the installer, double-click the `splunk.msi` file.

The Welcome panel is displayed.

2. To begin the installation, click **Next**.

Note: On each panel, you can click **Next** to continue, **Back** to go back a step, or **Cancel** to close the installer.

The licensing panel is displayed.

3. Read the licensing agreement and select "I accept the terms in the license agreement". Click **Next** to continue installing.

The **Destination Folder** panel is displayed.

Note: Splunk is installed by default into `\Program Files\Splunk` on the system drive. Splunk's installation directory is referred to as `$SPLUNK_HOME` or

`%SPLUNK_HOME%` throughout this documentation set.

4. Click **Change...** to specify a different location to install Splunk, or click **Next** to accept the default value.

The **Logon Information** panel is displayed.

Splunk installs and runs two Windows services, `splunkd` and `splunkweb`. These services will be installed and run as the user you specify on this panel. You can choose to run Splunk with Local System credentials, or provide a specific account. That account should have local administrator privileges, plus appropriate domain permissions if you are collecting data from other machines.

The user Splunk runs as must have permissions to:

- Run as a service.
- Read whatever files you are configuring it to monitor.
- Collect performance or other WMI data.
- Write to Splunk's directory.

Note: If you install as the Local System user, some network resources will not be available to the Splunk application. Additionally, WMI remote authentication will not work; this user has null credentials and Windows servers normally disallow such connections. Only local data collection with WMI will be available. Contact your systems administrator for advice if you are unsure what account to specify.

5. Select a user type and click **Next**.

If you specified the local system user, proceed to step 7. Otherwise, the **Logon Information: specify a username and password** panel is displayed.

6. Specify a username and password to install and run Splunk and click **Next**.

Note: This must be a valid user in your security context. Splunk cannot start without a valid username and password.

The pre-installation summary panel is displayed.

7. Click **Install** to proceed.

The installer runs and displays the **Installation Complete** panel.

Caution: If you specified the wrong user during the installation procedure, you will see two popup error windows explaining this. If this occurs, Splunk installs itself as the local system user by default. Splunk will 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.

8. If desired, check the boxes to **Launch browser with Splunk** and **Create Start Menu Shortcut** now. Click **Finish**.

The installation completes, Splunk starts, and Splunk Web launches in a supported browser if you checked the appropriate box.

Note: The first time you access Splunk Web after installation, login with the default username `admin` and password `changeme`.

Launch Splunk in a Web browser

To access Splunk Web after you start Splunk on your machine, you can either:

- Click the Splunk icon in **Start > Programs > Splunk**

or

- Open a Web browser and navigate to `http://localhost:8000`.

Log in using the default credentials: username: `admin` and password: `changeme`.

The first time you log into Splunk successfully, you'll be prompted right away to change your password. You can do so by entering a new password and clicking the **Change password** button, or you can do it later by clicking the **Skip** button.

Note: If you do not change your password, remember that anyone who has access to the machine can access your Splunk instance. Be sure to change the admin password as soon as possible and make a note of what you changed it to.

Avoid IE Enhanced Security pop-ups

If you're using Internet Explorer to access Splunk, add the following URLs to the allowed Intranet group or fully trusted group to avoid getting "Enhanced Security" pop-ups:

- `quickdraw.splunk.com`
- the URL of your Splunk instance

Change the Splunk Web or splunkd service ports

If you want the Splunk Web service or the splunkd service to use a different port, you can change the defaults.

To change the splunk web service port:

- Open a command prompt.
- Change to the `%SPLUNK_HOME%\bin` directory.
- Type in `splunk set web-port ####` and press Enter.

To change the splunkd port:

- Open a command prompt, if one isn't already.
- Change to the `%SPLUNK_HOME%\bin` directory.
- Type in `splunk set splunkd-port ####` and press Enter.

Note: If you specify a port and that port is not available, or if the default port is unavailable, Splunk will automatically select the next available port.

Install or upgrade license

If you are performing a new installation of Splunk or switching from one license type to another, you must install or update your license.

Uninstall Splunk

To uninstall Splunk, use the **Add or Remove Programs** option in the **Control Panel**.

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

What's next?

Now that you've installed Splunk, you can find out what comes next, or you can review these topics in the Getting Data In Manual for information on adding Windows data to Splunk:

- 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 via the command line

This topic describes the procedures for installing Splunk on Windows using the command line.

Note: You can only run one Splunk instance per Windows host.

Important: Running the 32-bit version of Splunk for Windows on a 64-bit platform is not recommended. If you run the 32-bit installer on a 64-bit system, the installer will warn you about this.

If you can run 64-bit Splunk on 64-bit hardware, we strongly recommend it. The performance is greatly improved over the 32-bit version.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding.

In particular, be aware that changing the management or HTTP port during an upgrade is not supported.

Splunk for Windows and anti-virus software

Splunk's indexing subsystem requires lots of disk I/O bandwidth. Anti-virus software - or any software with a device driver that intermediates between Splunk and the operating system - can rob Splunk of processing power, causing slowness and even an unresponsive system.

It's extremely important to configure such software to avoid on-access scanning of Splunk installation directories and processes, before starting a Splunk installation.

Choose the user Splunk should run as

When you run the Splunk Windows installer, you are given the option to select the user that Splunk will run as.

If you install as the Local System user, Splunk will have access to all or nearly all of the important information on your local machine. However, the Local System user has no privileges on other Windows machines by design.

If you intend to do any of the following things, you must give Splunk a domain account:

- read Event Logs remotely
- collect performance counters remotely
- read network shares for log files
- enumerate the Active Directory schema using Active Directory monitoring

The domain account you use must also be a member of the local Administrators group. This is particularly important for installation on versions of Windows prior to Windows Server 2008 - failure to give the Splunk account access to the local Administrators group can cause Splunk to fail to function properly.

If you're not sure which account to run Splunk under, speak with your Windows domain administrator about the best way to proceed. If you are the domain administrator, then start off by using an account that has at most the permissions described here, and add rights as needed until you get the results you want.

Important: If you decide to change the user Splunk runs as after you have installed, you must ensure that the new user has the necessary resource access rights, and Full Control permissions to the entire `%SPLUNK_HOME%` directory.

Managed service accounts on Windows Server 2008 and Windows 7

If you run WIndows Server 2008, Windows Server 2008 R2 or Windows 7, and your domain is properly configured or has at least one Windows Server 2008 R2 domain controller present, you can use managed server accounts (MSA) on your Splunk instance.

The major benefits of using a MSA are:

- Increased security from the isolation of accounts for services.
- Administrators no longer need to manage the credentials or administer service principle names (SPNs).

- Administrators can delegate the administration of these accounts to non-administrators.

Some important things to understand before installing Splunk under a MSA are:

- The MSA requires the same permissions as a domain account on the machine that runs Splunk.
- The MSA must be a local administrator on the machine that runs Splunk.
- You cannot use the same account on different computers, as you would with a domain account.
- You must correctly configure and install the MSA on the machine that runs Splunk **before** you install Splunk on the machine. For information and instructions on how to do this, review "Service Accounts Step-by-Step Guide" (http://technet.microsoft.com/en-us/library/dd548356%28WS.10%29.aspx) on MS Technet.

To install Splunk using a managed service account:

1. Ensure that the MSA you plan to use is properly installed and configured.

Important: The MSA must have appropriate rights configured for the Windows resources you need to monitor, and must also be a local Administrator on the machine that runs Splunk.

2. Install Splunk from the command line as the "Local System" user.

Important: You **must** use the `LAUNCHSPLUNK=0` flag to keep Splunk from starting after installation is completed.

3. After installation is complete, use the Windows Explorer or the `ICACLS` command line utility to grant the MSA "Full Control" permissions to the Splunk installation directory and all its sub-directories.

4. Follow the instructions in the topic "Correct the user selected during Windows installation" in this manual. In this instance, the correct user is the MSA you configured prior to installing Splunk.

Important: You **must** append a dollar sign ($) to the end of the username when completing Step 4 in order for the MSA to work correctly. For example, if the MSA is `SPLUNKDOCS\splunk1`, then you must enter `SPLUNKDOCS\splunk1$` in the appropriate field in the properties dialog for the service. You must do this for both the `splunkd` and `splunkweb` services.

5. Make sure that the MSA has the "**Log on as a service**" right.

6. Restart Splunk. Splunk will run as the MSA configured above, and will have access to all data the MSA has access to.

Security and remote access considerations

In the interests of security, Splunk strongly recommends that you create and place the Splunk account into a domain group, and then place that group into local groups on member servers, when assigning rights for the Splunk account. This helps maintain security integrity and makes it a lot easier to control access in the event of a security breach or site-wide change.

The following is a list of the minimum **local** permissions required for the two Splunk services. Depending on the sources of data you need to access, the Splunk account may need a significant amount of additional permissions.

Required basic permissions for the `splunkd` service:

- Full control over Splunk's installation directory
- Read access to any flat files you want to index

Required Local Security Policy user rights assignments for the `splunkd` service:

- 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

Required basic permissions for the `splunkweb` service:

- Full control over Splunk's installation directory

Required Local Security Policy user rights assignments for the `splunkweb` service:

- Permission to log on as a service

Using Group Policy to assign user rights domain-wide

If you want to assign the policy settings shown above to all member servers in your AD domain, you can define a Group Policy object (GPO) for these specific rights and deploy that GPO across the domain or forest using the **Domain Security Policy** MMC snap-in (use the **Domain Controller Security Policy** snap-in for domain controllers). The member servers in your domain will pick up the changes either during the next scheduled AD replication cycle (usually every 2-3 hours), or at the next boot time.

Remember that identical Local Security Policy user rights defined on a member server are overwritten by the rights inherited from a GPO, and you can't change this setting. If you wish to retain previously existing rights defined on your member servers, they'll also need to be assigned within the GPO.

If you accidentally specify the wrong user the first time you install

If you specified the wrong user during the installation procedure, you'll see two popup error dialogs telling you this. Complete the installation and then use these instructions to switch to the correct user. You must not start Splunk before doing this.

Troubleshooting permissions issues

The rights described above are the rights that the `splunkd` and `splunkweb` services specifically invoke. Other rights may be required depending on your usage and what data you want to access. Additionally, many user rights assignments and other Group Policy restrictions can prevent Splunk from running. If you have issues, consider using a tool such as Process Monitor to troubleshoot your environment.

You can use the `GPRESULT` command line tool or the Group Policy Management Console (GPMC) to troubleshoot issues related to GPO application in your enterprise. As a last resort, you can revert to running the `splunkd` service under a domain administrator or equivalent account.

How to use the Microsoft Installer on the command line

You can install Splunk for Windows using the Microsoft Installer (MSI) on the command line by typing the following:

```
msiexec.exe /i Splunk.msi
```

This section lists the available flags for doing this, and provides a few examples of doing this in various configurations.

You can specify

- which Windows event logs to index
- which Windows registry hive(s) to monitor
- which Windows Management Instrumentation (WMI) information to pull
- the user Splunk runs as (be sure the user you specify has the appropriate permissions to access the content you want Splunk to index)
- an included application configuration for Splunk to enable (such as the Splunk light forwarder)
- whether or not Splunk should start up automatically when the installation is completed

Note: The first time you access Splunk Web after installation, log in with the default username `admin` and password `changeme`.

Supported flags

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

Important: The Splunk universal forwarder is a separate executable, with its own installation flags. Review the supported installation flags for the universal forwarder in "Deploy a Windows universal forwarder from the command line" in the Distributed Deployment Manual.

Flag	What it's for	Default
`INSTALLDIR="<directory_path>"`	Use this flag to specify directory to install. Splunk's installation directory is referred to as `$SPLUNK_HOME` or `%SPLUNK_HOME%` throughout this documentation set.	`C:\Program Files\Splunk`
`SPLUNKD_PORT=<port number>`	Use these flags to specify alternate ports for `splunkd` and `splunkweb` to use. **Note:** If you specify a port and that port is not available, Splunk will automatically select the next available port.	`8089`
`WEB_PORT=<port number>`	Use these flags to specify alternate ports for `splunkd` and `splunkweb` to use.	`8000`

	Note: If you specify a port and that port is not available, Splunk will automatically select the next available port.	

`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 should index a particular Windows event log:
	Application log
	Security log
	System log
	Forwarder log
	Setup log
	Note: You can specify multiple flags.
		`0` (off)

`REGISTRYCHECK_U=1/0` `REGISTRYCHECK_BASELINE_U=1/0`	Use this flag to specify whether or not Splunk 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 this flag to specify whether or not Splunk should index events from capture a baseline snapshot of the Windows Registry machine hive (`HKEY_LOCAL_MACHINE`). **Note:** You can set both of these at the	`0` (off)

	same time.	

`WMICHECK_CPUTIME=1/0`
 	
`WMICHECK_LOCALDISK=1/0`
 	
`WMICHECK_FREEDISK=1/0`
 	
`WMICHECK_MEMORY=1/0`
	Use these flags to specify which popular WMI-based performance metrics Splunk should index:
	
	CPU usage
	
	Local disk usage
	
	Free disk space
	
	Memory statistics
	
	Caution: If you need this instance of Splunk to monitor remote Windows instances over WMI, then you must also specify the `LOGON_USERNAME` and `LOGON_PASSWORD` installation flags. Splunk will not collect any remote WMI-based data that it does not have explicit access to. Read "Choose the user Splunk should run as" in the "Install on Windows" topic in this manual for additional information about the required credentials.
	
	Note: There are many more WMI-based metrics that Splunk can index. Review "Monitor WMI Data" in the Getting Data In Manual for specific information.
		`0` (off)

`LOGON_USERNAME="<domain\username>"`
 	
`LOGON_PASSWORD="<pass>"`
	Use these flags to provide domain\username and password information for the user that Splunk will run as. The `plunked` 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`."
	
	These flags are required if you want this Splunk installation to monitor any remote WMI-based data. Review
		none

	"Choose the user Splunk should run as" in the "Install on Windows" topic in this manual for additional information about which credentials to use.	
SPLUNK_APP="<SplunkApp>"	Use this flag to specify an included Splunk application configuration to enable for this installation of Splunk. Currently supported options for <SplunkApp> are: SplunkLightForwarder and SplunkForwarder. These specify that this instance of Splunk will function as a light forwarder or heavy forwarder, respectively. Refer to the "About forwarding and receiving" topic in the Distributed Deployment manual for more information. **Important:** The universal forwarder is not enabled from full Splunk; it is a separate downloadable executable, with its own installation flags. **Note:** If you specify either the Splunk forwarder or light forwarder here, you must also specify FORWARD_SERVER="<server:port>". To install Splunk with no applications at all, simply omit this flag.	none
FORWARD_SERVER="<server:port>"	Use this flag *only* when you are also using 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. **Important:** This flag requires that the SPLUNK_APP flag also be set.	none
DEPLOYMENT_SERVER="<host:port>"	Use this flag to specify a deployment server for pushing configuration updates. Enter the deployment server's name (hostname or IP address) and port.	none
LAUNCHSPLUNK=0/1	Use this flag to specify whether or not Splunk should start up automatically on system boot.	1 (on)

	Important: If you enable the Splunk Forwarder by using the `SPLUNK_APP` flag, Splunk is configured to start automatically, and this flag is ignored.	

Silent installation

To run the installation silently, add `/quiet` to the end of your installation command string. If your system is running UAC (which is sometimes on by default) you must run the installation as Administrator. To do this: when opening a cmd prompt, right click and select "Run As Administrator". Then use this cmd window to run the silent install command.

Examples

The following are some examples of using different flags.

Silently install Splunk to run as the Local System user

```
msiexec.exe /i Splunk.msi /quiet
```

Enable SplunkForwarder and specify credentials for the user Splunk will run as

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

Enable SplunkForwarder, enable indexing of the Windows System event log, and run the installer in silent mode

```
msiexec.exe /i Splunk.msi SPLUNK_APP="SplunkForwarder"
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.

Launch Splunk in a Web browser

To access Splunk Web after you start Splunk on your machine, you can either:

- Click the Splunk icon in **Start>Programs>Splunk**

or

- Open a Web browser and navigate to `http://localhost:8000`.

Log in using the default credentials: username: `admin` and password: `changeme` . Be sure to change the admin password as soon as possible and make a note of what you changed it to.

Now that you've installed Splunk, **what comes next?**

Avoid IE Enhanced Security pop-ups

To avoid IE Enhanced Security pop-ups, add the following URLs to the allowed Intranet group or fully trusted group in IE:

- quickdraw.splunk.com
- the URL of your Splunk instance

Install or upgrade license

If you are performing a new installation of Splunk or switching from one license type to another, you must install or update your license.

Uninstall Splunk

To uninstall Splunk, use the **Add or Remove Programs** option in the **Control Panel**.

You can also use `msiexec` from the command line.

What's next?

Now that you've installed Splunk, **what comes next?**

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

Install on Linux

You can install Splunk on Linux using RPM or DEB packages, or a tarball.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migation considerations before proceeding.

RedHat RPM install

To install the Splunk RPM in the default directory `/opt/splunk`:

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

To install Splunk in a different directory, use the `--prefix` flag:

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

To upgrade an existing Splunk installation that resides in /opt/splunk using the RPM:

```
rpm -U splunk_package_name.rpm
```

To upgrade an existing Splunk installation that was done in a different directory, use the `--prefix` flag:

```
rpm -U --prefix=/opt/existing_directory splunk_package_name.rpm
```

Note: If you do not specify with `--prefix` for your existing directory, rpm will install in the default location of `/opt/splunk`.

For example, to upgrade to the existing directory of `$SPLUNK_HOME=/opt/apps/splunk` enter the following:

```
rpm -U --prefix=/opt/apps splunk_package_name.rpm
```

If you want to automate your RPM install with kickstart, add the following to your kickstart file:

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

Note: The second line is optional for the kickstart file.

Debian DEB install

To install the Splunk DEB package:

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

Note: You can only install the Splunk DEB package in the default location, `/opt/splunk`.

Tarball install

To install Splunk on a Linux system, expand the tarball into an appropriate directory. The default install directory is `/opt/splunk`.

When installing with the tarball:

- Splunk does not create the `splunk` user automatically. If you want Splunk to run as a specific user, you must create the user manually.
- Be sure the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.

What gets installed

Splunk package status:

```
dpkg --status splunk
```

List all packages:

```
dpkg --list
```

Start Splunk

Splunk can run as any user on the local system. If you run Splunk as a non-root user, make sure that Splunk has the appropriate permissions to read the inputs that you specify. Refer to the instructions for running Splunk as a non-root user for more information.

To start Splunk from the command line interface, run the following command from `$SPLUNK_HOME/bin` directory (where $SPLUNK_HOME is the directory into which you installed Splunk):

```
./splunk start
```

By convention, this document uses:

- `$SPLUNK_HOME` to identify the path to your Splunk installation.
- `$SPLUNK_HOME/bin/` to indicate the location of the command line interface.

Startup options

The first time you start Splunk after a new installation, you must accept the license agreement. To start Splunk and accept the license in one step:

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

Note: There are two dashes before the `accept-license` option.

Launch Splunk Web and log in

After you start Splunk and accept the license agreement,

1. In a browser window, access Splunk Web at `http://<hostname>:port`.

- `hostname` is the host machine.
- `port` is the port you specified during the installation (the default port is 8000).

2. Splunk Web prompts you for login information (default, username `admin` and password `changeme`) before it launches. If you switch to Splunk Free, you will bypass this logon page in future sessions.

What's next?

Now that you've installed Splunk, **what comes next?**

Uninstall Splunk

Use your local package management commands to uninstall Splunk. In most cases, files that were not originally installed by the package will be retained. These files include your configuration and index files which are under your installation directory.

If you can't use package management commands, follow the instructions for manually uninstalling Splunk components.

RedHat Linux

To uninstall from RedHat Linux

```
rpm -e splunk_product_name
```

Debian Linux

To uninstall from Debian Linux:

```
dpkg -r splunk
```

To purge (delete everything, including configuration files):

```
dpkg -P splunk
```

Install on Solaris

This topic provides the procedures for installing Splunk on Solaris.

Upgrading?

If you are upgrading, first review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding to the instructions in this topic.

Install Splunk

Splunk for Solaris is available as a PKG file or a tarball.

PKG file install

The PKG installation package includes a request file that prompts you to answer a few questions before Splunk installs.

```
pkgadd -d ./splunk_product_name.pkg
```

A list of the available packages is displayed.

- Select the packages you wish to process (the default is "all").

The installer then prompts you to specify a base installation directory.

- To install into the default directory, `/opt/splunk`, leave this blank.

PKG file upgrade

To upgrade an existing Splunk installation using a PKG file, you should use the `instance` parameter, either in the system's default package installation configuration file (`/var/sadm/install/admin/default`) or in a custom configuration file that you define and call.

In the default or custom configuration file, set `instance=overwrite`. This will prevent the upgrade from creating a second `splunk` package (with `instance=unique`), or failing (with `instance=quit`). For information about the `instance` parameter, see the Solaris man page (`man -s4 admin`).

To upgrade Splunk using the system's default package installation file, use the same command line as you would for a fresh install.

```
pkgadd -d  ./splunk_product_name.pkg
```

You will be prompted to overwrite any changed files, answer yes to every one.

To upgrade using a custom configuration file, type:

```
pkgadd -a conf_file -d ./splunk_product_name.pkg
```

To run the upgrade silently (and not have to answer yes for every file overwrite), type:

```
pkgadd -n -d  ./splunk_product_name.pkg
```

Tarball install

To install Splunk on a Solaris system, expand the tarball into an appropriate directory. By default, Splunk installs into `/opt/splunk/`.

When installing with the tarball:

- Splunk does not create the `splunk` user automatically. If you want Splunk to run as a specific user, you must create the user manually.
- Be sure the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.

What gets installed

Splunk package info:

```
pkginfo -l splunk
```

List all packages:

```
pkginfo
```

Start Splunk

Splunk can run as any user on the local system. If you run Splunk as a non-root user, make sure that Splunk has the appropriate permissions to read the inputs that you specify. For more information, refer to the instructions on running Splunk as a non-root user.

To start Splunk from the command line interface, run the following command from `$SPLUNK_HOME/bin` directory (where $SPLUNK_HOME is the directory into which you installed Splunk):

```
./splunk start
```

By convention, Splunk's documentation uses:

- `$SPLUNK_HOME` to identify the path to your Splunk installation.
- `$SPLUNK_HOME/bin/` to indicate the location of the command line interface.

Startup options

The first time you start Splunk after a new installation, you must accept the license agreement. To start Splunk and accept the license in one step:

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

Note: There are two dashes before the `accept-license` option.

Launch Splunk Web and log in

After you start Splunk and accept the license agreement,

1. In a browser window, access Splunk Web at http://mysplunkhost:port, where:

- `mysplunkhost` is the host machine.
- `port` is the port you specified during the installation (8000).

2. Splunk Web prompts you for login information (default, username `admin` and password `changeme`) before it launches. If you switch to Splunk Free, you will bypass this logon page in future sessions.

What's next?

Now that you've installed Splunk, **what comes next?**

Uninstall Splunk

Use your local package management commands to uninstall Splunk. In most cases, files that were not originally installed by the package are retained. These files include your configuration and index files which are under your installation directory.

```
pkgrm splunk
```

Install on Mac OS

This topic describes how to install Splunk on MacOS.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding.

Installation options

The Mac OS build comes in two forms: a DMG package and a tarball. Below are instructions for the:

- Graphical (basic) and command line installs using the DMG file.
- Tarball install.

Note: if you require two installations in different locations on the same host, use the tarball. The pkg installer cannot install a second instance. If one exists, it will remove it upon successful install of the second.

Graphical install

1. Double-click on the DMG file.

A **Finder** window containing splunk.pkg opens.

2. In the **Finder** window, double-click on splunk.pkg.

The Splunk installer opens and displays the **Introduction**, which lists version and copyright information.

3. Click **Continue**.

The **Select a Destination** window opens.

4. Choose a location to install Splunk.

- To install in the default directory, `/Applications/splunk`, click on the harddrive icon.
- To select a different location, click **Choose Folder...**

5. Click **Continue**.

The pre-installation summary displays. If you need to make changes,

- Click **Change Install Location** to choose a new folder, or
- Click **Back** to go back a step.

6. Click **Install**.

Your installation will begin. It may take a few minutes.

7. When your install completes, click **Finish**.

Command line install

1. To mount the dmg:

```
hdid splunk_package_name.dmg
```

2. To Install

- To the root volume:

```
installer -pkg splunk.pkg -target /
```

- To a different disk of partition:

```
installer -pkg splunk.pkg -target /Volumes\ Disk
```

`-target` specifies a target volume, such as another disk, where Splunk will be installed in `/Applications/splunk`.

To install into a directory other than `/Applications/splunk` on any volume, use the graphical installer as described above.

Tarball install

To install Splunk on a Mac OS, expand the tarball into an appropriate directory. The default install directory is `/Applications/splunk`.

When installing with the tarball:

- Splunk does not create the `splunk` user automatically. If you want Splunk to run as a specific user, you must create the user manually.
- Be sure the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.

Start Splunk

Splunk can run as any user on the local system. If you run Splunk as a non-root user, make sure that Splunk has the appropriate permissions to read the inputs that you specify.

To start Splunk from the command line interface, run the following command from `$SPLUNK_HOME/bin` directory (where $SPLUNK_HOME is the directory into which you installed Splunk):

```
./splunk start
```

By convention, this document uses:

- `$SPLUNK_HOME` to identify the path to your Splunk installation.
- `$SPLUNK_HOME/bin/` to indicate the location of the command line interface.

Startup options

The first time you start Splunk after a new installation, you must accept the license agreement. To start Splunk and accept the license in one step:

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

Note: There are two dashes before the `accept-license` option.

Launch Splunk Web and log in

After you start Splunk and accept the license agreement,

1. In a browser window, access Splunk Web at `http://<hostname>:port`

- `hostname` is the host machine.
- `port` is the port you specified during the installation (the default port is 8000).

2. Splunk Web prompts you for login information (default, username `admin` and password `changeme`) before it launches. If you switch to Splunk Free, you will bypass this logon page in future sessions.

What's next?

Now that you've installed Splunk, **what comes next?**

Manage your license

If you are performing a new installation of Splunk or switching from one license type to another, you must install or update your license.

Uninstall Splunk

Use your local package management commands to uninstall Splunk. In most cases, files that were not originally installed by the package will be retained. These files include your configuration and index files which are under your installation directory.

You can also simply go to `$SPLUNK_HOME/bin`, type `./splunk stop` on the command line and then delete the `$SPLUNK_HOME` directory and everything under it.

Install on FreeBSD

The FreeBSD builds comes in two forms: an installer (5.4-intel) and a tarball (i386). Both are gzipped tarball (.tgz) files.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding.

Prerequisites

For FreeBSD 8 , Splunk requires compatibility packages. To install the compatibility package:

1. Install the port:

```
portsnap fetch update

cd /usr/ports/misc/compat7x/ && make install clean
```

2. Add the package:

```
pkg_add -r compat7x-amd64
```

Basic install

To install FreeBSD using the intel installer:

```
pkg_add splunk_package_name-6.1-intel.tgz
```

Important: This installs Splunk in the default directory, `/opt/splunk`. If `/opt` does not exist, you will need to create it prior to running the install command. If you don't, you might receive an error message. Splunk recommends that you create a symbolic link to another filesystem and install Splunk there, since best practices for FreeBSD maintain a small root ("`/`") filesystem.

To install Splunk in a different directory:

```
pkg_add -v -p /usr/splunk splunk_package_name-6.1-intel.tgz
```

The FreeBSD package system does not have native upgrade support. There are some add-on utilities which try to manage it, but this is not explicitly tested. To upgrade a package on FreeBSD you can either uninstall the prior package, and install the new package, or you can upgrade the existing installation using a tarball install as below.

Tarball install

To install Splunk on a FreeBSD system, expand the tarball into an appropriate directory. The default install directory is `/opt/splunk`.

When installing with the tarball:

- Splunk does not create the `splunk` user automatically. If you want Splunk to run as a specific user, you must create the user manually.
- Be sure the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.

After you install

To ensure that Splunk functions properly on FreeBSD, you must:

1. Add the following to `/boot/loader.conf`

```
kern.maxdsiz="2147483648" # 2GB
kern.dfldsiz="2147483648" # 2GB
machdep.hlt_cpus=0
```

2. Add the following to `/etc/sysctl.conf`:

```
vm.max_proc_mmap=2147483647
```

A restart of the OS is required for the changes to effect.

If your server has less than 2 GB of memory, reduce the values accordingly.

What gets installed

To see the list of Splunk packages:

```
pkg_info -L splunk
```

To list all packages:

```
pkg_info
```

Start Splunk

Splunk can run as any user on the local system. If you run Splunk as a non-root user, make sure that Splunk has the appropriate permissions to read the inputs that you specify.

To start Splunk from the command line interface, run the following command from `$SPLUNK_HOME/bin` directory (where $SPLUNK_HOME is the directory into which you installed Splunk):

```
./splunk start
```

By convention, this document uses:

- `$SPLUNK_HOME` to identify the path to your Splunk installation.
- `$SPLUNK_HOME/bin/` to indicate the location of the command line interface.

Startup options

The first time you start Splunk after a new installation, you must accept the license agreement. To start Splunk and accept the license in one step:

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

Note: There are two dashes before the `accept-license` option.

Launch Splunk Web and log in

After you start Splunk and accept the license agreement,

1. In a browser window, access Splunk Web at `http://<hostname>:port`

- `hostname` is the host machine.
- `port` is the port you specified during the installation (the default port is 8000).

2. Splunk Web prompts you for login information (default, username `admin` and password `changeme`) before it launches. If you switch to Splunk Free, you will bypass this logon page in future sessions.

What's next?

Now that you've installed Splunk, **what comes next?**

Manage your license

If you are performing a new installation of Splunk or switching from one license type to another, you must install or update your license.

Uninstall Splunk

Use your local package management commands to uninstall Splunk. In most cases, files that were not originally installed by the package will be retained. These files include your configuration and index files which are under your

installation directory.

To uninstall Splunk from the default location:

```
pkg_delete splunk
```

To uninstall Splunk from a different location:

```
pkg_delete -p /usr/splunk splunk
```

Install on AIX

This topic will guide you through installing Splunk on the AIX platform.

Important: The user Splunk is installed as must have permission to read `/dev/urando` and `/dev/random` or the installation will fail.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding.

Install Splunk

The AIX install comes in tarball form.

When installing with the tarball:

- Splunk does not create the `splunk` user automatically. If you want Splunk to run as a specific user, you must create the user manually.
- Be sure the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.
- We recommend you use GNU tar to unpack the tarballs, as AIX tar can fail to unpack long file names, fail to overwrite files, and other problems. If you must use the system tar, be sure to check the output for error messages.

To install Splunk on an AIX system, expand the tarball into an appropriate directory. The default install directory is `/opt/splunk`.

For **AIX 5.3**, check to make sure your service packs are up to date. Splunk requires the following service level:

```
$ oslevel -r
5300-005
```

Start Splunk

Splunk can run as any user on the local system. If you run Splunk as a non-root user, make sure that Splunk has the appropriate permissions to read the inputs that you specify. Refer to the instructions for running Splunk as a non-root user for more information.

To start Splunk from the command line interface, run the following command from `$SPLUNK_HOME/bin` directory (where $SPLUNK_HOME is the directory into which you installed Splunk):

```
./splunk start
```

By convention, this document uses:

- `$SPLUNK_HOME` to identify the path to your Splunk installation.
- `$SPLUNK_HOME/bin/` to indicate the location of the command line interface.

Note: The AIX version of Splunk does not register itself to auto-start on reboot.

Startup options

The first time you start Splunk after a new installation, you must accept the license agreement. To start Splunk and accept the license in one step:

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

Note: There are two dashes before the `accept-license` option.

For more information, refer to "Splunk startup options" in this manual.

Launch Splunk Web and log in

After you start Splunk and accept the license agreement,

1. In a browser window, access Splunk Web at `http://<hostname>:port`

- `hostname` is the host machine.
- `port` is the port you specified during the installation (the default port is 8000).

2. Splunk Web prompts you for login information (default, username `admin` and password `changeme`) before it launches. If you switch to Splunk Free, you will bypass this logon page in future sessions.

What's next?

Now that you've installed Splunk, **what comes next?**

Manage your license

If you are performing a new installation of Splunk or switching from one license type to another, you must update your license.

Uninstall Splunk

Use your local package management commands to uninstall Splunk. In most cases, files that were not originally installed by the package will be retained. These files include your configuration and index files which are under your installation directory.

Install on HP-UX

To install Splunk on an HP-UX system, expand the tarball, using GNU tar, into an appropriate directory. The default install directory is `/opt/splunk`.

NOTE: The system default tar on HP-UX will not successfully extract the splunk tar. GNU tar is a pre-requisite, or you can unpack the tar on another platform.

When installing with the tarball:

- Splunk does not create the `splunk` user automatically. If you want Splunk to run as a specific user, you must create the user manually.
- Be sure the disk partition has enough space to hold the uncompressed volume of the data you plan to keep indexed.

Upgrading?

If you are upgrading, review the upgrade documentation later in this manual and check READ THIS FIRST for any migration considerations before proceeding.

Start Splunk

Splunk can run as any user on the local system. If you run Splunk as a non-root user, make sure that Splunk has the appropriate permissions to read the inputs that you specify.

To start Splunk from the command line interface, run the following command from `$SPLUNK_HOME/bin` directory (where $SPLUNK_HOME is the directory into which you installed Splunk):

```
./splunk start
```

By convention, this document uses:

- `$SPLUNK_HOME` to identify the path to your Splunk installation.
- `$SPLUNK_HOME/bin/` to indicate the location of the command line interface.

Note: The HP-UX version of Splunk does not register itself to auto-start on reboot.

Startup options

The first time you start Splunk after a new installation, you must accept the license agreement. To start Splunk and accept the license in one step:

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

Note: There are two dashes before the `accept-license` option.

Launch Splunk Web and log in

After you start Splunk and accept the license agreement,

1. In a browser window, access Splunk Web at `http://<hostname>:port`

- `hostname` is the host machine.

- `port` is the port you specified during the installation (the default port is 8000).

2. Splunk Web prompts you for login information (default, username `admin` and password `changeme`) before it launches. If you switch to Splunk Free, you will bypass this logon page in future sessions.

What's next?

Now that you've installed Splunk, **what comes next?**

Manage your license

If you are performing a new installation of Splunk or switching from one license type to another, you must install or update your license.

Uninstall Splunk

To uninstall Splunk on HPUX, you must stop splunk, disable boot-start (if you configured it), and then delete the Splunk installation.

Note: The $SPLUNK_HOME variable refers to the directory where you installed Splunk.

1. Stop Splunk:

```
$SPLUNK_HOME/bin/splunk stop
```

2. If you enabled boot-start, run the following command as root:

```
$SPLUNK_HOME/bin/splunk disable boot-start
```

3. Delete the Splunk installation directories:

```
rm -rf $SPLUNK_HOME
```

Other things you may want to delete:

- If you created any indexes and did not use the Splunk default path, you must delete those directories as well.

- If you created a user or group for running Splunk, you should also delete them.

Install a license

This topic discusses installing new licenses. Before you proceed, you may want to review these topics:

- Read "How Splunk licensing works" in the Admin Manual for an introduction to Splunk licensing.
- Read "Groups, stacks, pools, and other terminology" in the Admin Manual for more information about Splunk license terms.

Add a new license

To add a new license:

1. Navigate to **Manager > Licensing**.

2. Click **Add license**.

3. 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.

4. Click **Install**. If this is the first Enterprise license that you are installing, you must restart Splunk. Your license is installed.

License violations

Violations occur when you exceed the maximum indexing volume allowed for your license. If you exceed your licensed daily volume on any one calendar day, you will get a violation *warning*. The message persists for 14 days. **If you have 5 or more warnings on an Enterprise license or 3 warnings on a Free license**

in a rolling 30-day period, you are in *violation* of your license and search will be disabled. Search capabilities return when you have fewer than 5 (Enterprise) or 3 (Free) warnings in the previous 30 days, or when you apply a temporary reset license (available for Enterprise only). To obtain a reset license, contact your sales rep. The license comes with instructions on how to apply it.

Note: Summary indexing volume is not counted against your license.

If you get a violation warning, you have until midnight (going by the time on the license master) to resolve it before it counts against the total number of warnings within the rolling 30-day period.

During a license violation period:

- Splunk does not stop indexing your data. Splunk only blocks search while you exceed your license.
- Searches to the `_internal` index are not disabled. This means that you can still access the Indexing Status dashboard or run searches against `_internal` to diagnose the licensing problem.

Got license violations? Read "About license violations" in the Admin Manual or "Troubleshooting indexed data volume" from the Splunk Community Wiki.

More licensing information is available in the "Manage Splunk licenses" chapter in the Admin Manual.

Start Splunk for the first time

Start Splunk for the first time

To start Splunk:

On Windows

You can start Splunk on Windows using either the command line, or the Windows Services Manager. Using the command line offers more options, described later in this section. In a `cmd` window, go to `C:\Program Files\Splunk\bin` and type:

```
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 System Properties dialog's Advanced tab.)

On UNIX

Use the Splunk command-line interface (CLI):

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

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

Other start options

To accept the license automatically when you start Splunk for the first time, add the `accept-license` option to the `start` command:

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

The startup sequence displays:

```
Checking prerequisites...
Checking http port [8000]: open
Checking mgmt port [8089]: open
Verifying configuration.  This may take a while...
Finished verifying configuration.
Checking index directory...
Verifying databases...
Verified databases: _audit, _blocksignature, _internal, _thefishbucket,
history, main, sampledata, splunklogger, summary
Checking index files
All index checks passed.
All preliminary checks passed.
Starting splunkd...
Starting splunkweb...
Splunk Server started.
The Splunk web interface is at http://<hostname>:8000
```

Note: If the default ports are already in use (or are otherwise not available), Splunk will offer to use the next available port. You can either accept this option or specify a port for Splunk to use.

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

- If you run `$SPLUNK_HOME/bin/splunk start --no-prompt`, Splunk proceeds with startup until it requires you to answer a question. Then, it displays the question, why it is quitting, and quits.
- If you run `SPLUNK_HOME/bin/splunk start --answer-yes`, Splunk proceeds with startup and automatically answers "yes" to all yes/no questions. Splunk displays the question and answer as it continues.

If you run start with all three options in one line, for example:

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

- Splunk does not ask you to accept the license.
- Splunk answers yes to any yes/no question.
- Splunk quits when it encounters a non-yes/no question.

Start and disable individual processes

You can start and stop individual Splunk processes by adding the process as an object to the `start` command. The objects include:

- `splunkd`, the Splunk server daemon.

- `splunkweb`, Splunk's Web interface process.

For example, to start only `splunkd`:

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

To disable `splunkweb`:

```
$SPLUNK_HOME/bin/splunk disable webserver
```

For more information about `start`, refer to the CLI help page:

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

Launch Splunk Web

Navigate to:

```
http://mysplunkhost:8000
```

Use whatever host and port you chose during installation.

The first time you log in to Splunk Enterprise, the default login details are:
Username - *admin*
Password - *changeme*

Splunk Free does not have access controls.

Upgrade from an earlier version

About upgrading to 4.3 READ THIS FIRST

This topic is about what you need to know **before you upgrade to Splunk 4.3.**

What's new and awesome in 4.3?

Check out Meet Splunk 4.3 in the Release Notes for a full list of the new features we've delivered in 4.3.

Review the Known issues in the Release Notes for a list of issues and workarounds in this release.

Always back up your existing deployment first

Get into the habit of backing up your existing deployment before any upgrade or migration.

You can manage your risk by using technology that allows you to restore your Splunk install and data to a state prior to the upgrade, whether you use external backups, disk or file system snapshots, or other means. When backing up your Splunk data, consider the $SPLUNK_HOME directory, as well as any indexes outside of it.

For more information about backing up your Splunk deployment, read the topics beginning with "What you can back up" in the Admin Manual.

Upgrading from 4.0 and later

Splunk supports a direct upgrade from versions 4.0 and later to version 4.3.

If you're upgrading from 4.2.x or later, read the rest of this topic first before proceeding with the installation instructions linked below.

If you're upgrading from 4.1.x to 4.3, Splunk recommends that you also review the 4.2.x version of the topic you're reading now as well as this version before proceeding with the installation instructions:

- Upgrade to 4.3 on Linux, Solaris, FreeBSD, HP-UX, AIX, and MacOS

- Upgrade to 4.3 on Windows

Migrating from 3.4.x and earlier

Upgrading from version 3 of Splunk to version 4 is a significant, complex undertaking. This is because Splunk versions 4.0 and later have a completely different architecture than Splunk versions 3.x and earlier. For this reason, there is no direct upgrade path from version 3.x to version 4.3. Attempting to upgrade directly from 3.x to 4.3 is unsupported and is strongly discouraged. You should only upgrade to version 4.3 after you have determined that your migration to version 4.0 is stable and complete. Read "Upgrade from 3.x to 4.3" for additional information and instructions.

After you've migrated from 3.4.x to 4.0, you can then upgrade from 4.0 to 4.3 using the UNIX and Windows instructions in this manual.

Upgrade distributed deployments

If you're planning to upgrade your distributed Splunk environment, be sure to read "Upgrade your distributed environment" in the Distributed Deployment Manual for guidance on how to do so with minimal impact.

Upgrade universal forwarders

Upgrading universal forwarders is a different process than upgrading full Splunk. Before upgrading your universal forwarders, be sure to read the appropriate upgrade topic for your operating system:

- Upgrade the Windows universal forwarder
- Upgrade the Unix universal forwarder

To learn about interoperability and compatibility between indexers and universal forwarders, read "Indexer and universal forwarder compatibility" in the "Deployment Overview" topic of the Distributed Deployment Manual.

You want to know this stuff

Upgrading to 4.3 from 4.0 and later is pretty simple, but here are a few tips and gotchas:

Changes to how Splunk handles host names retrieved from scripted inputs

In version 4.2 of Splunk, various scripted inputs (in particular, the Windows-based scripted inputs defined in `inputs.conf`) would generate multiple hostnames that identified a single host. In version 4.3, we've changed that behavior to produce a single name for the computer running these inputs, by identifying various equivalent names for the local system. The inputs recognize when a system name matches one of the equivalent names such as the socket hostname, the NetBIOS-style `COMPUTERNAME`, or the name defined in `server.conf` when Splunk is first run on a machine. For all of these cases, the inputs allow the value to be controlled by your choice in `inputs.conf`.

We've also fixed a bug that sometimes causes invalid hostname generation.

This might impact you if you have searches defined that expect fully-qualified domain names. After you upgrade, if you need to retain the previous functionality, you can specify the desired hostname by using the 'host' attribute on the input in question.

No change to license management

Licenses issued for 4.0-4.1 as well as newer 4.2 XML licenses will both work in 4.3 without change. Regardless of whether you've migrated Splunk or reinstalled it, your existing 4.x licenses will work with 4.3.

If your license is not accepted by Splunk 4.3.x, it is possible that your upgrade rights have expired due to a lapsed maintenance agreement. Contact Splunk Support to reinstate your agreement and upgrade entitlement.

Splunk uses more Unix file descriptors

Splunk 4.3 uses more file descriptors on Unix filesystems than version 4.2 did when monitoring files.

Before you upgrade, consider increasing the number of open file descriptors your system can use with the `ulimit` command.

Changes to Windows Registry Monitor configuration files

We've changed how Splunk handles Windows Registry monitoring configurations.

On versions of Splunk prior to 4.3, Registry monitoring was configured with two configuration files - `sysmon.conf` and `regmon-filters.conf`. The `sysmon.conf` file contained global settings for which event types to monitor, and the `regmon-filters.conf` file contained the regular expressions used to filter the Registry keys that Splunk should monitor.

In 4.3, `sysmon.conf` is deprecated and is no longer shipped. All Registry monitoring configuration now resides in `regmon-filters.conf`. After you upgrade your Splunk instance and restart Splunk, the Splunk Registry monitoring utility (`splunk-regmon.exe`) will read your existing `sysmon.conf` and transfer any relevant configuration values to `regmon-filters.conf`. It will then mark `sysmon.conf` as "migrated" and ignore any further updates to this file.

Changes to Windows Registry monitor default behavior

In 4.3, we migrated the defaults for Windows Registry monitoring from core Splunk into the Windows app. If you currently use the Registry monitoring tools in Splunk 4.2, you might be impacted by an issue that arises when you upgrade. Be sure to read "Workaround for Registry monitoring configuration issue" to learn how to address this specific issue.

Note: The previous URL is hard-coded to take you to version 4.3 of the documentation because this issue is resolved in 4.3.1 and later; upgrading from 4.2.x to 4.3.1 and later will not experience this issue.

Changes to Active Directory monitoring configuration file defaults

We made a minor change to the case of the `targetDc` attribute in `admon.conf`. If you are using Splunk's Active Directory monitoring utility, you might see a warning that Splunk found a typo in this configuration file after you upgrade.

To prevent this notice from occurring, edit `admon.conf` and change any defined `targetDC` attributes to `targetDc` before upgrading.

Note that making this change does not affect how Active Directory monitoring works.

Changes to Windows event log monitoring configuration values

We corrected a minor issue with regard to the acceptable values for the `start_from` attribute for the `[WinEventLog]` stanza of `inputs.conf`. If you use Splunk's Windows event log monitoring inputs, you might be impacted by an issue related to case sensitivity of these values after you upgrade.

The two acceptable values for `start_from` are `oldest` and `newest`. These values must be specified in lower case. Before upgrading, review your `inputs.conf` to make sure that the values have the proper case.

The Windows app doesn't use the performance monitor collection features

The Windows app still does not use the Windows performance monitor collection features available in Splunk 4.3. While the app does work, and is supported, by default it will continue to gather local performance metrics using WMI-based scripted inputs.

If you want to use the new features, or you're using a universal forwarder to send data with the default performance monitoring data collections to an instance that's running the app, then you'll need to update the searches within the app, based on your defined performance monitoring collections.

You can follow the Windows app on Splunkbase for future updates.

Splunk expects canonical IP addresses in its configuration file stanzas

When editing configuration files such as `inputs.conf` to define stanzas that contain IP addresses, make sure to use canonical IP address formatting (for example, `[tcp://10.0.0.10:9995]`).

Do not put leading zeros in your IP addresses (`[tcp://010.000.000.010:9995]`, as doing so can cause problems with reading and correctly parsing the configuration files.

The accessibility defaults for Splunk Web's single sign-on (SSO) implementation have changed

We've made a change to how Splunk's SSO implementation works. Beginning with Version 4.3, the default mode of Splunk Web's SSO is 'strict' instead of 'permissive'.

The `SSOMode` attribute in `$SPLUNK_HOME/etc/system/local/web.conf` defines the mode of accessibility. Strict mode blocks all SSO requests except those that originate from a trusted IP address (See "Use single sign-on with Splunk" for more information.)

If you use Splunk's SSO implementation, and did not explicitly define `SSOMode` in `web.conf`, you might experience problems with SSO after you upgrade. To fix the problem, add the `SSOMode` attribute under the `[settings]` stanza as follows:

```
[settings]
SSOMode = permissive
```

Splunk timeline now displayed in HTML5 Canvas instead of Flash

We've included a new timeline histogram that is based on JavaScript and HTML5 Canvas by default. This replaces the Flash-based timeline that is present in version 4.2.

After you upgrade, the timeline is by default displayed using Canvas instead of Flash. You might see slight differences in the names of certain elements of the timeline, as well as how the timeline zooms in and out.

If you do not access Splunk Web with an HTML5-capable browser, timelines will continue to display in Flash.

Existing saved searches must be reconfigured to use the new report builder view

If you have saved searches that reference the `report_builder` view, Splunk might display an error "`Cannot find the 'report_builder_display' view`" when accessing this search after you upgrade.

If you experience this issue, you can edit `savedsearches.conf` and change the value of the `displayview` attribute from `report_builder` to `report_builder_display`.

You must restart Splunk in order for this change to take effect.

For more information about defining reports with Report Builder, read "Define reports" in the User Manual.

Dashboards with certain charting properties might not render or display as desired

We've added the ability to display charts based on JavaScript (JSChart) instead of Adobe Flash. This introduces some important considerations when creating or editing charts after an upgrade:

- Some charting properties are not yet supported by JSChart. As a result, after you upgrade, some charts might display differently than they did before the upgrade.

- Backward compatibility with Flash is maintained through the upgrade. This means that dashboards with charting properties that are not supported by JSChart will be kept as Flash-based during the upgrade.
- If Splunk encounters a charting property in a dashboard that is not supported by JSChart, it falls back to rendering the dashboard in Flash.
- If you view these Flash-based dashboards with a browser that is not capable of displaying Flash, the browser will not display the charts.
- If you make a change to an existing dashboard outside of Splunk Web, and use a charting property that is not supported by JSChart, you might experience undesirable results in how the dashboard is rendered.
- Even when the dashboard is properly rendered using JSChart, you might see some charts display incorrectly - particularly those charts that feature row-grouping or have large numbers of categories.

For additional information about how customizing charts in dashboards affects how dashboards are rendered, review "Chart customization and non-Flash chart displays" in the "Advanced charting options" topic of the Developer Manual.

The Panel Editor introduces a new editing workflow for some dashboards

The new Panel Editor changes how you edit some dashboards - in particular, those which use Simple XML. The new editor only edits dashboard properties that are supported by JSChart. There might also be instances where dashboards edited with the new Panel Editor display inconsistently or incorrectly.

For information on the new Panel Editor and how to use it, read "Edit dashboard panel visualizations" in the User Manual.

Simple XML now always has precedence when editing form or dashboard settings

We've fixed a bug where some Simple XML form and dashboard settings were not correctly prioritized when you edited a form or dashboard.

Simple XML settings now always have precedence over settings defined in viewstates.conf that have the same name. This is true whether you edit the form or dashboard using the new Panel Editor, or edit the form or dashboard XML directly.

If you want to use the settings defined in viewstates.conf, you can change the value of the desired setting within the form or dashboard XML, or remove it from the XML entirely.

Changes to LDAP authentication configuration for multi-domain support

We've added functionality for multi-domain LDAP configurations in 4.3.

- We've made changes to settings in `authentication.conf`:
  - The authSettings attribute is now a comma-separated list of LDAP strategies.
  - The roleMap attribute is no longer global, but is rather scoped to a specific LDAP strategy. When you upgrade, all 'roleMap' attributes will be renamed according to the strategy that they are defined in.

- We've made changes to the nomenclature of existing LDAP strategies. When you upgrade to 4.3, any existing LDAP strategies that contain commas (',') will be renamed.

Existing LDAP configurations will be automatically updated to work with these changes as part of migration.

For additional information about using Splunk with multiple LDAP strategies, review "Use multiple LDAP strategies" in the "Set up user authentication with LDAP" topic in the Admin Manual.

Users of LDAP version 2 might experience user configuration issues

We've fixed a bug that caused Splunk's LDAP configuration screen to display LDAP users that could not actually log into Splunk when Splunk was configured to map users to groups by distinguished name (DN). This problem occurred because those users fell outside the scope of the list of user base DNs, as defined by the `UserBaseDN` attribute in `authentication.conf`.

If you are using LDAP version 2, Splunk will populate the LDAP user list with only those users who *have already successfully logged in to Splunk,* as opposed to the full list of users who have login access.

If you are concerned about this, Splunk recommends that you utilize LDAP version 3 in your environment, as it does not cause Splunk to exhibit this anomaly.

Bloom filters get created after upgrading

When you start Splunk after upgrading to 4.3, Bloom filters are created for specific buckets in your Splunk indexes for the most recent 30 days of indexed data. Typically, these filters take up around 10% of the size of an index bucket on

disk. You might want to ensure that you have adequate disk space prior to upgrading, to address the increased overhead.

You can adjust the backfill time window (the amount of days for which bloom filters are generated for an index bucket) by editing the `maxBloomBackfillBucketAge` attribute in `indexes.conf`.

Splunk's database-checking utility might use more resources

After you upgrade to 4.3, Splunk's database consistency checking utility (fsck) might use more system resources (in particular, disk I/O) when they run, particularly if bloom filters are being created at the same time.

The deployment server's serverclass.conf has a new attribute "machineTypesFilter"

We've added a new attribute for `serverclass.conf` for distributed deployments. This new attribute allows you to specify which servers in a given server class are filtered from a whitelist or blacklist that is defined within that class.

The new `machineTypesFilter` attribute tells Splunk that, for a given whitelist or blacklist of servers in a server class definition, the machine type(s) defined by the attribute must be applied only to the list of servers defined within that whitelist or blacklist. Review the following `serverclass.conf` example:

```
[serverClass:Austin_Linux]
whiltelist.0=*.austin.ourcompany.com
machineTypesFilter=linux-i686, linux-x86_64
```

In this example, the "Austin_Linux" server class is defined as "all hosts in the *.austin.ourcompany.com domain whose machine type is either 'linux-i686' or 'linux-x86-64'."

This is different from the existing `machineTypes` attribute, which selects all hosts whose machine type matches what is defined by that attribute, regardless of what is defined in the whitelist or blacklist for a given server class.

Section 508 compliance on Mac OS X and Firefox

When you upgrade to 4.3, you might lose the ability to use the Tab key to toggle through various navigation menus and dashboards in Splunk when running Firefox on Mac OS X. To fix this problem, open System Preferences, select **Keyboard**, and under the **Full Keyboard Access** option at the bottom of the

page, select **All controls.** Close System Preferences and then restart Firefox. You should now be able to tab through the menus and dashboards.

Migrate your 4.1.x apps to 4.3

Check out this topic about how to migrate 4.1.x apps to 4.3.

Upgrade from 3.x to 4.3

This topic discusses the steps required to upgrade from version 3.x of Splunk to version 4.3. If you currently operate a Splunk deployment based on version 3.4 or earlier, read this topic to find out what to expect and how to do it properly.

Overview

Upgrading from version 3 of Splunk to version 4 is a significant, complex undertaking. This is because Splunk versions 4.0 and later have a completely different architecture than Splunk versions 3.x and earlier. The changes between versions 3.x and 4.0 are diverse and far-reaching, and present many challenges during the upgrade process.

For this reason, there is **no** direct upgrade path from version 3.x to version 4.3. Attempting to upgrade directly from 3.x to 4.3 is unsupported and is strongly discouraged. You should only upgrade to version 4.3 after you have determined that your migration to version 4.0 is stable and complete.

Depending on the complexity of your Splunk deployment, you might want to perform the upgrade from 3.x to 4.0 manually.

Before you begin the upgrade, be sure to back up your Splunk 3.x deployment completely first by using whatever backup tools are available to you.

Upgrade process

To upgrade from Splunk 3.x to Splunk 4.3, you must perform these steps in order:

1. Back up your Splunk 3.x deployment and have it available in case problems occur during the upgrade.

2. Carefully review "What to expect when upgrading to 4.0" in this manual for specifics on the differences between versions, what gets upgraded, what you must upgrade manually, licensing changes and other pertinent information.

3. Upgrade from version 3.x to version 4.0. Be sure to read the important migration notes in the following topics for additional information:

- Upgrade from version 3.x to 4.0 on UNIX
- Upgrade from version 3.x to 4.0 on Windows

Note: If you choose instead to upgrade Splunk manually from version 3 to version 4.0, read this topic:

- Steps for manual migration to Splunk 4.x

4. Confirm that your Splunk deployment works properly on version 4.0.

5. Once you have confirmed that your environment is working properly, upgrade from version 4.0 to 4.3.

- Upgrade to 4.3 on Linux, Solaris, FreeBSD, HP-UX, AIX, and MacOS
- Upgrade to 4.3 on Windows

Upgrade to 4.3 on UNIX

This topic describes the procedure for upgrading your Splunk instance from version 4.0.x or later to 4.3.

Before you upgrade

Make sure you've read this information before proceeding, as well as the following:

Back your files up

Before you perform the upgrade, we strongly recommend that you **back up all of your files**, including Splunk configurations, indexed data, and binaries. Splunk does not provide a means of downgrading to previous versions; if you need to revert to an older Splunk release, just reinstall it.

For information on backing up data, read "Back up indexed data".

For information on backing up configurations, read "Back up configuration information".

How upgrading works

After performing the installation of the new version, your configuration changes are not actually made until you start Splunk. You can run the migration preview utility at that time to see what will be changed before the files are updated. If you choose to view the changes before proceeding, a file containing the changes that the upgrade script proposes to make is written to `$SPLUNK_HOME/var/log/splunk/migration.log.<timestamp>`

Steps for upgrading

1. Execute the `$SPLUNK_HOME/bin/splunk stop` command.

Important: Make sure no other processes will start Splunk automatically (such as Solaris SMF).

2. To upgrade and migrate from version 4.0 and later, install the Splunk package over your existing Splunk deployment:

- If you are using a .tar file, expand it into the same directory with the same ownership as your existing Splunk instance. This overwrites and replaces matching files but does not remove unique files.
  Note: AIX tar will fail to correctly overwrite files when run as a user other than root. Use GNU tar (`gtar`) to avoid this problem.
- If you are using a package manager, such as RPM, type `rpm -U [--prefix <existing Splunk location>] splunk_package_name.rpm`
- If you are using a .dmg file (on Mac OS X), double-click it and follow the instructions. Be sure specify the same installation directory as your existing installation.

3. Execute the `$SPLUNK_HOME/bin/splunk start` command.

The following output is displayed:

```
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]
```

4. Choose whether you want to run the migration preview script to see what changes will be made to your existing configuration files, or proceed with the migration and upgrade right away.

5. If you choose to view the expected changes, the script provides a list.

6. Once you've reviewed these changes and are ready to proceed with migration and upgrade, run `$SPLUNK_HOME/bin/splunk start` again.

Note: You can complete Steps 3 to 5 in one line:

To accept the license and view the expected changes (answer 'n') before continuing the upgrade:

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

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

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

Upgrade to 4.3 on Windows

This topic describes the procedure for upgrading your Windows Splunk instance from version 4.0.x or later to 4.3. You can upgrade using the GUI installer, or by running the `msiexec` utility on the command line as described in "Install on Windows via the command line".

Before you upgrade

Make sure you've read this information before proceeding, as well as the following:

Make sure you specify the same domain user

When upgrading, you must explicitly specify the same domain user that you specified during first time install. If you do not specify the same user, Splunk will default to using the Local System User. If you accidentally specify the wrong user during your installation, use these instructions to switch to the correct user **before starting Splunk.**

Don't change the ports

Changing the management port and/or the HTTP port when upgrading is not supported.

Back your files up

Before you perform the upgrade, we strongly recommend that you back up all of your files, including Splunk configurations, indexed data and binaries. Splunk does not provide a means of downgrading to previous versions; if you need to revert to an older Splunk release, just reinstall it.

For information on backing up data, read "Back up indexed data".

For information on backing up configurations, read "Back up configuration information".

Note: When you upgrade to Splunk 4.3 on Windows, the installer will overwrite any custom certificate authority (CA) certificates you have created in `%SPLUNK_HOME%\etc\auth`. If you have custom CA files, make sure to back them up before you upgrade. After the upgrade, you can copy them back into `%SPLUNK_HOME%\etc\auth` to restore them. After you have restored the certificates, restart Splunk.

Upgrade using the GUI installer

1. Stop Splunk by either using the Services control panel or executing the `%SPLUNK_HOME%\bin\splunk stop` command.

2. Download the new MSI file from the Splunk download page.

3. Double-click the MSI file. The Welcome panel is displayed. Follow the on-screen instructions to upgrade Splunk. For information about each panel, refer to the installation instructions.

4. Splunk will start up by default when you complete the installation.

A log of the changes made to your configuration files during the upgrade is placed in %TEMP%.

Upgrade using the command line

1. Stop Splunk either by using the Services control panel or executing the `%SPLUNK_HOME%\bin\splunk stop` command.

2. Download the new MSI file from the Splunk download page.

3. Use the instructions in "Install on Windows via the command line".

- If Splunk is running as a user other than the Local System user, you must explicitly specify this user in your command-line instruction.
- You can use the LAUNCHSPLUNK option to specify whether Splunk should start up automatically or not when you're finished, but you cannot change any other settings.
- DO NOT change the ports (SPLUNKD_PORT and WEB_PORT) at this time.

4. Depending on your specification, Splunk may start automatically when you complete the installation.

A log of the changes made to your configuration files during the upgrade is placed in `%TEMP%`.

Start Splunk

On Windows, Splunk is installed by default into `%SYSTEMDRIVE%\Program Files\Splunk` and is started by default.

You can start and stop the following Splunk processes via the Windows Services control panel:

- Server process: `splunkd`
- Web interface process: `splunkweb`

You can also start, stop, and restart both processes at once by going to `%SYSTEMDRIVE%\Program Files\Splunk\bin` and typing

```
#  splunk [start|stop|restart]
```

Migrate searches for local performance monitoring metrics in the Windows app

The Windows app currently does not make use of the Windows performance monitor collection features available in Splunk 4.3. While the app does work, and is supported, by default it will continue to gather local performance metrics using WMI-based inputs.

If you're using the Windows app, and want to use the new features, or you're using a universal forwarder to send data with the default performance monitoring data collections to an instance that's running the app, then you'll need to update the searches within the app, based on your defined performance monitoring collections.

You can follow the Windows app on Splunkbase for future updates.

Other setup tasks

Configure PDF printing for Splunk Web

Splunk Web users can generate PDF output from any dashboard, view, search or report. To enable this functionality, you must download the PDF Report Server App from Splunkbase. Next install this into a Splunk instance on a single Linux host. The PDF Report Server App for that Splunk instance will then accept requests from all other Splunk hosts on the network.

Note: Splunk PDF printing is not yet available on the Windows platform. Currently you must have a Linux instance of Splunk running on your network to support PDF printing.

PDF printing requires a Splunk Enterprise license. For more information about Splunk licenses, refer to Types of Splunk licenses in the Admin Manual.

Note: To configure a machine to act as a PDF server, you can use the Forwarder license as described in Types of Splunk licenses - Forwarder license in the Admin Manual as long as that Splunk instance does no indexing.

System requirements

Only the instance of Splunk hosting the PDF Server functionality has to be running Linux; it will then serve PDFs to any Splunk instance on the network, regardless of the platform.

The PDF Server is supported on any Linux server that:

- is running Splunk 4.1 or later
- has the Xvfb X server, xauth, and base fonts installed
- has glibc 2.3 or later installed (glibc 2.4 required for 64 bit systems such as Centos 5.0 or later, needed by Flash)

How to configure PDF printing

Download and install the PDF Report Server app

1. Install Splunk 4.1 or later on a Linux host. For instructions, refer to "Install on Linux" in this manual.

2. Download and install the PDF Report Server app. You can click **Browse more apps** in Launcher, or you can download it separately from Splunkbase here. (Manual app installation instructions can be found here.)

3. Ensure that the Xvfb X server, xauth and fonts for your Linux distribution are installed. These are included with most Linux distributions, but not installed by default.

- On Redhat/CentOS/Fedora, type: `yum install Xvfb xauth bitstream-vera-fonts` (note the capital X for Xvfb).
- On Debian/Ubuntu, type: `apt-get install xvfb xauth fontconfig libxrender1 libxinerama1`
- Note that the names of font packages tend to change more than other package names.

4. Launch Splunk Web on the Linux host and navigate to **Manager**.

5. Navigate to **System Settings > Email Alert Settings**.

6. Check the **Use PDF Report Server** box.

7. Click **Save**.

This host is now configured to serve PDFs to all Splunk hosts with PDF printing enabled as described in the next section.

Note: If the hostname of the Splunk Web instance that this PDF Report Server will talk to is not resolvable in DNS, enter its IP address or a hostname that resolves to that IP in the **Link hostname** field. This will ensure that Splunk Web can contact the PDF Report Server, and that links sent in emailed PDF reports work correctly. If the field is left empty, Splunk will try to autodetect the hostname.

Enable PDF printing on a search head

On each search head you want to have access to the PDF printing feature:

1. Launch Splunk Web and navigate to **Manager > System settings > Email alert settings**.

2. Ensure that a valid SMTP server hostname is entered in the **Mail host** field.

3. Check the **Use PDF Report Server** box. Some more options are displayed.

4. Enter the **Report Server URL** if the PDF Server is not installed on your search head.

This host can now use the PDF server you configured in the previous section to generate PDF reports.

You can also set the same option by enabling the `reportServerEnabled` option in `$SPLUNK_HOME/etc/system/local/alert_actions.conf` for each search head.

Use PDF printing

For information on scheduling PDF reports to be sent by email, refer to "Schedule delivery of dashboard printouts via email" in the User Manual.

Status page

Check that you can display a test PDF with the status page.

In **Manager > System settings > Email alert settings**, click on the small blue link for "status page".

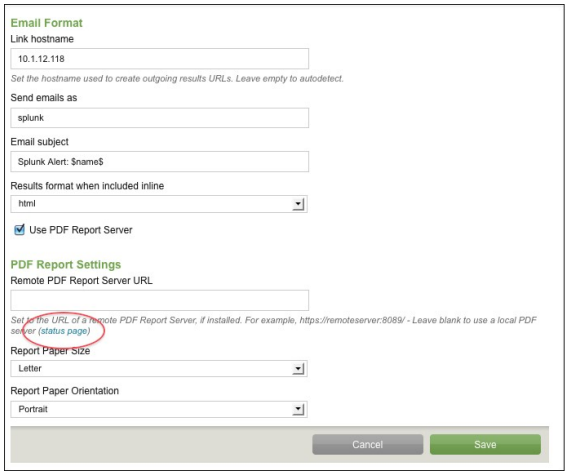

Clicking on "status page" takes you to a new screen:

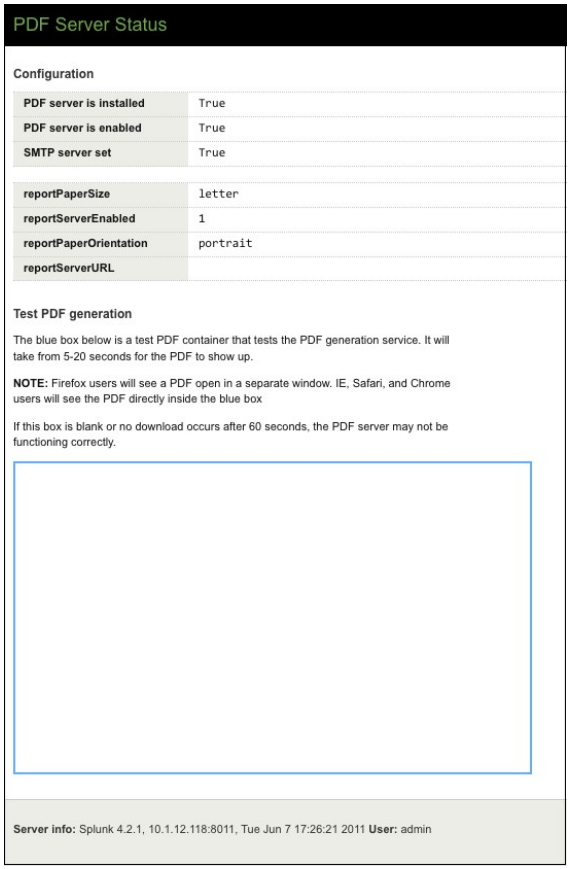

Check that the top three settings read "True." Within about a minute, a PDF should be sent to your browser. Be sure to read the message written on the PDF.

In case of trouble

If you run into problems while configuring the PDF server, such as errors in the UI, you may want to refer to `$SPLUNK_HOME/var/log/splunk/python.log` which may help clarify the problem.

Note: The stanza created in savedsearches.conf for the scheduled PDF will show "search = | noop". This is normal.

For more debugging, edit `$SPLUNK_HOME/etc/log.cfg` by setting DEBUG on

[python]

splunk = DEBUG

If the host running Splunk Web does not have a valid DNS entry for its hostname,

make sure you have set the **Link Hostname** field in **Manager > System settings > Email alert settings** to a hostname that resolves to the IP address of the host running Splunk Web, or that you have specified its IP address directly.

For more specifics, see "I'm having problems with the Splunk PDF Server app" in the Troubleshooting Manual.

Run Splunk as a different or non-root user

You can run Splunk as any user on the local system. If you run Splunk as a non-root user, make sure Splunk has the appropriate permissions to:

- Read the files and directories it is configured to watch. Some log files and directories may require root or superuser access to be indexed.
- Write to Splunk's directory and execute any scripts configured to work with your alerts or scripted input.
- Bind to the network ports it is listening on (ports below 1024 are reserved ports that only root can bind to).

Note: Because ports below 1024 are reserved for root access only, Splunk will only be able to listen on port 514 (the default listening port for syslog) if it is running 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.

Instructions

To run Splunk as a non-root user, you need to first install Splunk as `root`. Then, **before you start Splunk for the first time**, change the ownership of the `splunk` directory to the desired user. The following are instructions to install Splunk and run it as a non-root user, `splunk`.

Note: In the following examples, `$SPLUNK_HOME` represents the path to the Splunk installation directory.

1. Create the user and group, `splunk`.

For Linux, Solaris, and FreeBSD:

```
useradd splunk
groupadd splunk
```

For Mac OS:

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

2. As `root` and using one of the packages (not a tarball), run the installation.

Important: Do not start Splunk yet.

3. Use the `chown` command to change the ownership of the `splunk` directory and everything under it to the desired user.

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

4. Start Splunk.

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

Also, if you want to start Splunk 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
```

This example command assumes:

- If Splunk is installed in an alternate location, update the path in the command accordingly.
- Your system may not have `sudo` installed. If this is the case, you can use `su`.
- If you are installing using a tarball and want Splunk to run as a particular user (such as `splunk`), you must create that user manually.
- The `splunk` user will need access to `/dev/urandom` to generate the certs for the product.

Solaris 10 privileges

When installing on Solaris 10 as the `splunk` user, you must set additional privileges to start `splunkd` and bind to reserved ports.

To start `splunkd` as the `splunk` user on Solaris 10, run:

```
# usermod -K defaultpriv=basic,net_privaddr,proc_exec,proc_fork splunk
```

To allow the `splunk` user to bind to reserved ports on Solaris 10, run (as root):

```
# usermod -K defaultpriv=basic,net_privaddr splunk
```

Correct the user selected during Windows installation

If you have selected "other user" during the Windows GUI installation, and that user does not exist or perhaps you mistyped the information, you can go into the Windows Service Control Manager and specify the correct information, **as long as you have not started Splunk yet**.

If you specified an invalid user during the Windows GUI installation process, you will see two popup error windows.

To change the user:

1. In Control Panel > Administrative Tools > Services, find the `Splunkd` and `SplunkWeb` services. You'll notice that they are not started and are currently owned by the Local System User.

2. Right click on each service and choose **Properties**. The properties dialog for that service is displayed.

3. Select the **Log On** tab.

4. Select the **This account** radio button and fill in the correct domain\username and password.

5. Click **Apply**.

6. Click **OK**.

6. Repeat for the second service (you must do this for both Splunkd and Splunk Web).

7. You can now either start both services from the Service Manager or from the Splunk command line interface.

Uninstall Splunk

Before you uninstall, stop Splunk. Navigate to `$SPLUNK_HOME/bin` and type `./splunk stop` (or just `splunk stop` on Windows).

Use your local package management commands to uninstall Splunk. In most cases, files that were not originally installed by the package will be retained. These files include your configuration and index files which are under your installation directory.

Note: `$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`; for Mac OS, it is `/Applications/splunk`.

RedHat Linux

To uninstall Splunk on RedHat:

```
rpm -e splunk_product_name
```

Debian Linux

To uninstall Splunk on Debian:

```
dpkg -r splunk
```

To purge (delete everything, including configuration files) on Debian:

```
dpkg -P splunk
```

FreeBSD

To uninstall Splunk from the default location on FreeBSD:

```
pkg_delete splunk
```

To uninstall Splunk from a different location on FreeBSD:

```
pkg_delete -p /usr/splunk splunk
```

Solaris

To uninstall Splunk on Solaris:

```
pkgrm splunk
```

Windows

To uninstall Splunk on Windows:

Use the **Add or Remove Programs** option in the Control Panel. In Windows 7 and Windows Server 2008, that option is available under **Programs and Features.**

Note: 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 manually

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

Note: These instructions will not remove any `init` scripts that have been created.

1. Stop Splunk.

```
$SPLUNK_HOME/bin/splunk stop
```

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

For Linux and Solaris:

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

For FreeBSD and Mac OS

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

3. Remove the Splunk installation directory, `$SPLUNK_HOME`. For example:

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

Note: For Mac OS, you can also remove the installation directory by dragging the folder into the trash.

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

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

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

For Linux, Solaris, and FreeBSD:

```
userdel splunk
groupdel splunk
```

For Mac OS: You can 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 installed.

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

Configure a standalone 3.4.x deployment server

If you are planning to migrate to Splunk 4.x, but do not want to migrate your deployment clients until a later time, you can set up a stripped-down, standalone 3.4.x deployment server to serve your deployment clients until you're ready to migrate them (Splunk 4.x deployment server is incompatible with clients older than 4.x).

This procedure assumes the following:

- You have an existing deployment server at fflanda.splunk.com listening on port 8089 for deployment clients.

- Your deployment clients are all 3.x and are all polling this deployment server.
- The deployment classes are in $SPLUNK_HOME/etc/modules/distributedDeployment/classes.
- This Splunk instance is also an index server that must be upgraded.

Given the above, the procedure is as following:

1. Download the latest Splunk 3.4.x build for your architecture.

2. Back up the existing $SPLUNK_HOME/etc using `tar -zxvf $SPLUNK_HOME/etc > /tmp/splunk_old_etc.tgz`

3. Stop Splunk, remove deployment.conf and deployment classes.

4. If $SPLUNK_HOME = /opt/splunk, mv to /opt/splunk_old. Otherwise, install the 3.4.x tarball or rpm in the default location

5. Extract the splunk_old_etc.tgz over top of the fresh installation.

6. Remove/rename any inputs.conf/outputs.conf files in /opt/splunk_depserver/etc/system/local or /opt/splunk_depserver/etc/apps/. You will probably want to keep authentication.conf, server.conf, /opt/splunk/etc/passwd, /opt/splunk/etc/auth/* - pretty much anything but inputs.conf, outputs.conf and the unchanged splunk-launch.conf.

7. Disable Splunk Web on the newly installed instance using the CLI or web.conf.

8. Execute `mv /opt/splunk /opt/splunk_depserver`

9. Edit /opt/splunk_depserver/etc/splunk-launch.conf to change $SPLUNK_HOME to /opt/splunk_depserver. If $SPLUNK_DB is also set, comment out this variable so that the new instance does not try to write to the old data store.

10. To ensure that this deployment server remains functional, switch its license out for a 3.x forwarder license. Copy $SPLUNK_HOME/etc/splunk-forwarder.license to $SPLUNK_HOME/etc/splunk.license .

11. Execute `/opt/splunk_depserver/bin/splunk start`

12. Execute `mv /opt/splunk_old /opt/splunk`, then perform migration.

13. During post-migration start up, Splunk will notice that the old management port is bound, and will prompt the admin to change the management port. **Keep track of this new port as you must update it in any distributed search or REST configurations**.

14. Execute `/opt/splunk_depserver/bin/splunk list deploy-clients -auth admin:changeme` and verify that the deployment clients have been in touch with the deployment server.

14. Review /opt/splunk_depserver/var/log/splunk/splunkd.log and /opt/splunk/var/log/splunk/splunkd.log for errors.

What comes next?

Ready to start using Splunk?

Now that you've got Splunk installed on one server, here are some links to get you started:

- Learn what Splunk is, what it does, and how it's different.
- Learn how to add your data to Splunk.
- Add and manage users.
- Learn how to search, monitor, report, and more
- One of Splunk's biggest differences from traditional technologies is the fact that it **classifies and interprets data at search-time**. We call this Splunk knowledge: learn how it works and how to use it.
- Plan your Splunk deployment, from gigabytes to terabytes per day.
- If you downloaded Splunk packaged with an **app** (for example, Splunk + WebSphere), go to Splunk Web and select the app in Launcher to go directly to the app?s setup page. To see more information about the setup and deployment for a packaged app, search for the app name on Splunkbase.

Reference

PGP Public Key

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.4.1 (GNU/Linux)

mQGiBEbE21QRBADEMonUxCV2kQ2oxsJTjYXrYCWCtH5/OnmhK5lT2TQaE9QUTs+w
nM3sVInQqwRwBDH2qsHgqjJS0PIE867n+lVuk0gSVzS5SOlYzQjnSrisvyN452MF
2PgetHq8Lb884cPJnxR6xoFTHqOQueKEOXCovz1eVrjrjfpnmWKa/+5X8wCg/CJ7
pT7OXHFN4XOseVQabetEbWcEAIUaazF2i2x9QDJ+6twTAlX2oqAquqtBzJX5qaHn
OyRdBEU2g4ndiE3QAKybuq5f0UM7GXqdllihVUBatqafySfjlTBaMVzd4ttrDRpq
Wya4ppPMIWcnFG2CXf4+HuyTPgj2cry2oMBm2LMfGhxcqM5mpoyHqUiCn7591Ra/
J2/FA/0c2UAUh/eSiOn89I6FhFOicT5RPtRpxMoEM1Di15zJ7EXY+xBVF9rutqhR
5OI9kdHibYTwf4qjOOPOA7237N1by9GiXY/8s+rDWmSNKZB+xAaLyl7cDhYMv7CP
qFTutvE8BxTsF0MgRuzIHfJQE2quuxKJFs9lkSFGuZhvRuwRcrQgS2ltIFdhbGxh
Y2UgPHJlbGVhc2VAc3BsdW5rLmNvbT6IXgQTEQIAHgUCRsTbVAIbAwYLCQgHAwID
FQIDAxYCAQIeAQIXgAAKCRApYLH9ZT+xEhsPAKDimP8sdCr2ecPm8mre/8TK3Bha
pQCg3/xEickiRKKlpKnySUNLR/ZBh3m5Ag0ERsTbbRAIAIdfWiOBeCj8BqrcTXxm
6MMvdEkjdJCr4xmwaQpYmS4JKK/hJFfpyS8XUgHjBz/7zfR8Ipr2CU59Fy4vb5oU
HeOecK9ag5JFdG2i/VWH/vEJAMCkbN/6aWwhHt992PUZC7EHQ5ufRdxGGap8SPZT
iIKY0OrX6Km6usoVWMTYKNm/v7my8dJ2F46YJ7wIBF7arG/voMOg1Cbn7pCwCAtg
jOhgjdPXRJUEzZP3AfLIc3t5iq5n5FYLGAOpT7OIroM5AkgbVLfj+cjKaGD5UZW7
SO0akWhTbVHSCDJoZAGJrvJs5DHcEnCjVy9AJxTNMs9GOwWaixfyQ7jgMNWKHJp+
EyMAAwYH/RLNK0HHVSByPWnS2t5sXedIGAgm0fTHhVUCWQxN3knDIRMdkqDTnDKd
qcqYFsEljazI2kx1ZlWdUGmvU+Zb8FCH90ej8O6jdFLKJaq50/I/oY0+/+DRBZJG
3oKu/CK2NH2VnK1KLzAYnd2wZQAEja4O1CBV0hgutVf/ZxzDUAr/XqPHy5+EYg96
4Xz0PdZiZKOhJ5g4QjhhOL3jQwcBuyFbJADw8+Tsk8RJqZvHfuwPouVU+8F2vLJK
iF2HbKOUJvdH5GfFuk6o5V8nnir7xSrVj4abfP4xA6RVum3HtWoD7t//75gLcW77
kXDR8pmmnddm5VXnAuk+GTPGACj98+eISQQYEQIACQUCRsTbbQIbDAAKCRApYLH9
ZT+xEiVuAJ9INUCilkgXSNu9p27zxTZh1kL04QCg6YfWldq/MWPCwa1PgiHrVJng
p4s=
=Mz6T
-----END PGP PUBLIC KEY BLOCK-----
```

Installing the key

Copy and paste the key into a file. Install the key using:

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

File manifest

A complete inventory of the files and permissions that ship with your Splunk installation can be found in the root directory of your Splunk installation. The file will end with `-manifest`.