


# Bare Bones Splunk



Feeling lost? A bit Splunk-n00bish? Come along with me as we go through some of the main concepts of Splunk. We'll use only a basic, local installation of Splunk to do this. You won't need to do anything other than [install the Splunk Enterprise package](#) onto your local machine (e.g. Windows, Mac OS, or Linux computer).

Also, I link to a lot of things in this article. Except for downloading & installing Splunk, you really don't have to follow the links unless you want to dive deeper and read more context about Splunk terms and concepts I touch on.

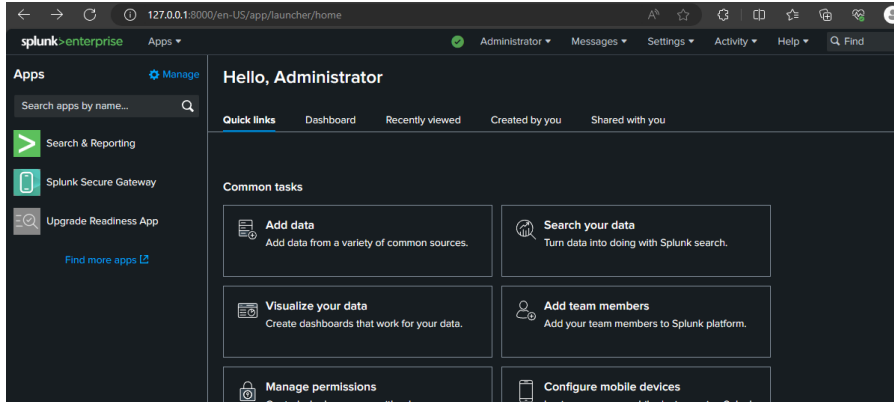
## Install Splunk!

Do this on your local machine - just a single Splunk instance, because guess what? That install on your local machine is already going to be ingesting data and allowing you to analyze it - we won't even need to get your data in, but I'll walk you through some [SPL](#) queries and we can pretend together.

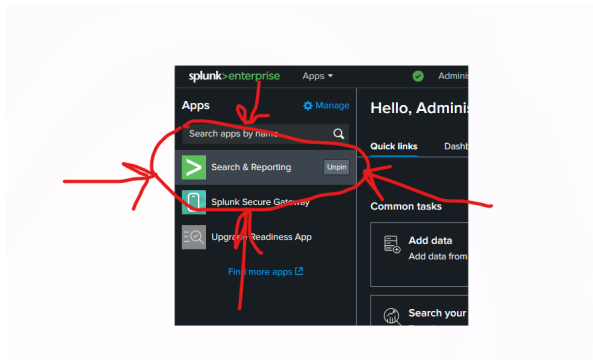
Splunk already has [documentation on how to install Splunk Enterprise](#) from the [download](#), but if you follow that documentation you end up being left at the [What Happens Next?](#) That provides links to use case articles, how to configure your environment for additional users, but you're here just to try it out. Let's get going faster - you can complete the [Search Tutorial](#) and install even more pieces for that later...

For now, let's just learn about how Splunk can ingest and analyze data using what's already happening under the hood with your freshy install. Do the install, and meet me back here.

I'll assume you've gotten to this screen, where you are logged into your newly installed Splunk instance via your browser (<https://127.0.0.1:8000>):




We will be making a lot of use of the **Search** app built into Splunk. You get there by clicking on Search & Reporting in the left menu after you log into Splunk. It's circled in red in this screenshot in case you don't see it:



This is important - you'll use this a lot in your Splunk career, even if that career just spans this document.

## Splunking the Splunk's Splunk

After installation Splunk is already writing to its own log files. Guess what? These are already being ingested by this locally installed Splunk instance. This really isn't much different than how Splunk Enterprise works, nor is the data being ingested probably too much different than some of the data you want to get into Splunk. Sure, it might be a totally different format, but the meta-concepts are the same. A web server is typically a web server. An application server is typically an application server. All of these have security implications.



- Is some executive breathing down your neck to give them visibility to all their crap? Then you better go through each section to see what Splunk can do for the whole enterprise.
- Are you wanting to use Splunk to monitor a web server? Jump to the **Monitor Your Web Server** section.
- Do you want to ingest your application logs to observe what your application and its components are doing? Jump to the **Observe Your App** section.
- Is security your business? Jump to the **Security for Security's Sake** section.

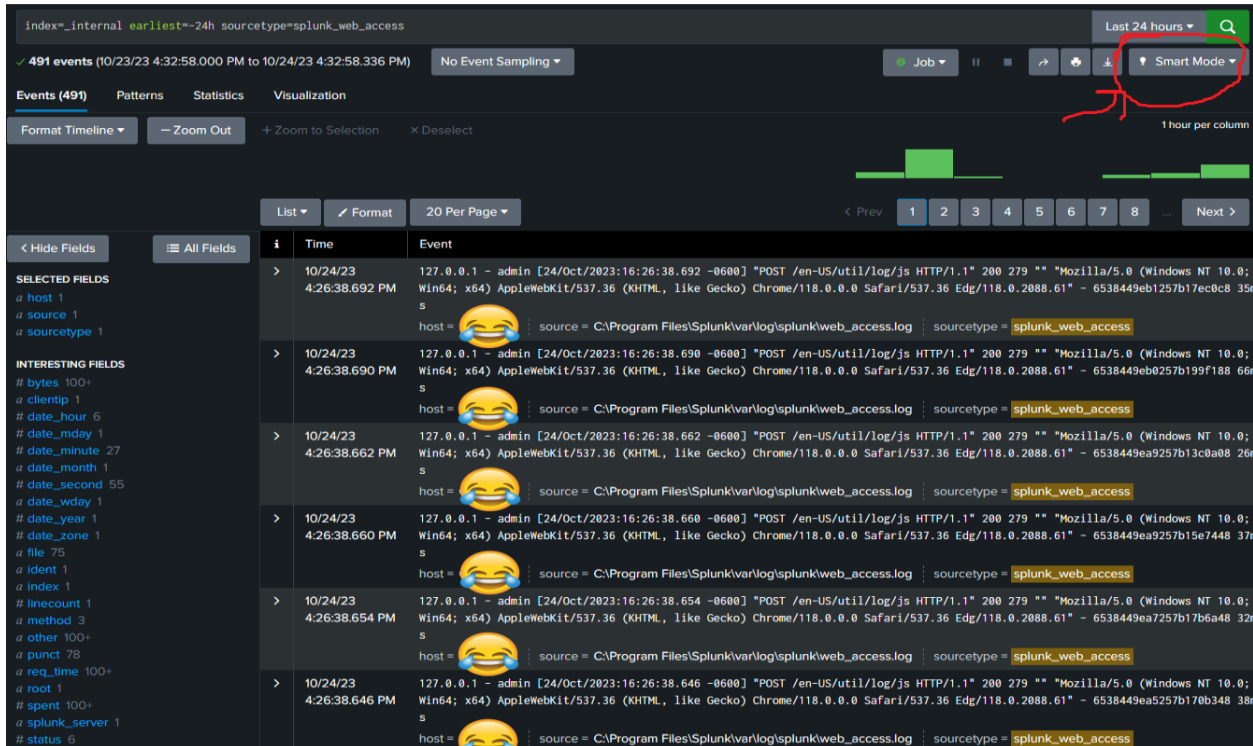
## Monitor Your Web Server

So you have a web server. Splunk has a web server - it's called splunkweb, and it's already running with your local installation. To see what it is doing let's search some of the data Splunk is already logging about its own web server.

1. Go to the Search app
2. Paste this SPL into the search box

```
index=_internal earliest=-24h sourcetype=splunk_web_access
```
3. Click the magnifying glass button way on the right

Look at all the pretty results! All this SPL query did was tell Splunk to return the last 24 hours of data in the `_internal` [index](#) for the [sourcetype](#) `splunk_web_access`. You should see something like this screenshot. If you have the same timestamps then let's chat time travel last Wednesday:

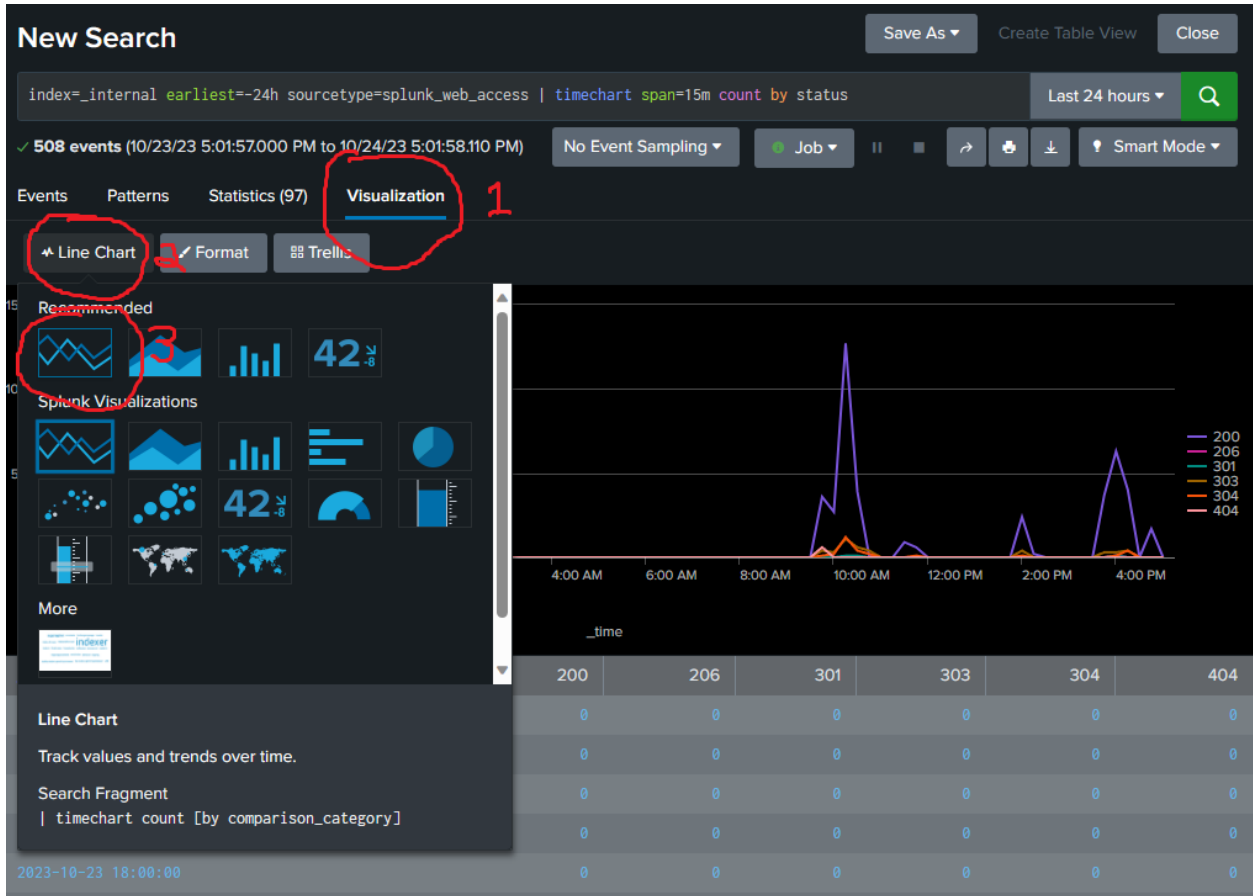


Gosh - that sure looks like web server logs for http access data. They are probably really similar to what your web server barfs out, but maybe in a slightly different format. There's an HTTP method, path, status. And if you have Smart Mode turned on for the search (see the red circle again?), then on the left side of your screen you'll see all of these [fields](#) that have names like what you might expect to find in the http access data.

Let's run a new query. Paste the following in your search box and run it again with that magnifying glass icon:


```
index=_internal earliest=-24h sourcetype=splunk_web_access |
timechart span=15m count by status
```

This search is summarizing the http status codes over the past 24 hours for splunkweb. It's giving us a table of data in 15 minute increments with a simple count of how many times that status occurred in that time window. I know, I know, a table of numbers is super cool. We all love Excel. To make this more interesting, click the Visualization tab (1), and if it isn't already a Line Chart then click the chart type (2) and then select the one that has multiple lines (3):



Do you think we're done? We are going to build a [dashboard](#):

- Click on Save As, then Dashboard Panel



- On the new screen fill out a couple of fields to define the dashboard:
  - Dashboard should be **New** (if you already did another section, select **Existing**)
  - Dashboard Title should be *My Enterprise* (if doing **Existing**, select the *My Enterprise* dashboard from the dropdown)
  - Dashboard ID will auto-populate, and that value is fine
  - Panel Title should be *Web Server Statuses*
- It should look like this if **New**:

**Save As Dashboard Panel** [X]

Dashboard:  New  Existing

Dashboard Title: My Enterprise

Dashboard ID <sup>?</sup>: my\_enterprise  
The dashboard ID can only contain letters, numbers, dashes, and underscores. Do not start the dashboard ID with a period.

Dashboard Description: optional

Dashboard Permissions:  Private  Shared in App

---

Panel Title: Web Server Statuses

Panel Powered By <sup>?</sup>: Q Inline Search

Drilldown <sup>?</sup>: No action

Panel Content:  Statistics  Line Chart

[Cancel] [Save]

7. Or like this if **Existing**:

**Save As Dashboard Panel** [X]

Dashboard:  New  Existing

My Enterprise ▾

---

Panel Title: Web Server Statuses

Panel Powered By <sup>?</sup>: Q Inline Search

Drilldown <sup>?</sup>: No action

Panel Content:  Statistics  Line Chart

[Cancel] [Save]

8. Click Save

You can View Dashboard to see what it looks like. Tada! You're a splunker now: SPL, dashboards, you're doing it!

## Observe Your App

The Splunk Enterprise you just installed has its own application server. It's called [splunkd](#). Hooray! This little thing is so similar to your enterprise. Let's now pretend splunkd is your app server, and we'll run some queries as if you're generating this data:

9. Go to the Search app
10. Paste this SPL into the search box

```
index=_internal earliest=-24h sourcetype=splunkd
```
11. Click the magnifying glass button way on the right

Nifty. There's a bunch of application data of a whole bunch of components doing things. Splunk is logging all this just like your server probably does. Also notice all of those [fields](#) on the left side of the screen. If you don't see a bunch of them, make sure Smart Mode is turned on for your search (see red circle):

The screenshot shows the Splunk Search interface. At the top, the search bar contains the query `index=_internal earliest=-24h sourcetype=splunkd`. Below the search bar, there are buttons for 'Save As', 'Create Table View', and 'Close'. A 'Smart Mode' button is circled in red. The search results show 206,932 events. The interface includes a timeline visualization and a table of results. The table has columns for 'Time' and 'Event'. The 'Event' column contains log entries from the Splunk application server.

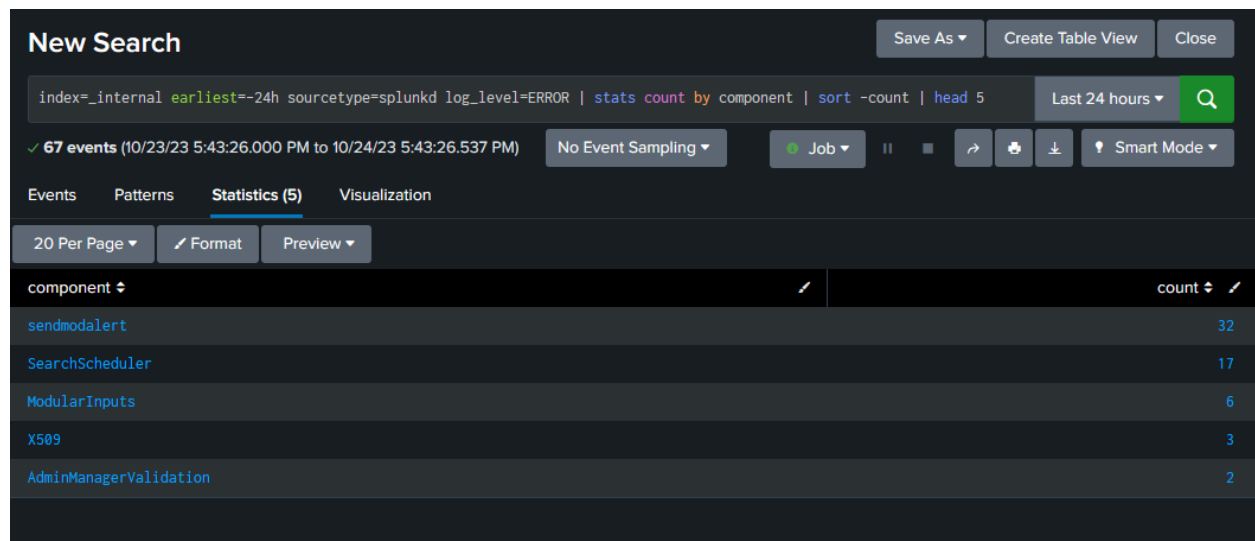
i	Time	Event
>	10/24/23 5:26:22.497 PM	10-24-2023 17:26:22.497 -0600 INFO NoahSearchPeerFetcher [26652 SearchPipelineExecutor] - Fetch requested. sid=ta_1698189982.554 use_cache=1 host = 🤪 source = C:\Program Files\Splunk\var\log\splunk\splunkd.log sourcetype = 🤪 splunkd
>	10/24/23 5:26:21.085 PM	10-24-2023 17:26:21.085 -0600 INFO NoahSearchPeerFetcher [17700 SearchPipelineExecutor] - Fetch requested. sid=ta_1698189981.553 use_cache=1 host = 🤪 source = C:\Program Files\Splunk\var\log\splunk\splunkd.log sourcetype = 🤪 splunkd
>	10/24/23 5:26:20.705 PM	10-24-2023 17:26:20.705 -0600 INFO Metrics - group=thruput, name=thruput, instantaneous_kbps=4.816, instantaneous_eps=18.038, average_kbps=6.048, total_k_processed=143924.000, kb=149.252, ev=559 host = 🤪 source = C:\Program Files\Splunk\var\log\splunk\metrics.log sourcetype = 🤪 splunkd
>	10/24/23 5:26:20.705 PM	10-24-2023 17:26:20.705 -0600 INFO Metrics - group=thruput, name=syslog_output, instantaneous_kbps=0.000, instantaneous_eps=0.000, average_kbps=0.000, total_k_processed=0.000, kb=0.000, ev=0 host = 🤪 source = C:\Program Files\Splunk\var\log\splunk\metrics.log sourcetype = 🤪 splunkd

Let's make some sense of all of this. Within the [index](#) \_internal, we are searching over the [sourcetype](#) splunkd, which is our application server logs. If you look at those fields on the left you will see names like component, log\_level, and thread\_id (again, probably similar to data your app server is logging).

Paste this in as a new search and run it with the magnifying glass icon:

```
index=_internal earliest=-24h sourcetype=splunkd log_level=ERROR | stats count by component | sort -count | head 5
```

This search is looking for all ERROR logs for the application server, counting the number of errors by component, and keeping the top 5 results with the most errors:




The screenshot shows the Splunk Search interface. At the top, the search query is: `index=_internal earliest=-24h sourcetype=splunkd log_level=ERROR | stats count by component | sort -count | head 5`. The results show 67 events from 10/23/23 5:43:26.000 PM to 10/24/23 5:43:26.537 PM. The search is in Smart Mode. Below the search bar, there are tabs for Events, Patterns, Statistics (5), and Visualization. The Statistics tab is active, showing a table with 20 items per page. The table has two columns: component and count. The results are as follows:

component	count
sendmodalert	32
SearchScheduler	17
ModularInputs	6
X509	3
AdminManagerValidation	2

Let's save this as a [dashboard](#).

1. Click on Save As, then Dashboard Panel



2. On the new screen fill out a couple of fields to define the dashboard:
  - a. Dashboard should be **New** (if you already did another section, select **Existing**)
  - b. Dashboard Title should be *My Enterprise* (if doing **Existing**, select the *My Enterprise* dashboard from the dropdown)
  - c. Dashboard ID will auto-populate, and that value is fine
  - d. Panel Title should be *App Server Errors by Component*
3. It should look like this if **New**:



### Save As Dashboard Panel

Dashboard  New  Existing

Dashboard Title

Dashboard ID <sup>?</sup>   
The dashboard ID can only contain letters, numbers, dashes, and underscores. Do not start the dashboard ID with a period.

Dashboard Description

Dashboard Permissions  Private  Shared in App

---

Panel Title

Panel Powered By <sup>?</sup>

Drilldown <sup>?</sup>

Panel Content

4. Or like this if **Existing**:

### Save As Dashboard Panel

Dashboard  New  Existing

---

Panel Title

Panel Powered By <sup>?</sup>

Drilldown <sup>?</sup>

Panel Content

5. Click Save

You can View Dashboard to see what it looks like. Tada! You're a splunker now: SPL, dashboards, you're doing it!

## Security for Security's Sake

There's security data that gets logged by Splunk, but since I'm assuming you have a shiny-new install and everything was done right, then you don't necessarily have any security incident data. Let's start by creating some:

1. Logout of your Splunk instance (i.e. <https://127.0.0.1:8000>)
2. Try some bogus usernames and fake passwords. Here's a few usernames for inspiration:
  - a. admin
  - b. administrator
  - c. root
  - d. sam
3. Log back into your instance for real now

Let's now find these security events and do something with them.

1. Go to the Search app
2. Paste this SPL into the search box

```
index=_internal earliest=-24h component=UiAuth
```
3. Click the magnifying glass button way on the right

Uh oh. There's the hackers doing their hacking. Your results should look similar to these:

**New Search** [Save As] [Create Table View] [Close]

index=\_internal earliest=-24h component=UiAuth [Last 24 hours] [Search]

✓ 21 events (10/23/23 6:06:46.000 PM to 10/24/23 6:06:47.175 PM) [No Event Sampling] [Job] [Pause] [Refresh] [Download] [Smart Mode]

Events (21) | Patterns | Statistics | Visualization

[Format Timeline] [Zoom Out] [+ Zoom to Selection] [Deselect] 1 hour per column

[List] [Format] [20 Per Page] [Prev] [1] [2] [Next]

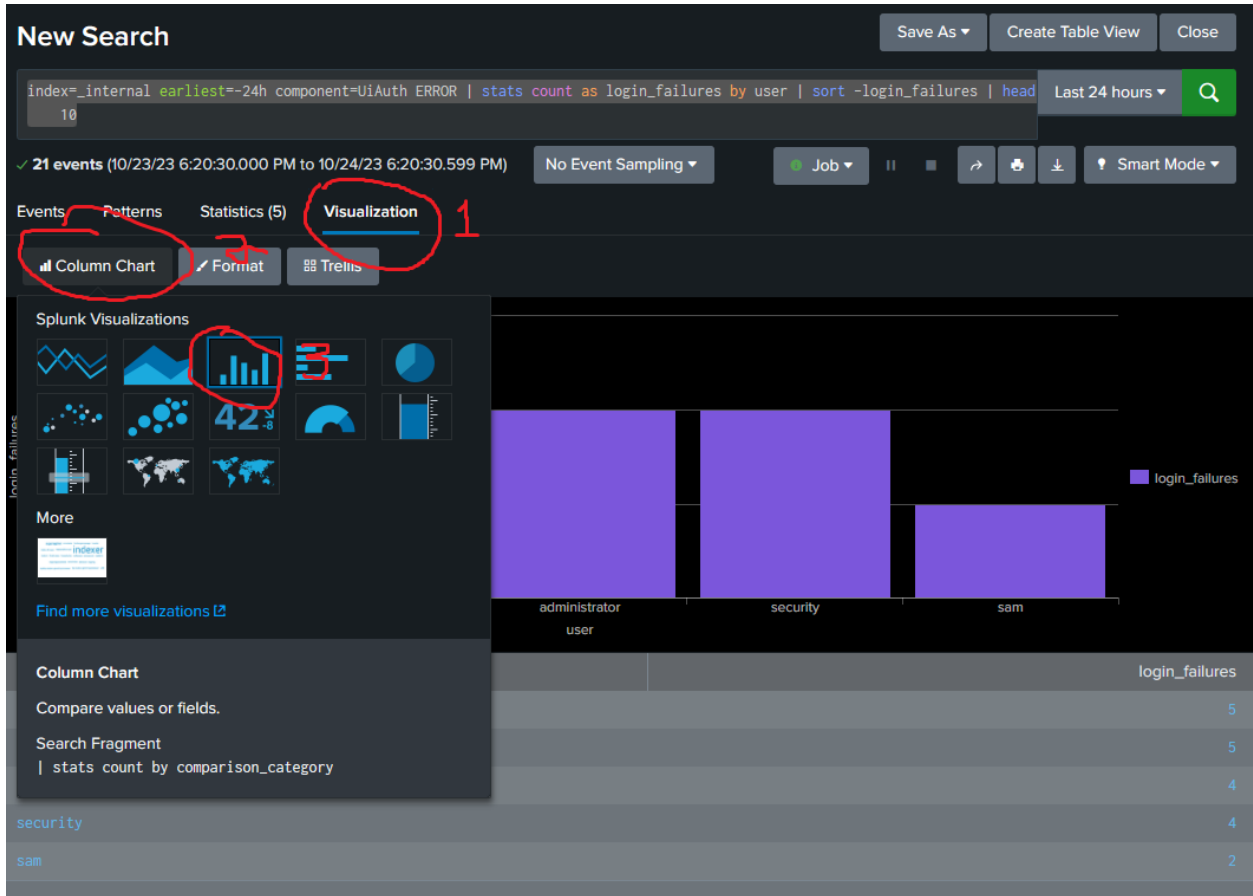
< Hide Fields	All Fields	i	Time	Event
SELECTED FIELDS a host 1 a source 1 a sourcetype 1	INTERESTING FIELDS a action 1 a clientip 1 a component 1 # date_hour 2 # date_mday 1 # date_minute 4 a date_month 1 # date_second 14 a date_wday 1 # date_year 1 # date_zone 1 a event_message 6 a eventtype 1 a index 1	>	10/24/23 6:06:42.482 PM	10-24-2023 18:06:42.482 -0600 ERROR UiAuth [28076 TcpChannelThread] - user=sam action=login status=failure session= reason=user-initiated useragent="Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/118.0.0.0 Safari/537.36 Firefox/118.0.2088.61" clientip=127.0.0.1 host = [redacted] source = C:\Program Files\Splunk\var\log\splunk\splunkd.log sourcetype = splunkd
		>	10/24/23 6:06:40.663 PM	10-24-2023 18:06:40.663 -0600 ERROR UiAuth [28076 TcpChannelThread] - user=sam action=login status=failure session= reason=user-initiated useragent="Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/118.0.0.0 Safari/537.36 Firefox/118.0.2088.61" clientip=127.0.0.1 host = [redacted] source = C:\Program Files\Splunk\var\log\splunk\splunkd.log sourcetype = splunkd
		>	10/24/23 6:03:13.661 PM	10-24-2023 18:03:13.661 -0600 ERROR UiAuth [28076 TcpChannelThread] - user=root action=login status=failure session= reason=user-initiated useragent="Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/118.0.0.0 Safari/537.36 Firefox/118.0.2088.61" clientip=127.0.0.1 host = [redacted] source = C:\Program Files\Splunk\var\log\splunk\splunkd.log sourcetype = splunkd

You'll also see [fields](#) on the left extracted by Splunk. If you don't see a bunch of them, make sure you're searching in Smart Mode (see red circle). This SPL query is searching the [index](#) \_internal, and we are interested in component=UiAuth because that's where we will see our failed login attempts.

Let's do a fancier search to get a better view of the failed login attempts (remember, this was you pretending to be *Hax0r the Hacker*):


```
index=_internal earliest=-24h component=UiAuth ERROR | stats count as login_failures by user | sort -login_failures | head 10
```

Our boss said they like bar charts, so let's turn this one into a chart by clicking on Visualization (1), and ensure the chart type (2) is set to be a Column Chart (3) like this:



We're not done yet. The big boss is getting this on their own [dashboard](#):

1. Click on Save As, then Dashboard Panel



2. On the new screen fill out a couple of fields to define the dashboard:
  - a. Dashboard should be **New** (if you already did another section, select **Existing**)
  - b. Dashboard Title should be *My Enterprise* (if doing **Existing**, select the *My Enterprise* dashboard from the dropdown)
  - c. Dashboard ID will auto-populate, and that value is fine
  - d. Panel Title should be *Security Hacking Attempts*
3. It should look like this if **New**:

### Save As Dashboard Panel

Dashboard:  New  Existing

Dashboard Title:

Dashboard ID <sup>?</sup>:   
The dashboard ID can only contain letters, numbers, dashes, and underscores. Do not start the dashboard ID with a period.

Dashboard Description:

Dashboard Permissions:  Private  Shared in App

---

Panel Title:

Panel Powered By <sup>?</sup>:

Drilldown <sup>?</sup>:

Panel Content:  Statistics  Column Chart

4. Or like this if **Existing**:

### Save As Dashboard Panel

Dashboard:  New  Existing

---

Panel Title:

Panel Powered By <sup>?</sup>:

Drilldown <sup>?</sup>:

Panel Content:  Statistics  Column Chart

5. Click Save

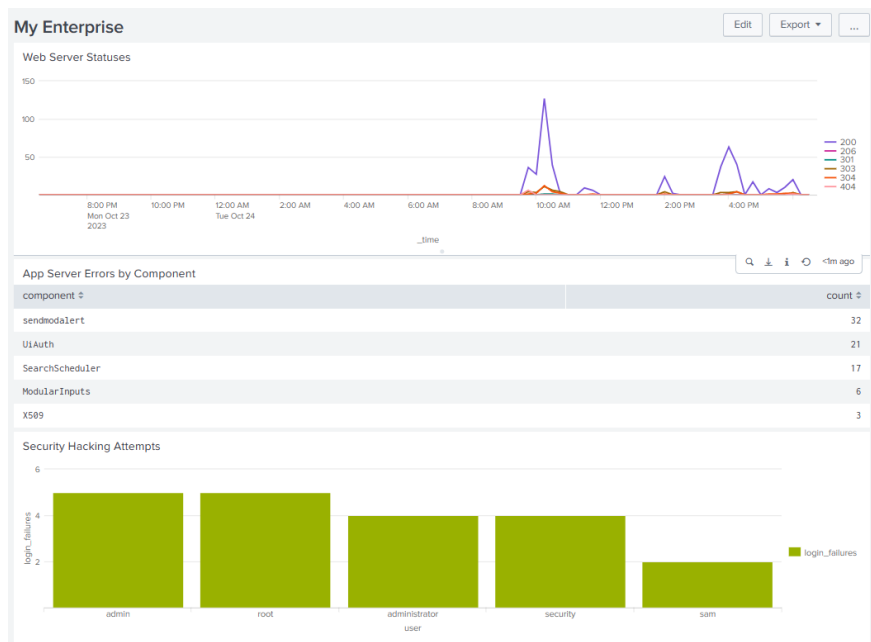
You can View Dashboard to see what it looks like. Tada! You're a splunker now: SPL, dashboards, you're doing it!

## Conclusion

There you have it. That's how Splunk can watch your enterprise, and you can build dashboards correlating things together. This was all really simple, and using built-in data. Just wait until you start getting your own data in and doing even more correlations. You'll soon find yourself writing gnarly SPL that scares children and small animals.



If you went through all of the sections, then you should end up with a pretend enterprise dashboard like this one:



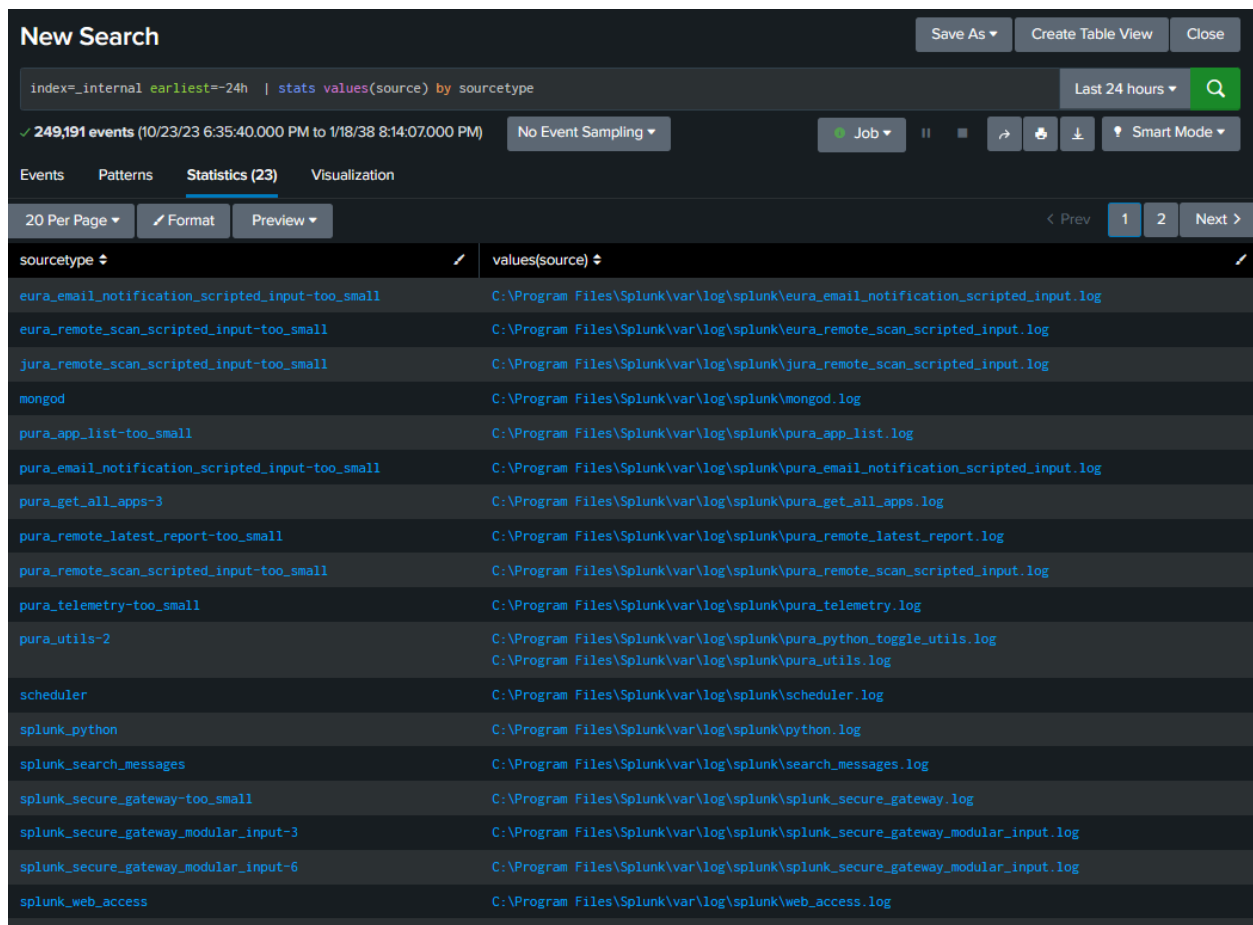
## Appendix: Where are these logs coming from?

Just like your web server, application server, and other servers/devices in your enterprise creates logs, so does Splunk. Under the hood within the Splunk install is a directory where Splunk is writing these, and then indexing them, just like you would configure it to do for all of your stuff. Splunk also takes care of rotating/limiting the amount of disk these internal logs use so you typically don't have to worry about managing them unless you have really specific storage requirements. If you want to poke around these files take a look at the *source* field for the data in the searches.

You can also run this SPL to see what log files (the *source*) correspond to what *sourcetypes* within your environment:

```
index=_internal earliest=-24h | stats values(source) by sourcetype
```

Here's a screenshot of my results:



The screenshot shows a Splunk search interface with the following details:

- Search Title:** New Search
- Search Query:** `index=_internal earliest=-24h | stats values(source) by sourcetype`
- Results:** 249,191 events (10/23/23 6:35:40.000 PM to 1/18/38 8:14:07.000 PM)
- Navigation:** 20 Per Page, Format, Preview, < Prev, 1, 2, Next >
- Table Columns:** sourcetype, values(source)
- Table Rows:** 23 rows of data, each showing a sourcetype and its corresponding log file path.

sourcetype	values(source)
aura_email_notification_scripted_input-too_small	C:\Program Files\Splunk\var\log\splunk\aura_email_notification_scripted_input.log
aura_remote_scan_scripted_input-too_small	C:\Program Files\Splunk\var\log\splunk\aura_remote_scan_scripted_input.log
jura_remote_scan_scripted_input-too_small	C:\Program Files\Splunk\var\log\splunk\jura_remote_scan_scripted_input.log
mongod	C:\Program Files\Splunk\var\log\splunk\mongod.log
pura_app_list-too_small	C:\Program Files\Splunk\var\log\splunk\pura_app_list.log
pura_email_notification_scripted_input-too_small	C:\Program Files\Splunk\var\log\splunk\pura_email_notification_scripted_input.log
pura_get_all_apps-3	C:\Program Files\Splunk\var\log\splunk\pura_get_all_apps.log
pura_remote_latest_report-too_small	C:\Program Files\Splunk\var\log\splunk\pura_remote_latest_report.log
pura_remote_scan_scripted_input-too_small	C:\Program Files\Splunk\var\log\splunk\pura_remote_scan_scripted_input.log
pura_telemetry-too_small	C:\Program Files\Splunk\var\log\splunk\pura_telemetry.log
pura_utils-2	C:\Program Files\Splunk\var\log\splunk\pura_python_toggle_utils.log C:\Program Files\Splunk\var\log\splunk\pura_utils.log
scheduler	C:\Program Files\Splunk\var\log\splunk\scheduler.log
splunk_python	C:\Program Files\Splunk\var\log\splunk\python.log
splunk_search_messages	C:\Program Files\Splunk\var\log\splunk\search_messages.log
splunk_secure_gateway-too_small	C:\Program Files\Splunk\var\log\splunk\splunk_secure_gateway.log
splunk_secure_gateway_modular_input-3	C:\Program Files\Splunk\var\log\splunk\splunk_secure_gateway_modular_input.log
splunk_secure_gateway_modular_input-6	C:\Program Files\Splunk\var\log\splunk\splunk_secure_gateway_modular_input.log
splunk_web_access	C:\Program Files\Splunk\var\log\splunk\web_access.log

*This page intentionally left blank Except for this text.*