Running Splunk in an Air-gapped environment

Steve Schohn
Staff Sales Engineer | Splunk

splunk> .conf19

Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or plans of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results may differ materially. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, it may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements made herein.

In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only, and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionalities described or to include any such feature or functionality in a future release.

The Challenge

Nobody said it was easy

splunk> .conf19

Are you in the right place?

You have an architecture that looks like this:

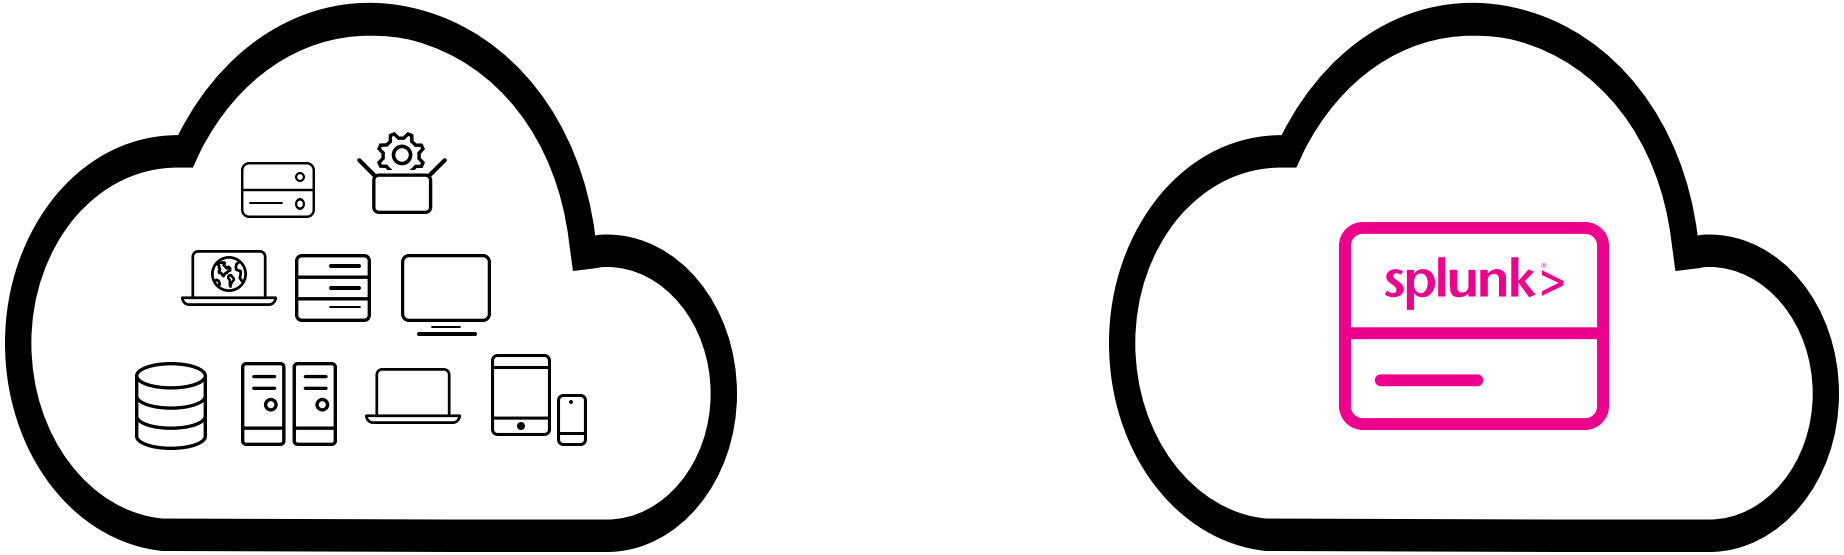

Data sources are on one network. **Splunk is on another.** You cannot use forwarders to get data from one to another.

Are you in the right place?

Or more likely this:

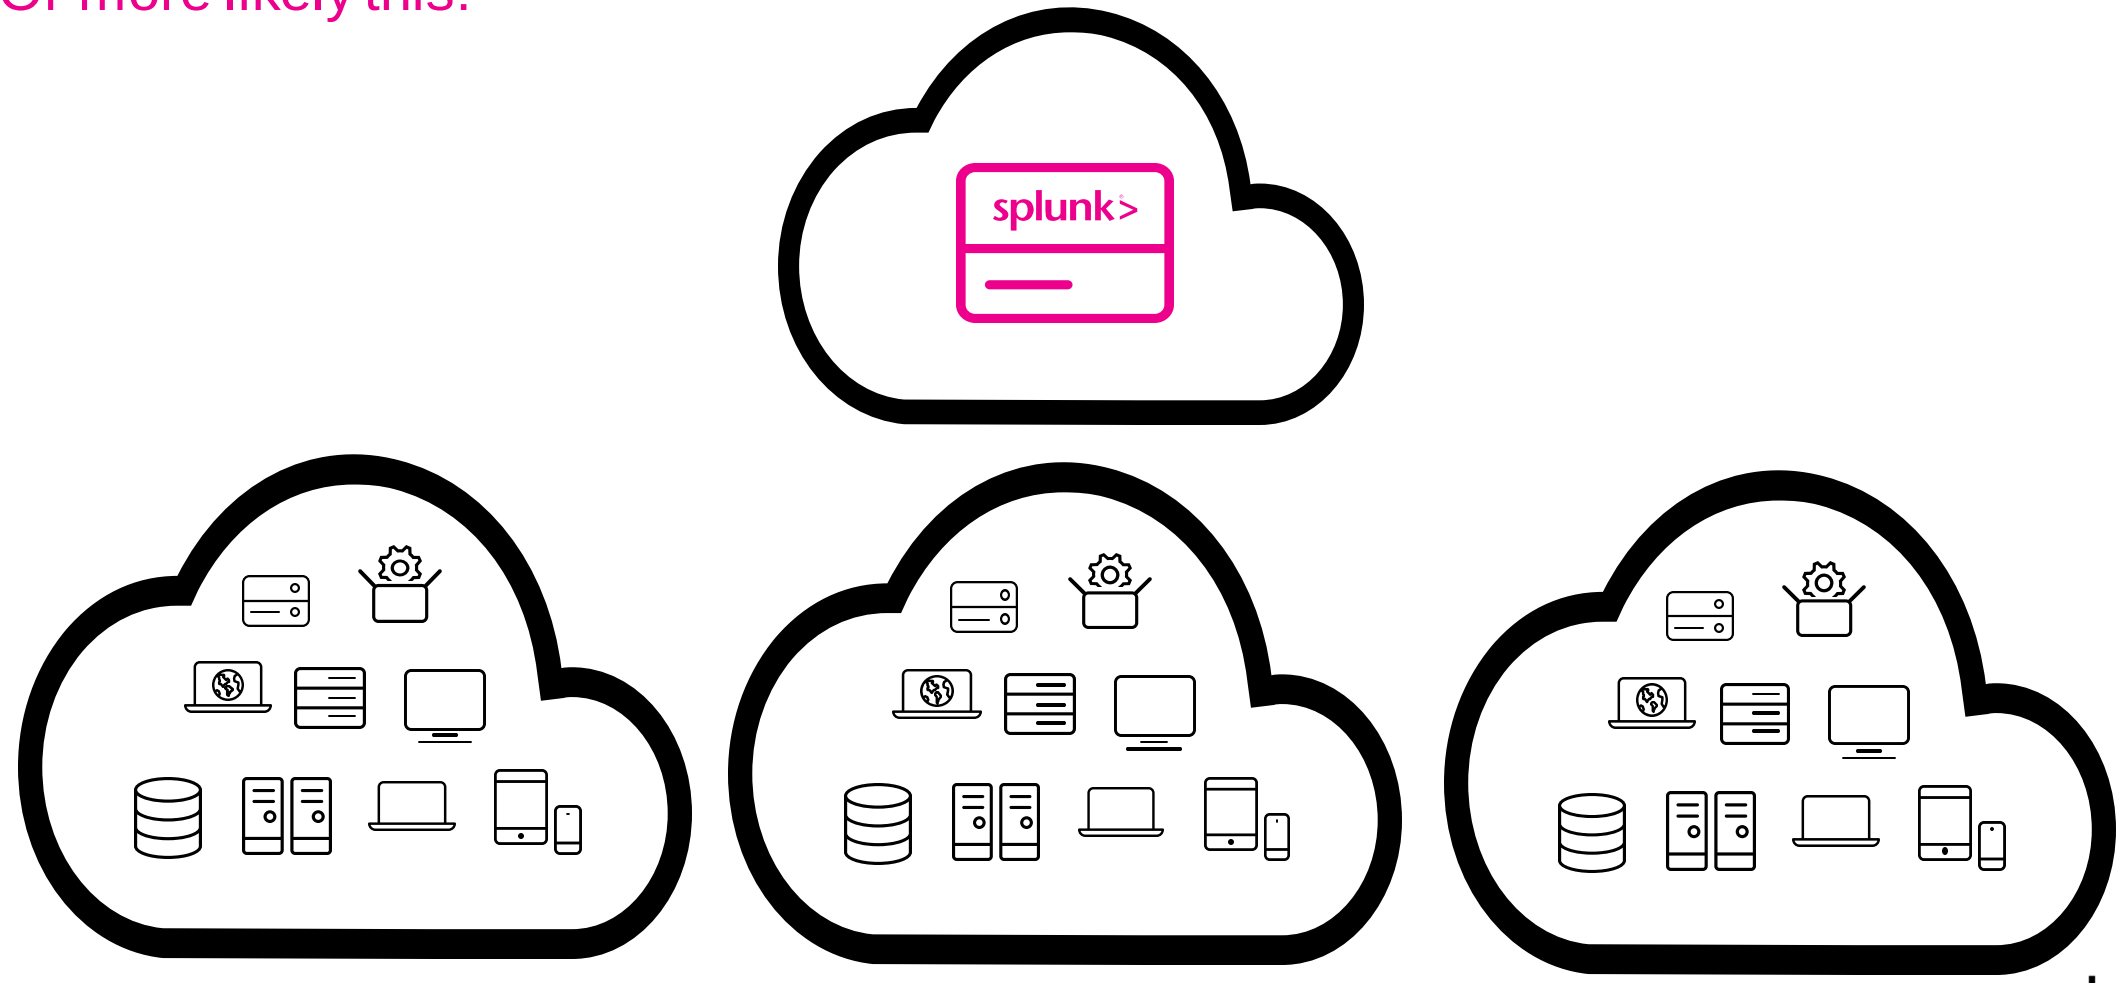

A Land Without TCP

- There is one-way connectivity from the network with data sources to the network running Splunk.
- The connectivity is limited to a few unidirectional protocols: FTP, UDP, etc
- If a true sneakernet: that way lies madness, and is outside the scope of this talk.
- TCP is specifically not an option: the SYN-SYNACK-ACK is unavailable, and therefore so is the UF->Indexer connection

Data can go up the stack, but never down

The Goal

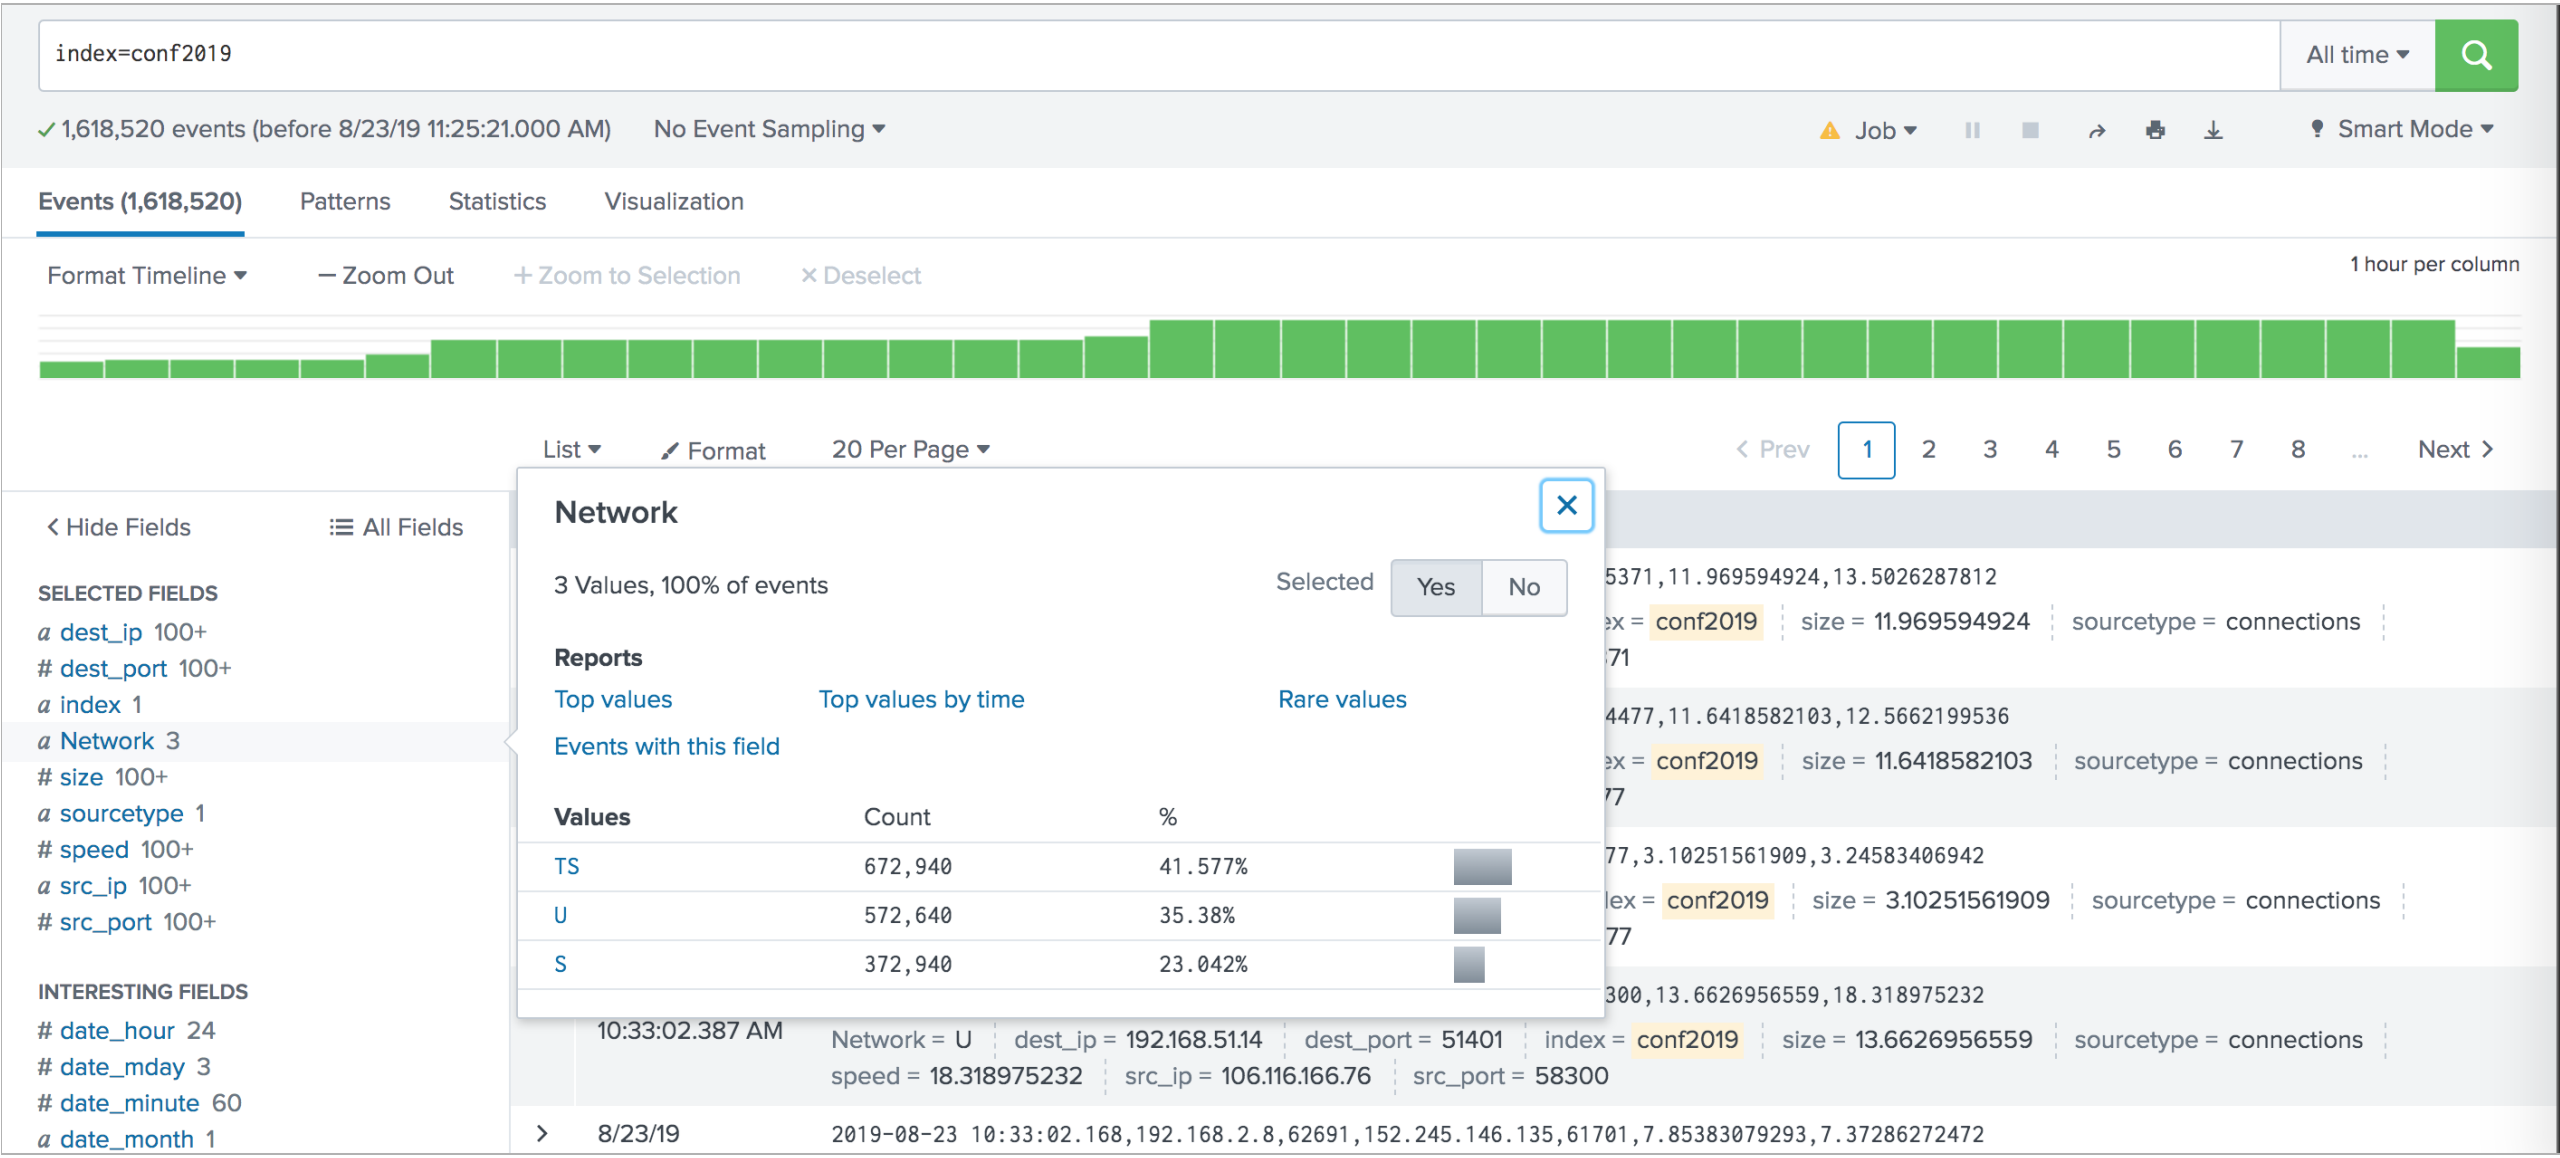

splunk> .conf19

Some quick definitions

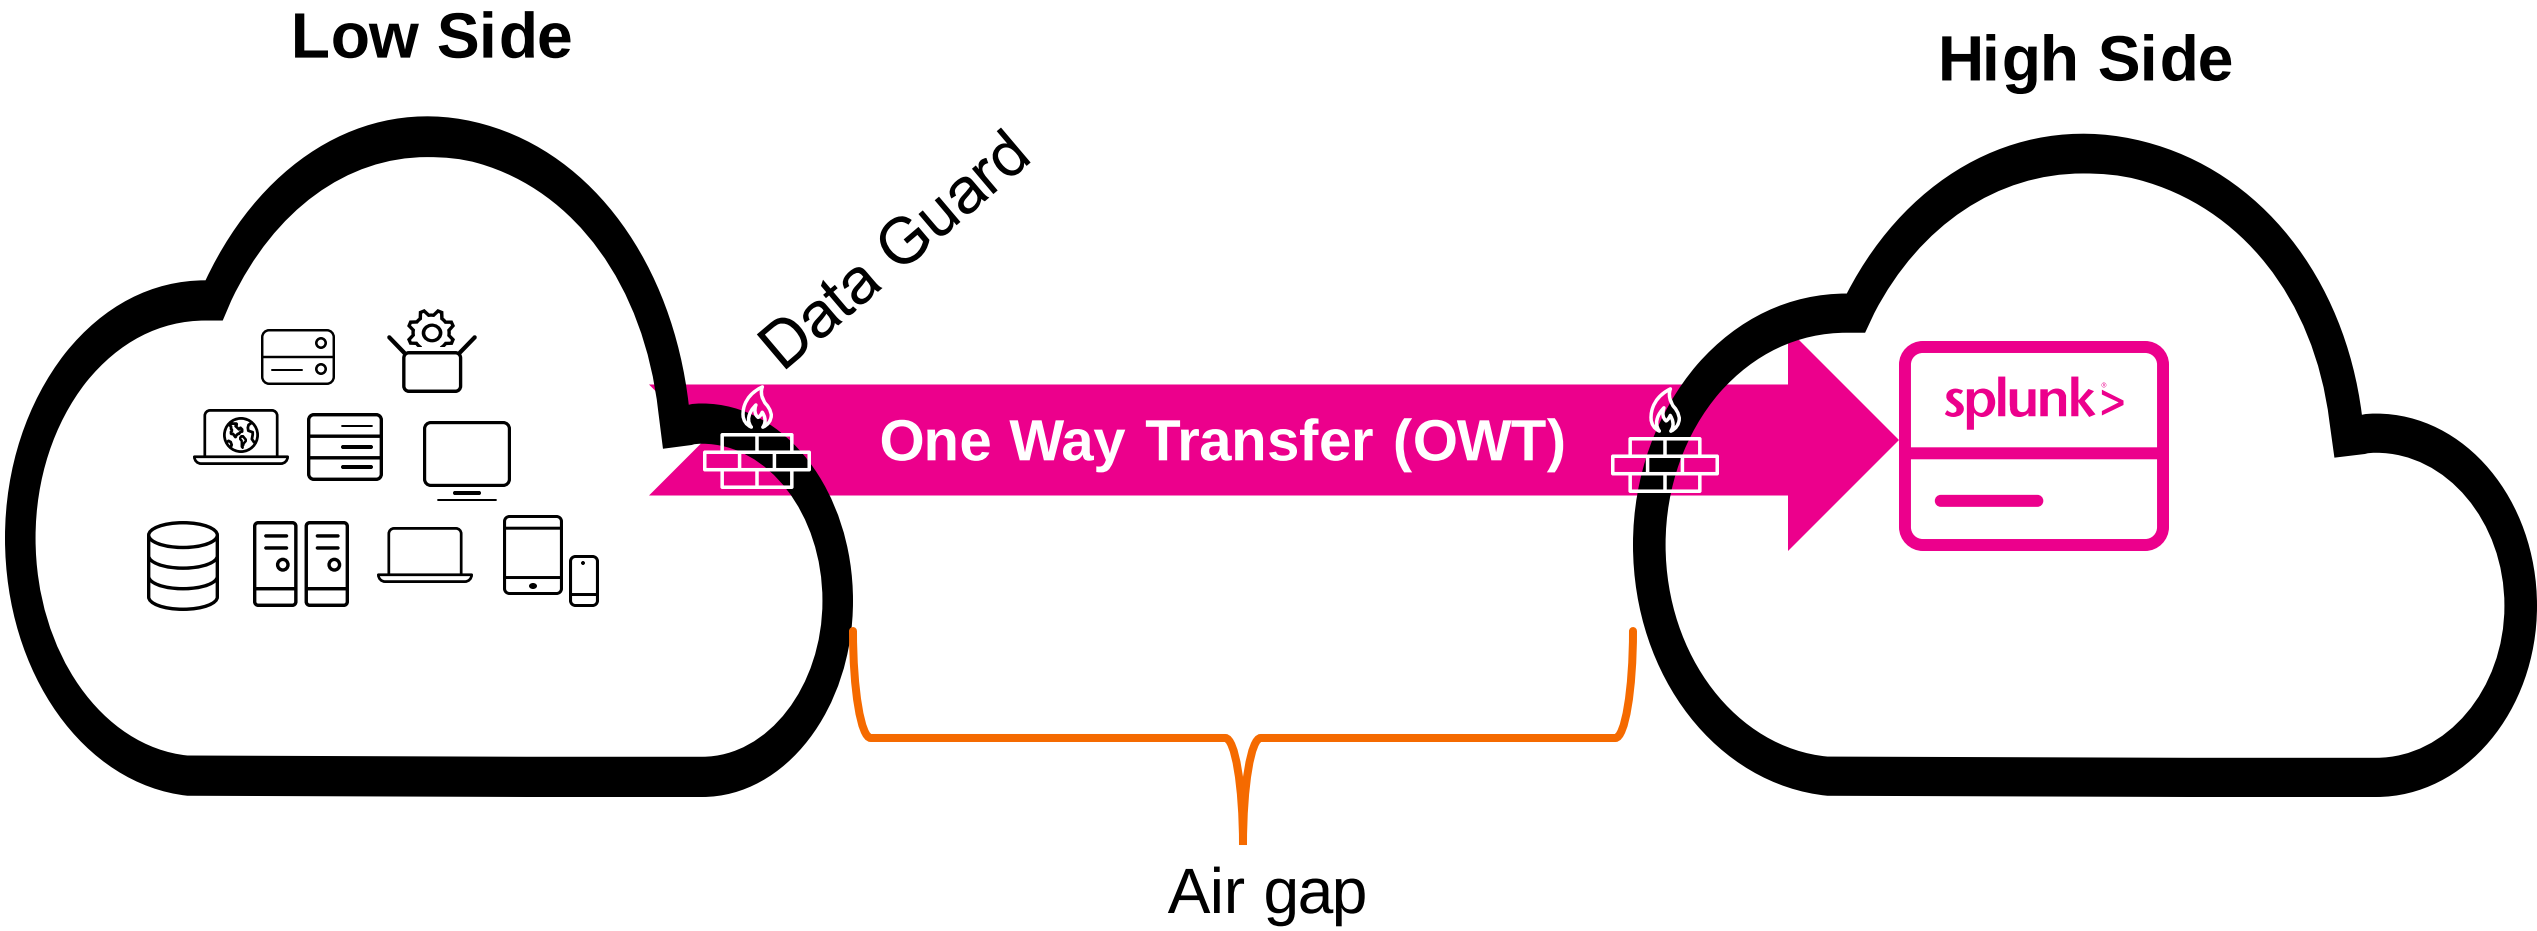

Getting Data In, Across the Air-gap

A Tour of What I've Seen

splunk> .conf19

Method 1: Save data low as files, move files through OWT and index high

Use a combination of syslog, WEF, and native formats to push individual files through OWT

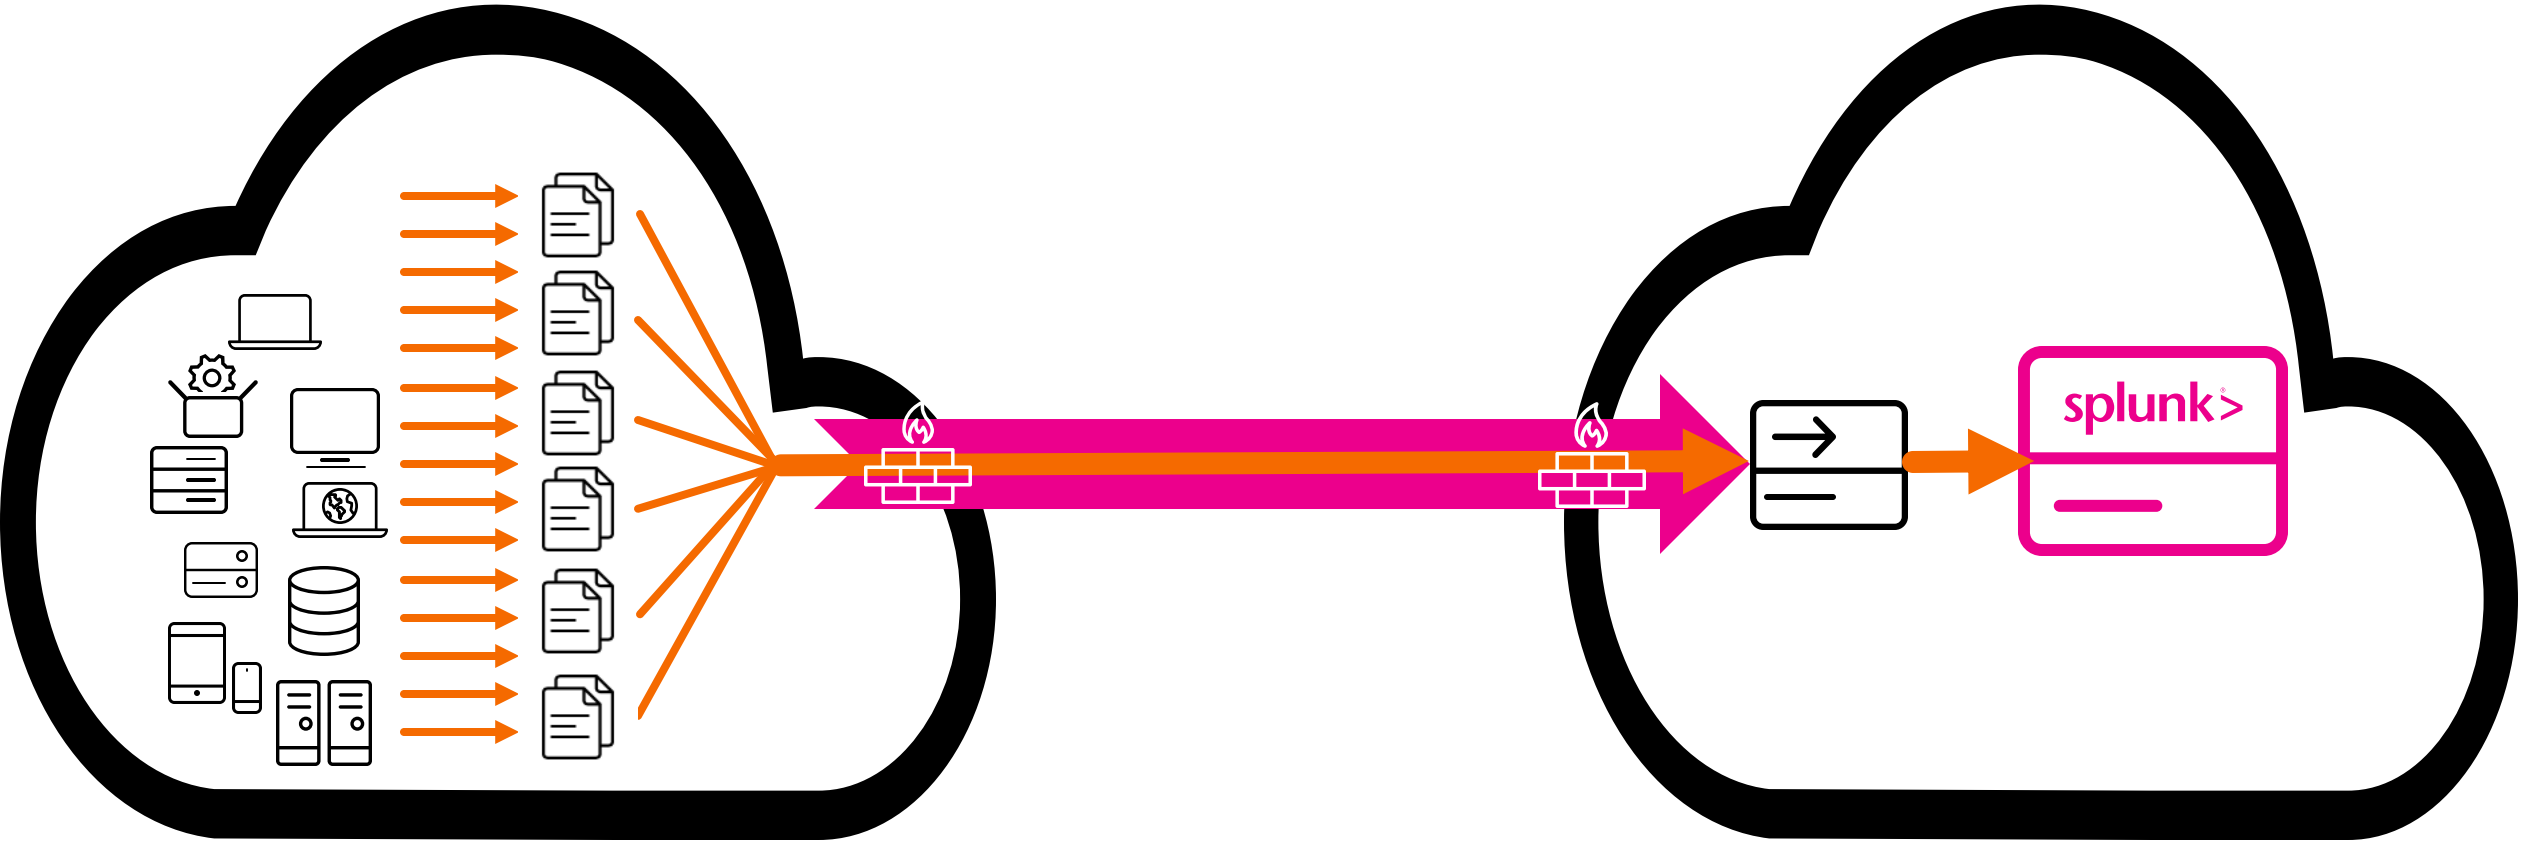

Should you do this?

No!
A resounding no!
Please no!

splunk> .conf19

Method 1: Save data low as files, move files through OWT and index high

Pros

Often the first thing people try?

Cons

I don't care how good your naming convention is: at scale, eventually you will mess up the originating metadata

- Data will end up on high associated with the wrong host, source, and/or sourcetype

It's a nightmare to maintain, let alone continually scale

- Imagine having thousands of hosts, and trying to copy up all the files in each's /var/log/ while keeping the host correct

Windows events are especially ugly this way

Method 2: UDP data across OWT, index it high

Send data via syslog/UDP across OWT. Receive on the high side like standard UDP inputs.

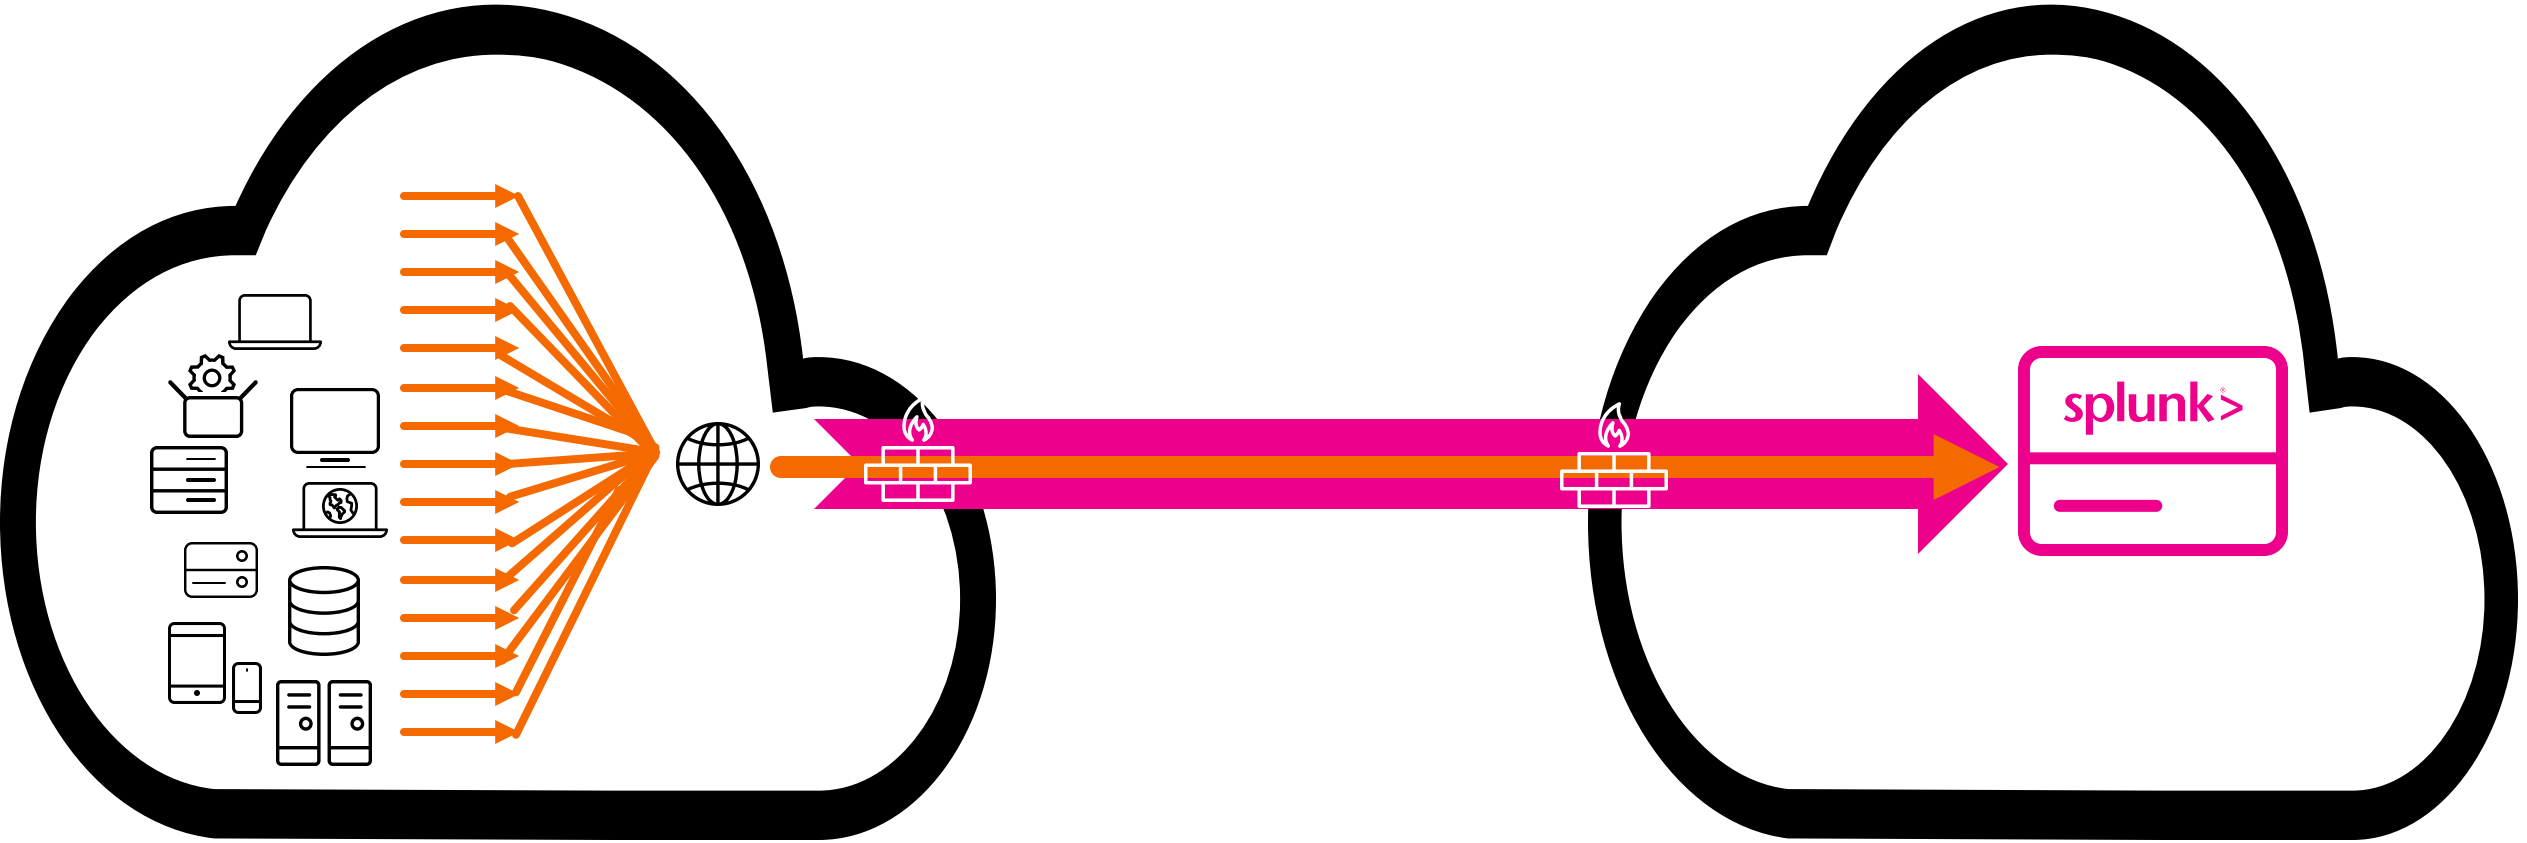

Should you do this?

Almost definitely not!

splunk> .conf19

Method 2: UDP data across OWT, index it high

In short: You will lose data

All the reasons not to use UDP as described in Jeff Champagne's "Worst Practices and How to Avoid Them" past .conf talks come into play here:

Lossless data transmission over UDP does not exist

- UDP lacks error control AND flow control
  - Delivery cannot be guaranteed
  - Packets may be lost
    - They never arrived due to network issues
    - They were dropped due to a busy destination
  - Retransmits can result in duplicates

Method 2: UDP data across OWT, index it high

Pros:
The one use case where this does make sense: when you care about real time way more than you care about potential data loss

If you are ok with getting less than 100% of your data, but you want it searchable immediately, this can be made to work

No low side hardware deployment is necessary

Cons:
Works best when you have mostly network device data. Endpoint host data will get messy in syslog form – you'll have whitespace and truncation issues.

Again, you're probably going to have host/source/sourcetype pain: make sure each feed comes in on its own port

And again, **YOU WILL LOSE DATA**

Method 3: Index data low, move the bucket files up

Set up indexers on both low side and high side. Index low side data as normal. On a schedule, move up the low side bucket files into the high side Splunk indexes.

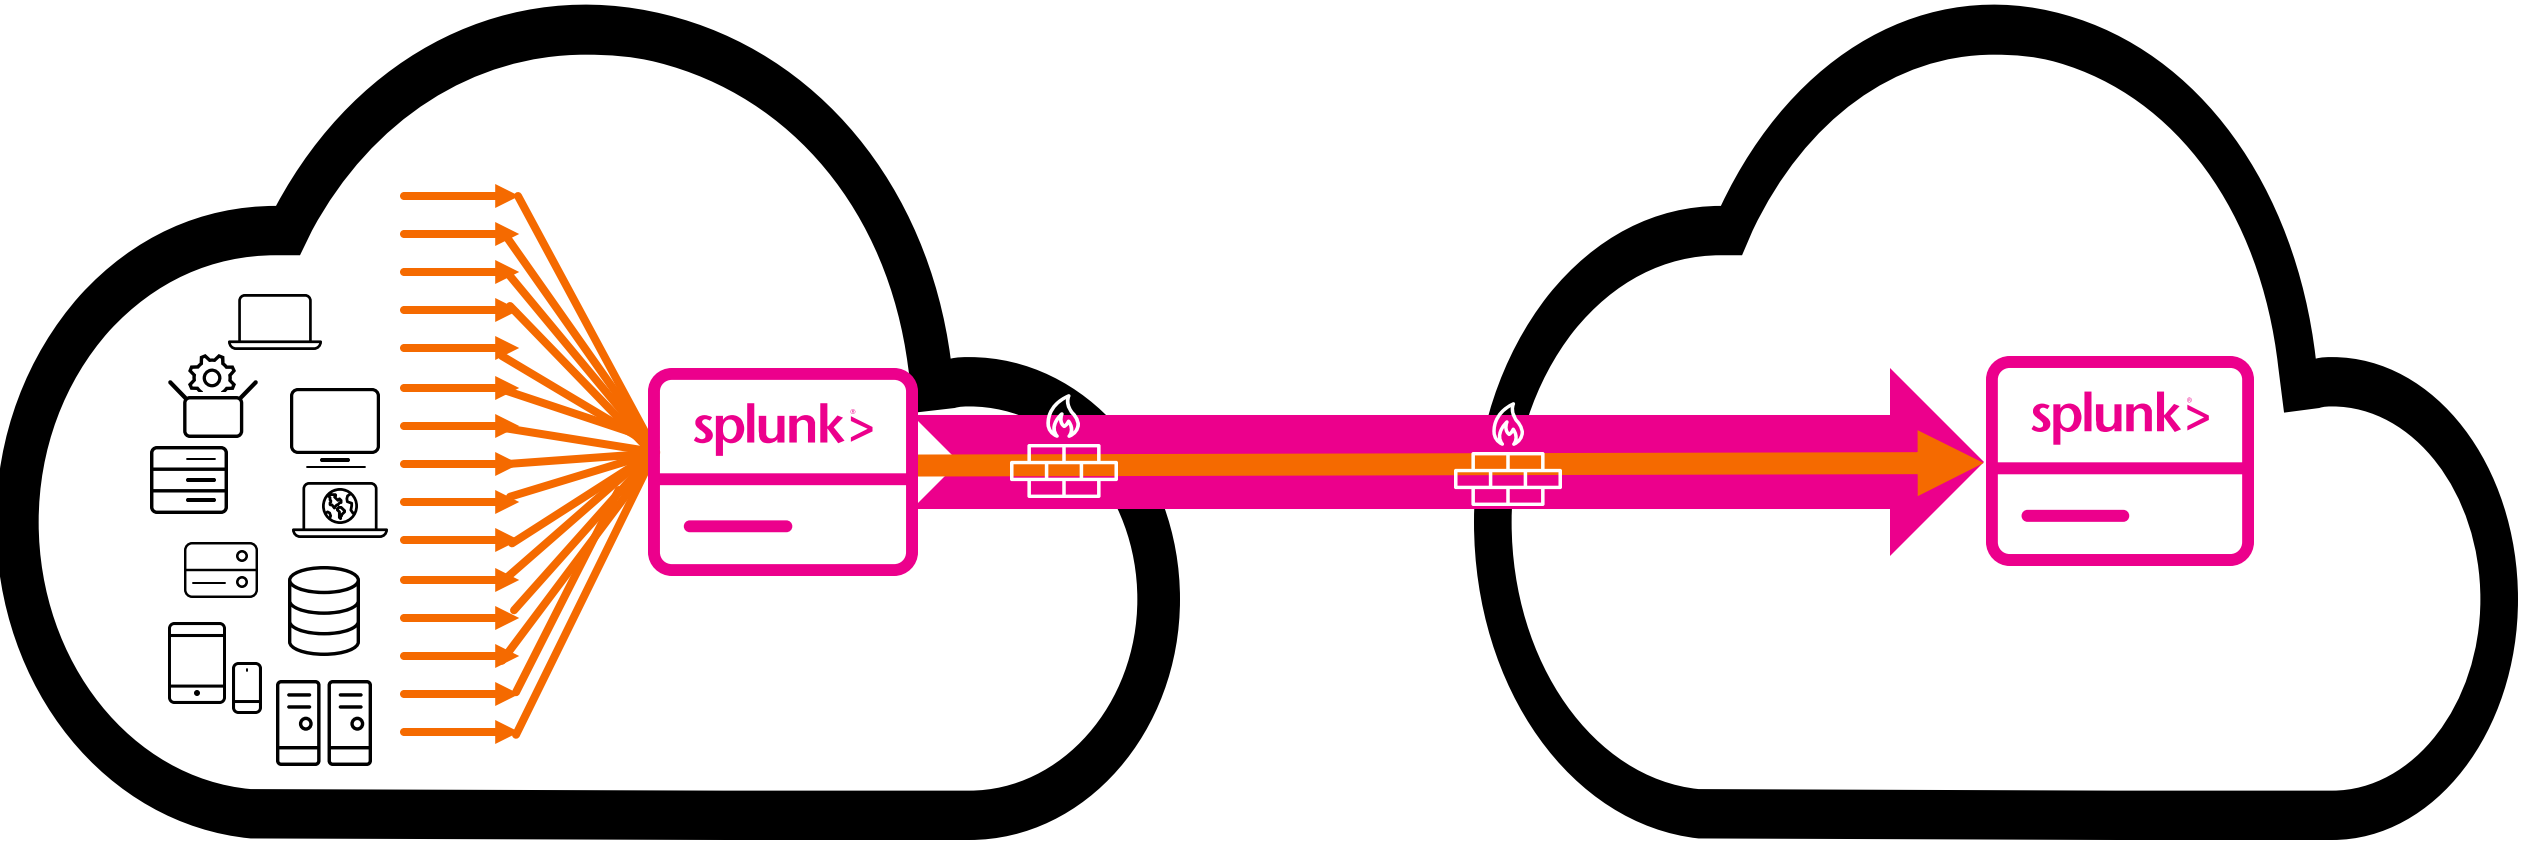

Should you do this?

Probably not!

splunk> .conf19

Method 3: Index data low, move buckets up

Pros:

You're guaranteed fidelity of the data

Conceptually it's the most obvious to explain and maintain

If you don't need real time data, and are ok with a daily restart, this can be a workable model

Cons:

You need to wait for low-side buckets to roll to warm before moving them

The buckets moved to high won't be searchable without a restart of splunkd

You need to ensure you don't have bucket id collisions: each bucket must have a unique id number. So prepare to rename the bucket, or have separate indexes on low and high.

Very difficult to move buckets from clustered systems -> non-clustered systems, and vice versa

Method 3: A Quick Addendum!

But but but Splunk Answers: https://answers.splunk.com/answers/838/how-can-you-add-move-a-bucket-without-restarting-splunkd.html

How can you add/move a bucket without restarting splunkd?

6

Rough first take:

1. In the <indexname>/db directory, delete the file .bucketmanifest
2. In the <indexname>/db directory, create the file (0 bytes works) meta.dirty

If we get into goat sacrifice territory, try also deleting .metamanifest. Step 2 should render that unnecessary. These files and their associated data should get rebuilt on need by search activity.

Answer by jrodman [Splunk]

Mar 23, 2010 at 11:32 PM

This is totally unsupported! One day a small upgrade will break this hack, and you'll be in a world of pain.

Method 4: Index Low, Read Data out to a File

Instead of moving the buckets up, use a Splunk search to output a flat file of all _raw data on a set interval. Move the file up for index.

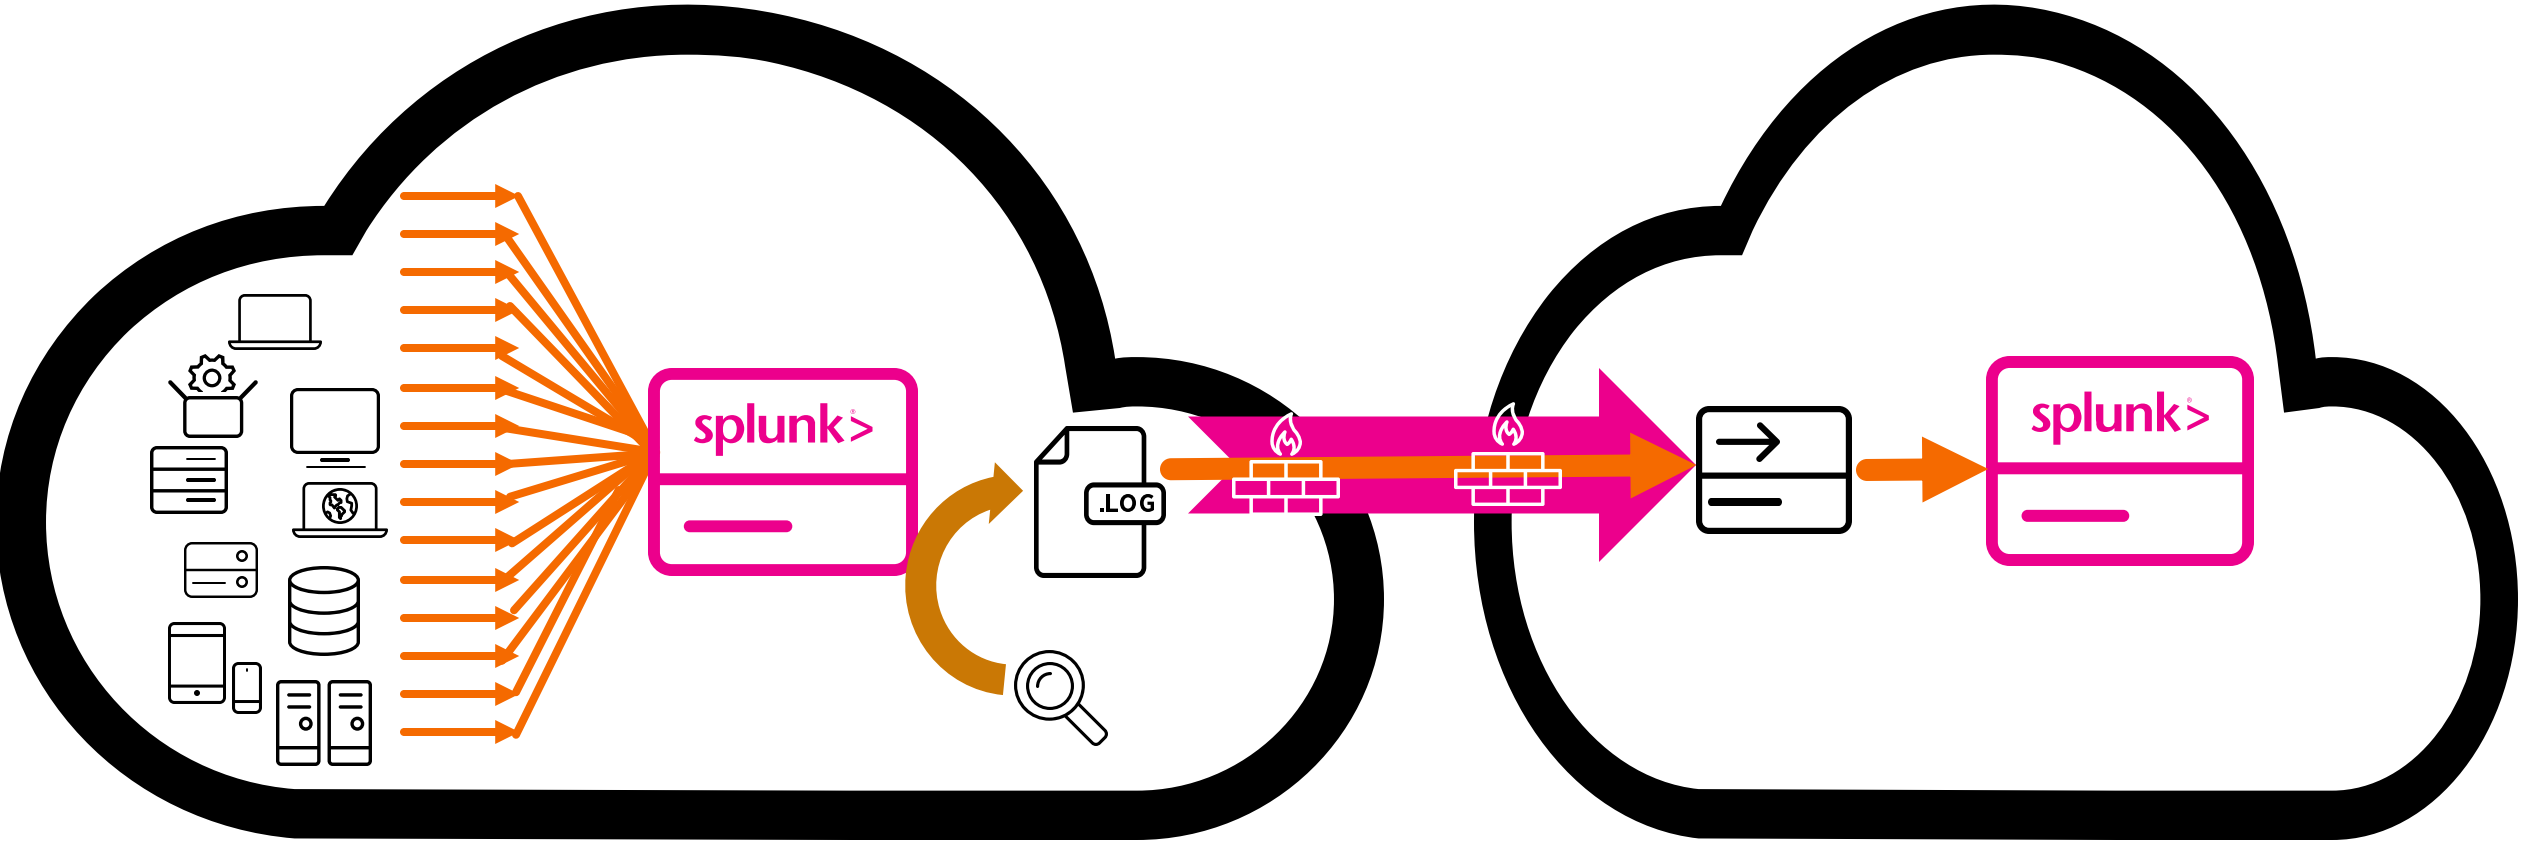

Should you do this?

Maybe!

Method 4: Index Low, Read Data out to a File

Using an example search like this… (done from CLI, as it's normally part a cron'd script):

```
$ bin/splunk search 'index=* earliest=-2h@h latest=-1h@h

| eval headeroutput = "***SPLUNK*** _time=\"".  _time."\" host=\"".host."\"
sourcetype=\"".sourcetype."\" source=\"".source."\" index=\"".index."\"-
splitonthis-"._raw

| table headeroutput

| makemv delim="-splitonthis-" headeroutput

| mvexpand headeroutput' -maxout 0 -header F
```

(Note: You'll need to tweak earliest/latest to match your cron schedule)

Method 4: Index Low, Read Data out to a File

- Low-side search output looks like:

```
***SPLUNK*** _time="1566568055.538" host="TS_Network" sourcetype="connections" source="connections_ts.log" index="conf2019"
2019-08-23 09:47:35.538,192.168.63.35,50563,127.194.146.171,64217,6.1696424199,7.30401128882
***SPLUNK*** _time="1566568055.658" host="TS_Network" sourcetype="connections" source="connections_ts.log" index="conf2019"
2019-08-23 09:47:35.658,45.13.47.14,52237,192.168.16.37,49262,7.51560592883,8.01660710924
***SPLUNK*** _time="1566568055.837" host="TS_Network" sourcetype="connections" source="connections_ts.log" index="conf2019"
2019-08-23 09:47:35.837,45.69.134.72,56663,192.168.186.45,63592,6.596969584,9.06063788109
***SPLUNK*** _time="1566568056.196" host="TS_Network" sourcetype="connections" source="connections_ts.log" index="conf2019"
2019-08-23 09:47:36.196,111.234.157.160,52066,192.168.168.163,55095,4.58686937187,4.56265725655
```

- On high-side input, use the HEADER_MODE = ALWAYS mode in props.conf (https://docs.splunk.com/Documentation/Splunk/latest/admin/Propsconf#Header_Processor_configuration)

```
HEADER_MODE = <empty> | always | firstline | none
* Determines whether to use the inline ***SPLUNK*** directive to rewrite
  index-time fields.
  * If "always", any line with ***SPLUNK*** can be used to rewrite
    index-time fields.
  * If "firstline", only the first line can be used to rewrite
    index-time fields.
  * If "none", the string ***SPLUNK*** is treated as normal data.
  * If <empty>, scripted inputs take the value "always" and file inputs
    take the value "none".
* This setting applies at input time, when data is first read by Splunk
  software, such as on a forwarder that has configured inputs acquiring the
  data.
* Default: <empty>
```

Method 4: Index Low, Read Data out to a File

Pros:

Moving events eliminates the challenges we discussed previously: cluster agnostic, bucket_ids a non-issue, metadata properly captured

- Your generating search will return the raw event and preserve the host/source/sourcetype/_time in an inserted header row for each event

There's nothing additional to install or maintain, outside the cron script

Cons:

This will eventually struggle to keep up at high enough ingest, at which point the low-side Splunk (which lived-reality shows tends to be underpowered) cannot read out the data fast enough over the specified interval

This can be brittle: you need to think about how to handle if/when the low-side server skips a generating search

This works very well at lower volumes if your use cases allow you to tolerate data coming in as a batch on a search schedule

Method 5: Index Low, NiFi out and up, NiFi back in

Instead of using a a Splunk search to output a log file of low side data, use Apache Niagra Files (NiFi).

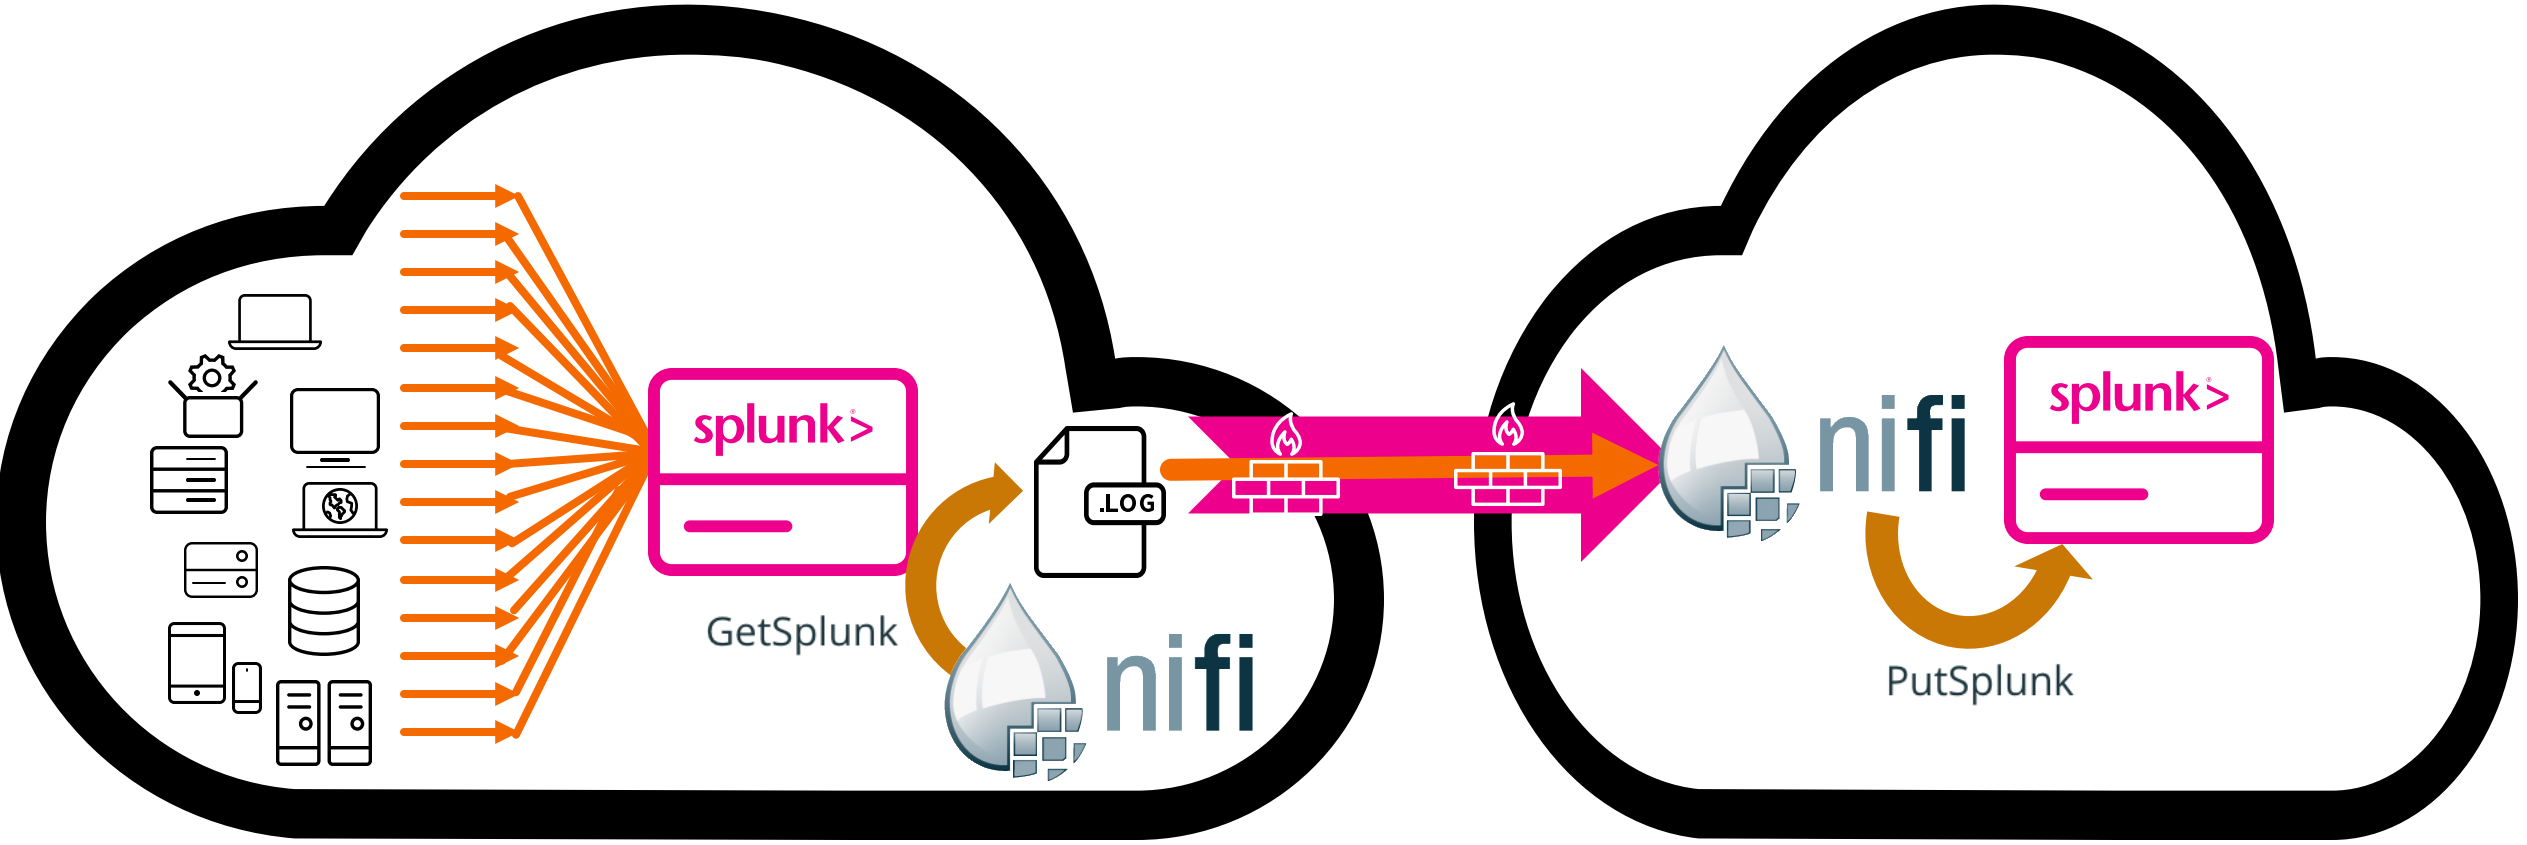

Should you do this?

If you can!

Method 5: Index Low, NiFi out and up, NiFi back in

Pros:
NiFi is scalable, resilient, and highly-available

Customers are using NiFi to push multiple TB through OWTs each day

Biggest customers have started using NiFi on low, but a HF receiver on high, using HEC to push into their indexers

Cons:
With the added capabilities comes added admin complexity

NiFi is a separate system to learn, install, and maintain

If your team is capable enough to manage a multi-TB deployment of Splunk, it's also capable enough to manage NiFi

Method 6: Data Stream Processor!

Instead of NiFi, use Data Stream Processor to write to a file

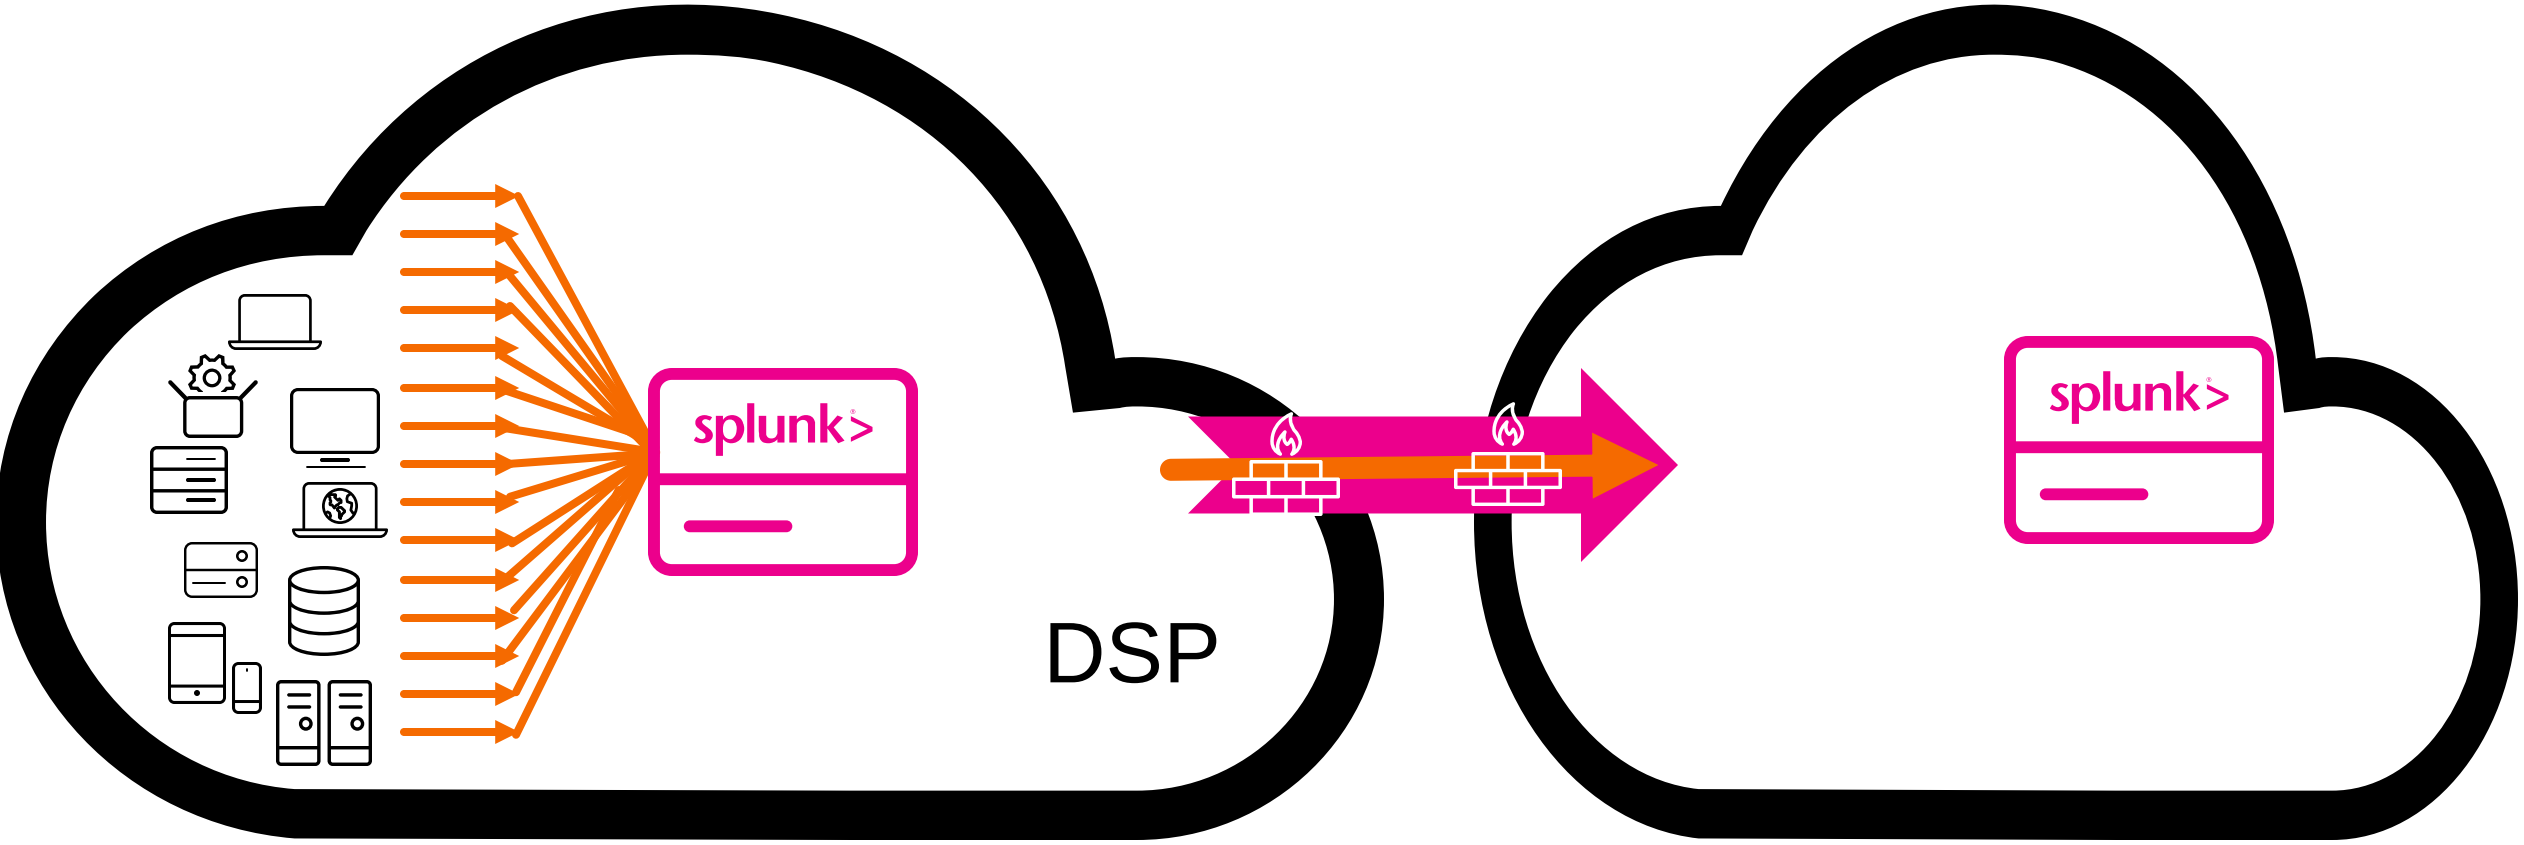

Should you do this?


It’s the holy grail.
THEN THE MIRACLE CAN HAPPEN TO YOU!

Method 6: Data Stream Processor

Pros:
Splunk built and supported: one system to rule them all

Resilient, fast, highly available stream processing engine

WYSIWYG pipelines for real-time data manipulation

Heterogenous ingress/egress to non-Splunk endpoints

Cons:
Just became generally available this week

While it's obviously the future, most customers are only now getting started and establishing best practices

Version released this week can't write directly to a file – send data to Kafka and pipe to file there – but that functionality is coming soon

Key Takeaways

No one ever said it would be this hard

1. Move events instead of raw files or buckets
2. The most mature customers use NiFi today
3. DSP is the future

Offline Resources

splunk> .conf19

Offline Documention

Splunk Documentation is fantastic, but the offline experience leaves something to be desired

There’s got to be a better way!

DEMO

Offline Documentation for Splunk App

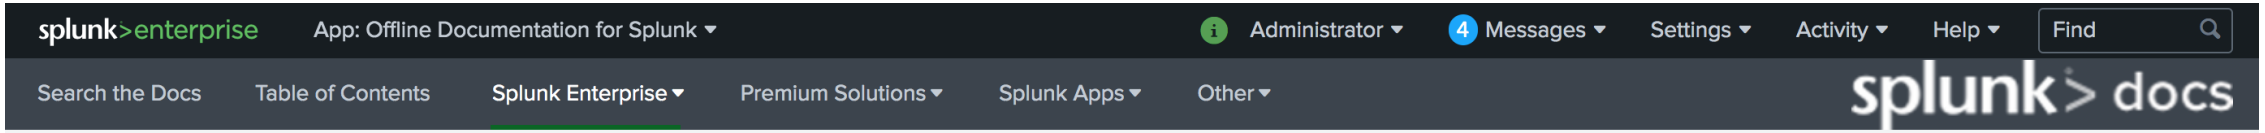

Move the index database

Edit | Export | ...

Managing Indexers and Clusters of Indexers

Indexing overview

1. Indexes, indexers, and indexer clusters
2. How indexing works
3. Index time versus search time
4. Install an indexer
5. Indexers in a distributed deployment

Manage indexes

1. About managing indexes
2. Create custom indexes
3. Remove indexes and indexed data
4. Manage pipeline sets for index parallelization
5. Optimize indexes
6. Use the monitoring console to view indexing performance

Manage index storage

1. How the indexer stores indexes
2. Configure index storage
3. Move the index database
4. Use multiple partitions for index data
5. Configure maximum index size
6. Set limits on disk usage
7. Reduce tsidx disk usage
8. Configure bloom filters
9. Determine which indexes.conf changes require restart

Move the index database

You can move the index database from one location to another. You do this by changing the path definition of `SPLUNK_DB` through the command-line interface of your operating system.

The procedures in this topic assume that the index database is in the default location, created during installation.

If you move individual indexes or parts of an index to separate locations, the procedures in this topic are not valid. For information on the structure of Splunk Enterprise indexes, see How the indexer stores indexes. For information on how to change the location for a single index, see Configure index storage.

Note: Although you can use Splunk Web to change the locations of individual indexes or index volumes, you cannot use it to change the default storage location of indexes, `SPLUNK_DB`.

For *nix users

Prerequisties

Make sure the target file system has at least 1.2 times the size of the total amount of raw data that you plan to index.

Steps

1. Create the target directory with write permissions for the user that Splunk Enterprise runs as. For example, if Splunk Enterprise runs as user "splunk", give it ownership of the directory:

```
mkdir /foo/bar
chown splunk /foo/bar/
```

For information on setting the user that Splunk Enterprise runs as, see Run Splunk Enterprise as a different or non-root user in the *Installation Manual*.

2. Stop the indexer:

```
splunk stop
```

3. Copy the index file system to the target directory:

```
cp -rp $SPLUNK_DB/* /foo/bar/
```

Offline Documentation for Splunk App

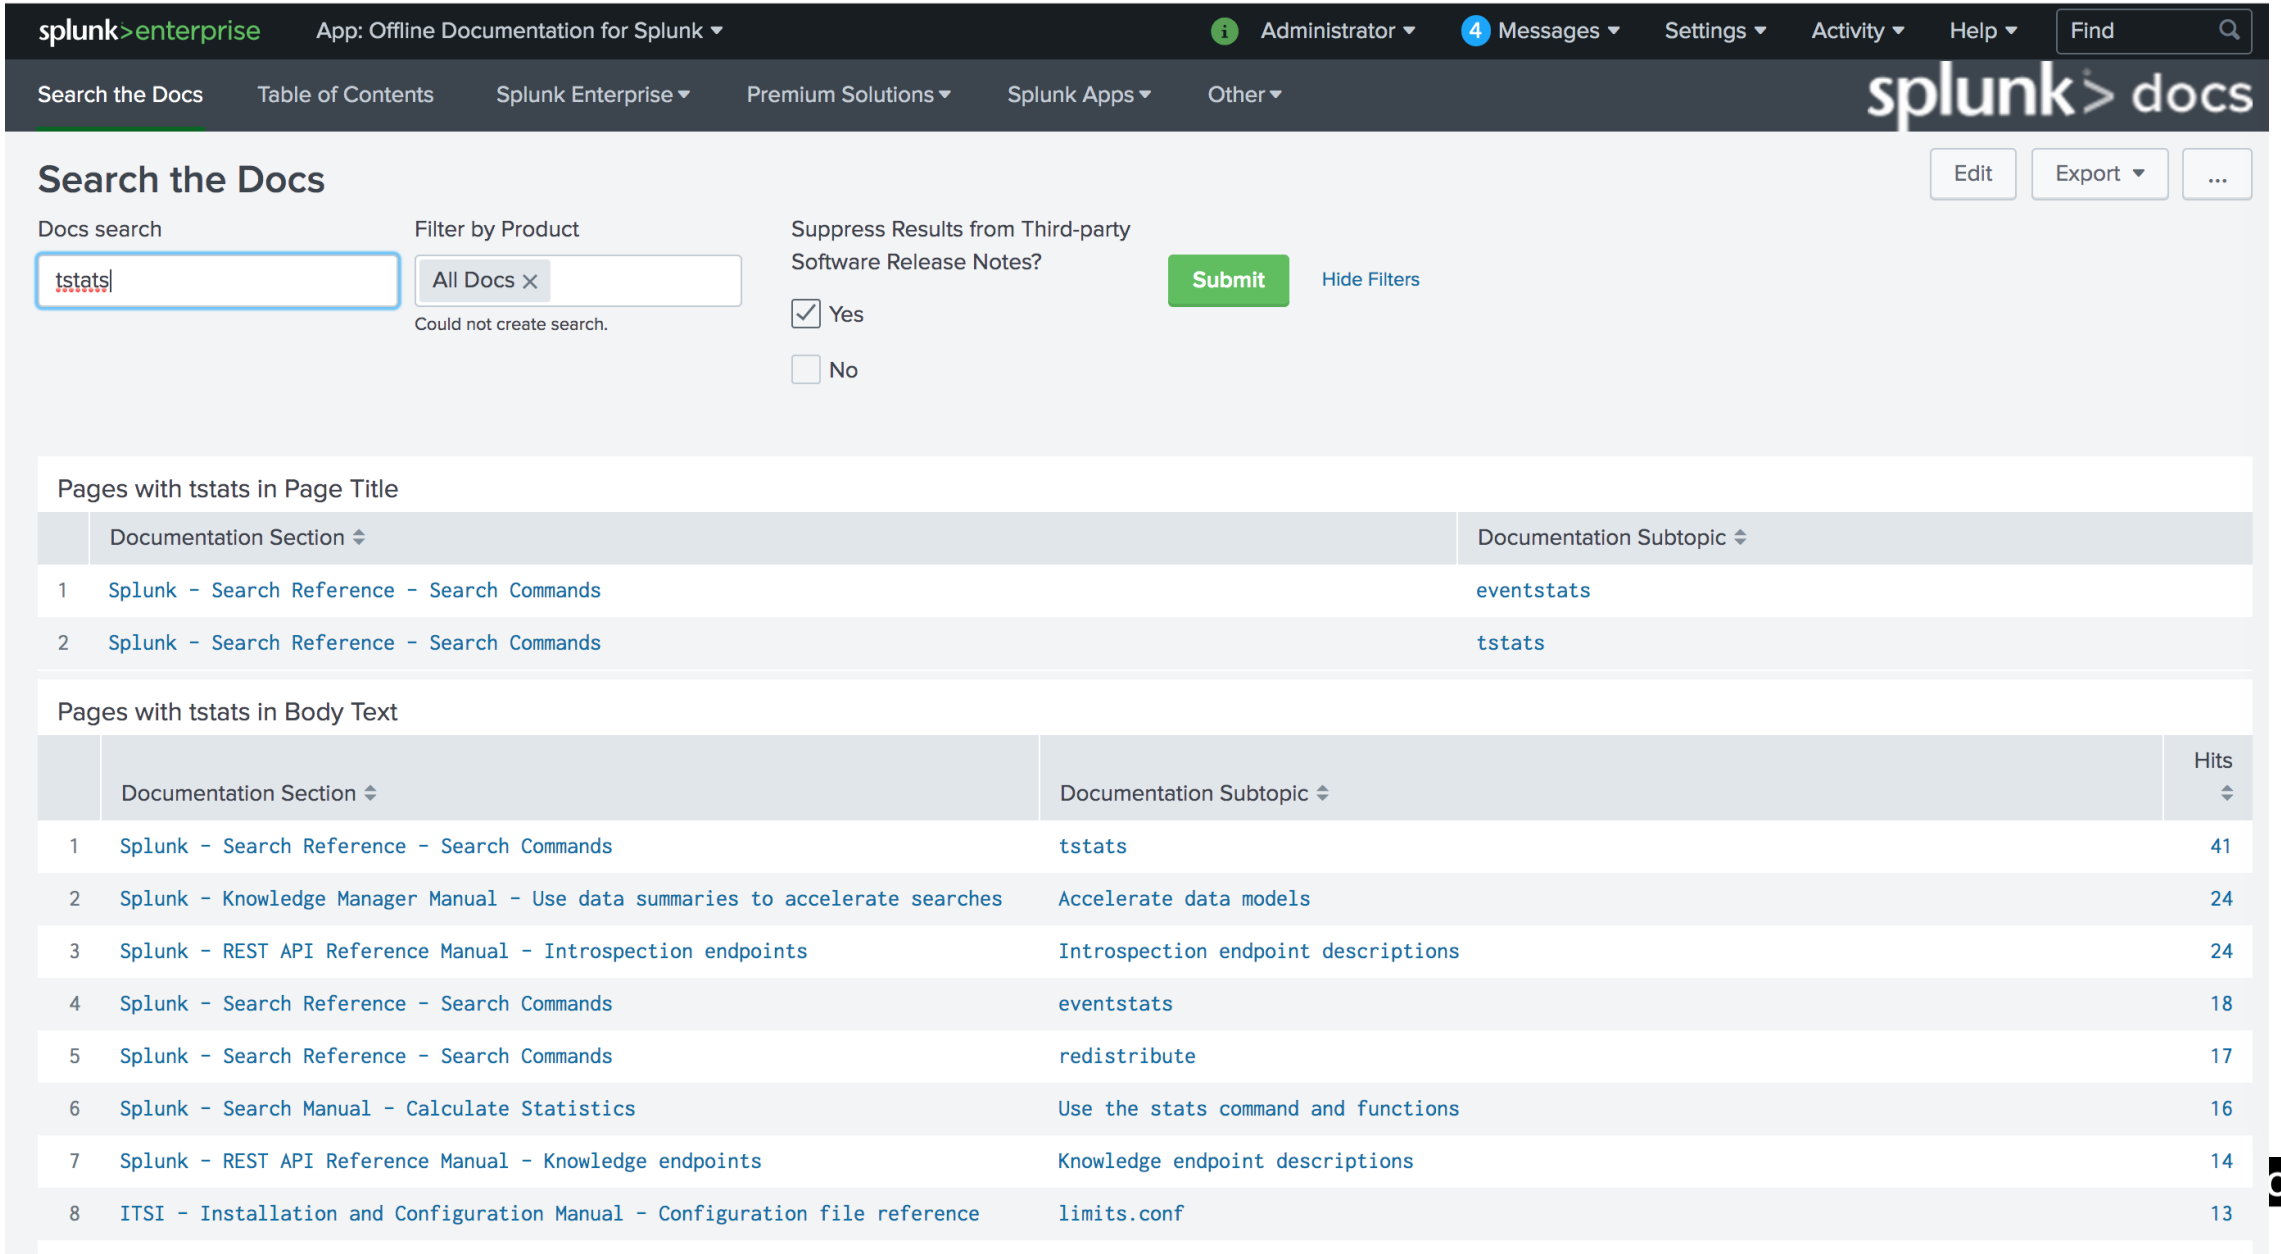

Offline Documentation for Splunk App

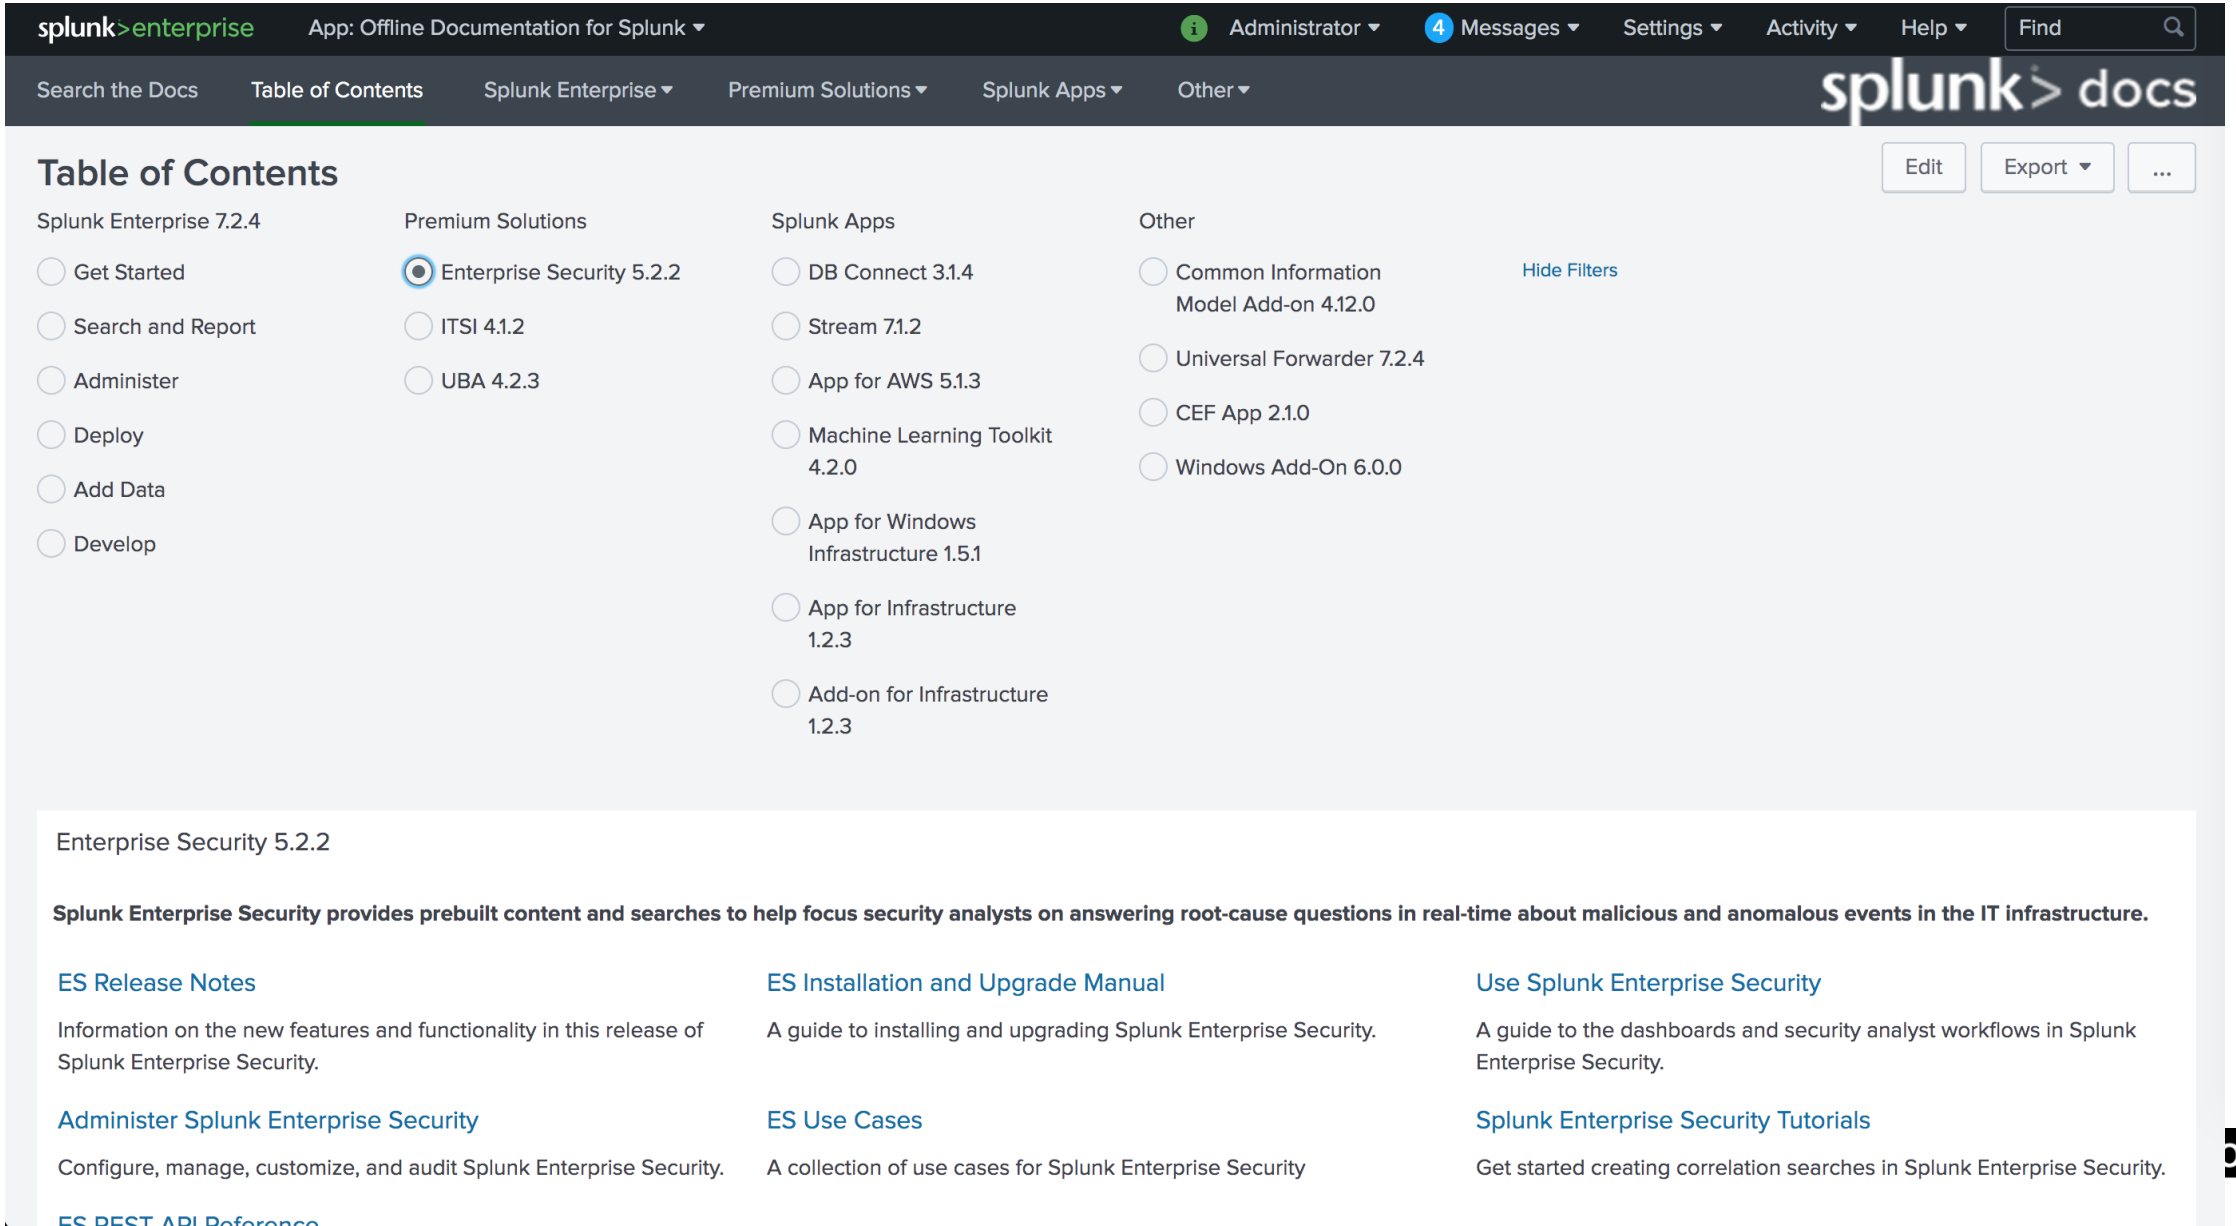

Offline Documentation

Not a Splunk product

Cannot guarantee maintenance schedule

Built by me to help out a few of my customers.

Hosted at http://docsapp.splunk-nsp.com

I've tried to stay on top of it, updating it after the first maintenance release after every major release (so I'll update it for 8.0.1)

Q&A

Steve Schohn | Staff Sales Engineer

splunk> .conf19

.conf19

splunk>

Thank You!

Go to the .conf19 mobile app to

RATE THIS SESSION