

Best Practices for Workload or Infrastructure Licensing

Jeff Meyers & Burch

Splunk, Inc.

Latest Slides: <http://splk.it/conf20-PLA1520>



Jeff Meyers

Director of Sales Engineering | Splunk. Inc.



Burch

Lead of Technical Guidance | Splunk, Inc.



Learning Objectives

Learning Objectives

1. When this licensing model is a proper fit

Learning Objectives

1. When this licensing model is a proper fit
2. How this licensing model will adjust your incentives and usage of Splunk

Learning Objectives

1. When this licensing model is a proper fit
2. How this licensing model will adjust your incentives and usage of Splunk
3. Best Practices that will save you money AND clean up your platform

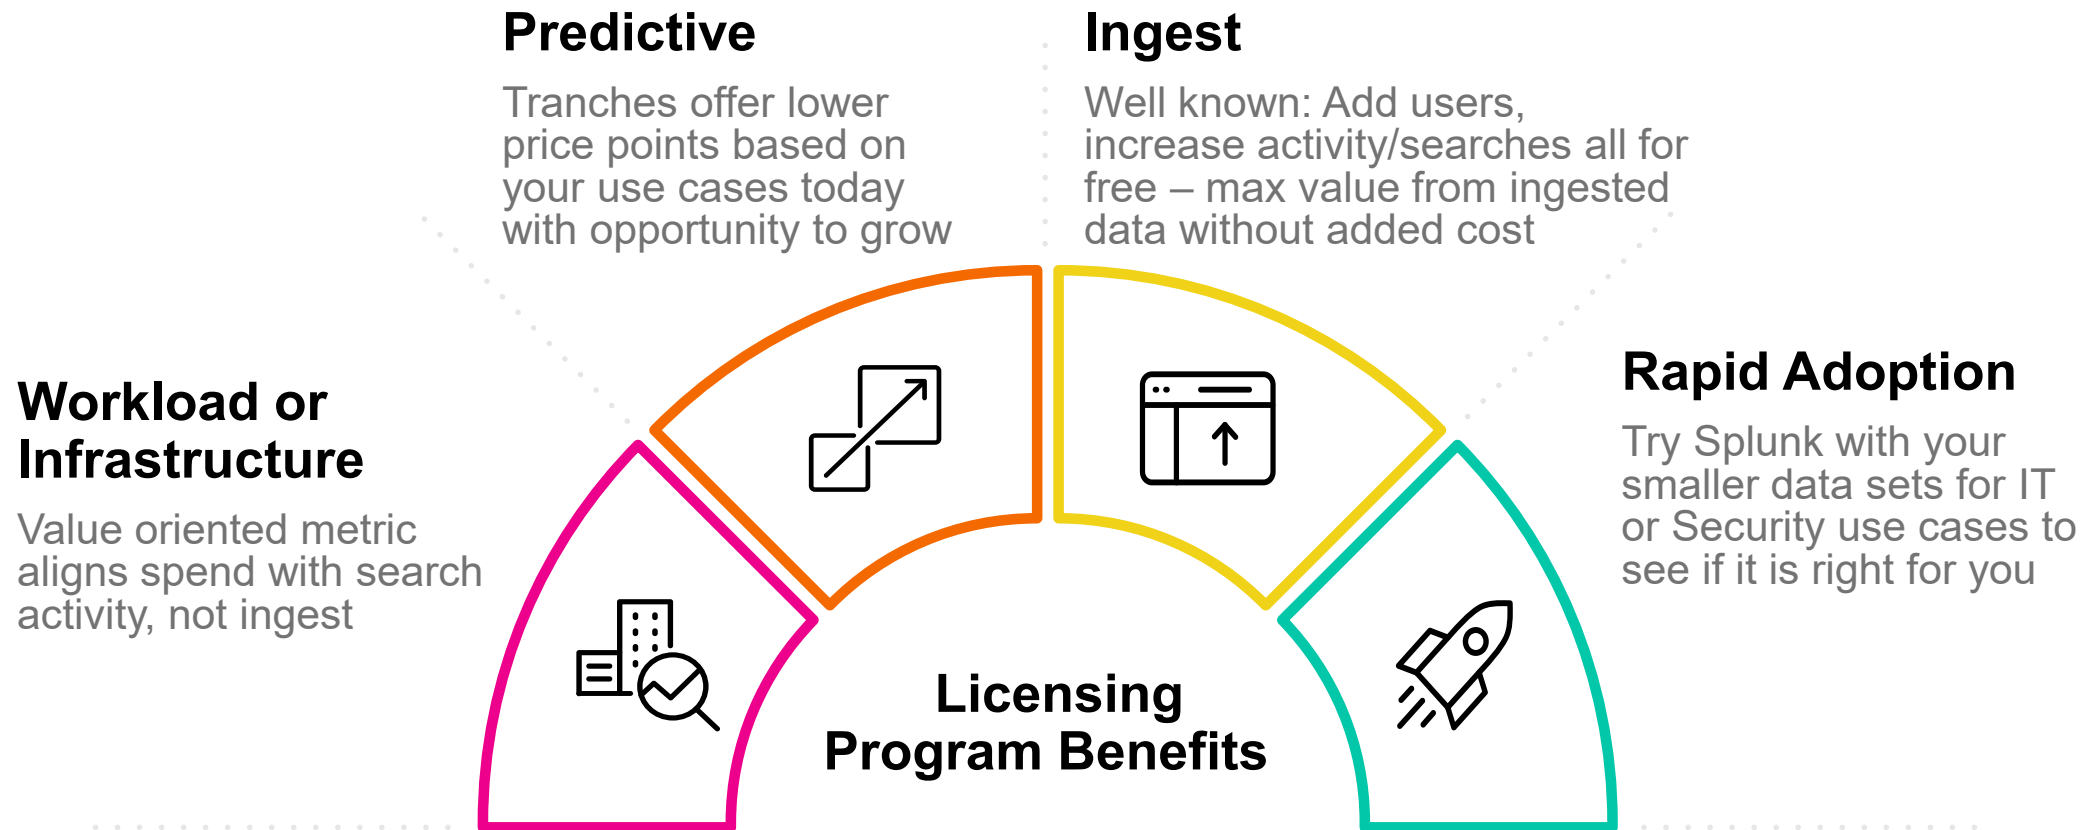


Background




Splunk's Flexible Licensing Options

Get the most out of Splunk depending on your data needs and capabilities




Workload or Infrastructure Licensing Model

Key things to know



Cloud: Workload Licensing
(SVC)




On Prem: Infrastructure Licensing
(vCPU)


Workload Licensing is an alternative for Splunk Cloud, Splunk Enterprise & Premium Solutions that is **based on the compute required to run Splunk**

Workload or Infrastructure Licensing Model

Key things to know



Cloud: Workload Licensing
(SVC)




On Prem: Infrastructure Licensing
(vCPU)

Workload Licensing is an alternative for Splunk Cloud, Splunk Enterprise & Premium Solutions that is **based on the compute required to run Splunk**

Available to new and current **Cloud or On Prem customers** – measured in SVCs for Cloud and vCPUs for On Prem


Workload or Infrastructure Licensing Model

Key things to know



Cloud: Workload Licensing
(SVC)

Utilization



On Prem: Infrastructure Licensing
(vCPU)

Allocation

For you if...

For you if...

**You do not want ingest limits to force
selection between important data**

For you if...

You do not want ingest limits to force selection between important data

You want a cost model tied to value (search, analytics and insights)

For you if...

You do not want ingest limits to force selection between important data

You want a cost model tied to value (search, analytics and insights)

You want the flexibility to adjust workloads based on compute capacity

For you if...

You do not want ingest limits to force selection between important data

You want a cost model tied to value (search, analytics and insights)

You want the flexibility to adjust workloads based on compute capacity

You want to index all or most of your datasets and source types in Splunk

For you if...

You do not want ingest limits to force selection between important data

You want a cost model tied to value (search, analytics and insights)

You want the flexibility to adjust workloads based on compute capacity

You want to index all or most of your datasets and source types in Splunk

You have datasets outside of Splunk that can help build additional use cases

Metric Drivers

Metrics
Deliver:

This is possible because...

Metric Drivers

Metrics
Deliver:

Customer control



This is possible because...

Splunk CPU utilization is primarily driven by search load and not ingest amount

Metric Drivers

Metrics
Deliver:

Customer control

Value orientation



This is possible because...

Splunk CPU utilization is primarily driven by search load and not ingest amount

Value is *typically* achieved at search time, not index time

Metric Drivers

Metrics
Deliver:

Customer control

Value orientation

**Logical way to access high
and low priority data**



Splunk CPU utilization is primarily driven by search load and not ingest amount



Value is *typically* achieved at search time, not index time



Many customers have a lot more barely used data than frequently used data

Metric Drivers

Metrics Deliver:

Customer control




Splunk CPU utilization is primarily driven by search load and not ingest amount

Value orientation



Value is *typically* achieved at search time, not index time

Logical way to access high and low priority data



Many customers have a lot more barely used data than frequently used data

Require:

Dedicated passion for Splunk administration



Bring in as much data as you want and pay based on how much compute your searches require.

Splunk Licenses Logical Cores (Not Usage)

On Prem only, not Cloud

Your Configuration

What Splunk Counts

Splunk Licenses Logical Cores (Not Usage)

On Prem only, not Cloud

Your Configuration

Bare Metal Server with no
hyperthreading

(no. of logical cores = no. of physical cores)



What Splunk Counts

Physical CPU cores

Splunk Licenses Logical Cores (Not Usage)

On Prem only, not Cloud

Your Configuration

Bare Metal Server with no
hyperthreading

(no. of logical cores = no. of physical cores)



Physical CPU cores

Bare Metal Server with hyperthreading

(no. of logical cores = no. of virtual cores)



Virtual CPUs

Splunk Licenses Logical Cores (Not Usage)

On Prem only, not Cloud

Your Configuration

What Splunk Counts

Bare Metal Server with no
hyperthreading

(no. of logical cores = no. of physical cores)



Physical CPU cores

Bare Metal Server with hyperthreading

(no. of logical cores = no. of virtual cores)



Virtual CPUs

Public Cloud/Virtualization

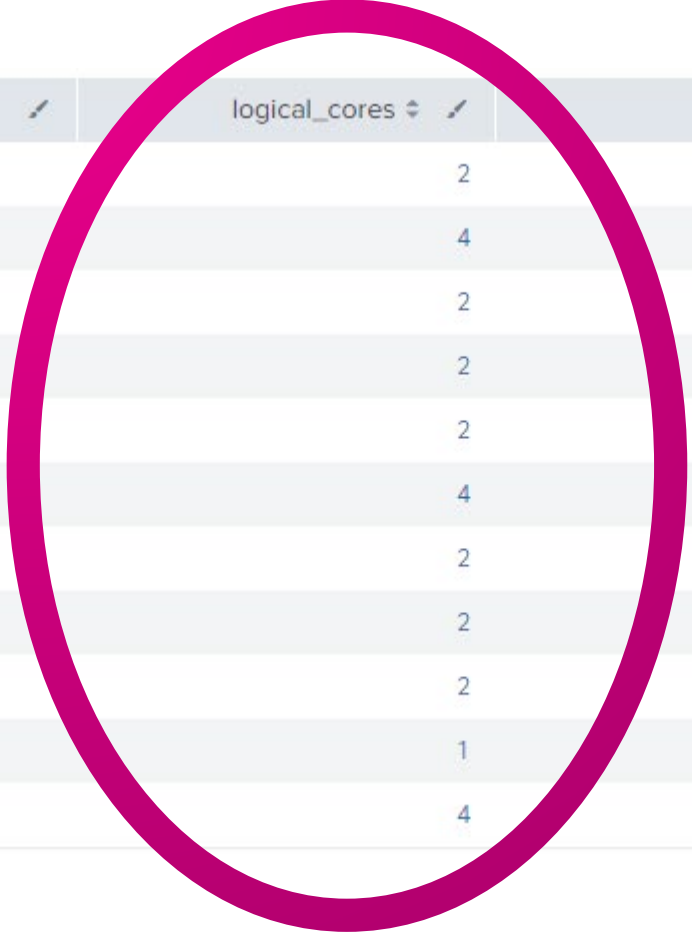


Virtual CPUs

In pseudo-SPL:

```
| rest splunk_server=* /services/server/sysinfo  
| eval logical_cores = coalesce(numberOfVirtualCores ,  
numberOfCores)  
| table splunk_server logical_cores numberOfVirtualCores  
numberOfCores  
| search splunk_server IN ("searchhead", "indexer" )
```

In pseudo-SPL:



| splunk_server ↕ | logical_cores ↕ | numberOfVirtualCores ↕ | numberOfCores ↕ |
|-----------------|-----------------|------------------------|-----------------|
| indexer | 2 | 2 | 1 |
| searchhead | 4 | 4 | 2 |
| indexer | 2 | 2 | 1 |
| indexer | 2 | 2 | 1 |
| indexer | 2 | 2 | 1 |
| indexer | 4 | 4 | 2 |
| indexer | 2 | 2 | 1 |
| searchhead | 2 | 2 | 1 |
| indexer | 2 | 2 | 1 |
| indexer | 1 | 1 | 1 |
| indexer | 4 | 4 | 2 |

Sizing Premium Solutions



Splunk IT Service
Intelligence™



Splunk Enterprise
Security™

Sizing Premium Solutions



Premium percentage =
premium data volume / total data volume

Sizing Premium Solutions



Premium percentage =
premium data volume / total data volume

Premium Logical Cores =
Premium search head cores +
(**premium percentage** * total logical cores)

Sizing Premium Solutions



Premium percentage =
premium data volume / total data volume

Premium Logical Cores =
Premium search head cores +
(**premium percentage** * total logical cores)

Example:

- Near max of about 1TB of total data
- 250GB is for Security
- Therefore, premium percentage is 25%

Sizing Premium Solutions



Premium percentage =
premium data volume / total data volume





Example:

- Near max of about 1TB of total data
- 250GB is for Security
- Therefore, premium percentage is 25%

Premium Logical Cores =
Premium search head cores +
(**premium percentage** * total logical cores)


- 125 total cores
- 1 ES host at 24 cores
- = 24 premium cores + (25% * 125 cores)
- = 57

Keep Everything You Love about Splunk






Keep Everything You Love about Splunk

**Compatible
across
portfolio**




Use for Enterprise,
ES, ITSI



Keep Everything You Love about Splunk

**Compatible
across
portfolio**





Use for Enterprise,
ES, ITSI

**How it's
licensed**




Based on entitlement:
cloud – utilization or
on prem – deployed
logical cores



Keep Everything You Love about Splunk

Compatible across portfolio



Use for Enterprise,
ES, ITSI

How it's licensed




Based on entitlement:
cloud – utilization or
on prem – deployed
logical cores

Account for...



Prod, non-prod,
DR/failover....



Keep Everything You Love about Splunk

Compatible across portfolio



Use for Enterprise,
ES, ITSI

How it's licensed




Based on entitlement:
cloud – utilization or
on prem – deployed
logical cores

Account for...



Prod, non-prod,
DR/failover....

Entitlement Hardware




Only SH and IDX
count towards
entitlement



Licensing



Hey Jeff, How Do I Size?



License Files: Don't Trip

Ingest license reporting is wrong and messages can be ignored

Splunk Enterprise - vCPU License - Term stack [Learn more](#)

| Licenses | Volume | Expiration | Status | |
|---|------------|-------------------------|-----------------|--------|
| Splunk Enterprise - vCPU License - Term | 512,000 MB | Jun 1, 2021, 1:59:59 AM | valid | Delete |
| Splunk Enterprise - vCPU License - Term | 512,000 MB | Jun 1, 2022, 1:59:59 AM | FROM_THE_FUTURE | Delete |
| Splunk Enterprise - vCPU License - Term | 512,000 MB | Jun 1, 2023, 1:59:59 AM | FROM_THE_FUTURE | Delete |
| Effective daily volume | | 512,000 MB | | |

```
<?xml version="1.0" encoding="UTF-8"?>
<license>
  <signature>
    bIann7lCcSG305oJtirAls97v9sa3mCC75GcCiR5en4R7PIz3tIZxKF6+rDCDMHrZnOI
    1m+VOnm4ys10Z1EWKoS10YfY9iB0URQZfxWH15XU17q8fi3cPxavAQI/zj27y+saqqQ
    +21VNIIm0MHK7CBRGDyrexFQDFisNcZxfxHiegFLPyp4KONPiaNlityxO/SUV4f1nQoE
  </signature>
  <payload>
    <type>enterprise</type>
    <group_id>Enterprise</group_id>
    <quota>429496729600</quota>
    <max_violations>5</max_violations>
    <window_period>30</window_period>
    <creation_time>1597561199</creation_time>
    <label>Splunk Enterprise Term Non-Production vCPU License</label>
    <expiration_time>1597561199</expiration_time>
    <subgroup_id>DevTest</subgroup_id>
    <features>
      <feature>Auth</feature>
      <feature>FwdData</feature>
      <feature>RcvData</feature>
      <feature>LocalSearch</feature>
      <feature>DistSearch</feature>
      <feature>RcvSearch</feature>
      <feature>ScheduledSearch</feature>
      <feature>Alerting</feature>
      <feature>DeployClient</feature>
      <feature>DeployServer</feature>
      <feature>SplunkWeb</feature>
      <feature>SigningProcessor</feature>
      <feature>SyslogOutputProcessor</feature>
      <feature>CanBeRemoteMaster</feature>
      <feature>SubgroupId</feature>
    </features>
    <optional_features>
      <feature>DisableQuotaEnforcement</feature>
    </optional_features>
    <sourcetypes/>
    <guid>8949BC12-DEC0-4A83-88DC-C53D6904B42F</guid>
  </payload>
</license>
```

License Combinations



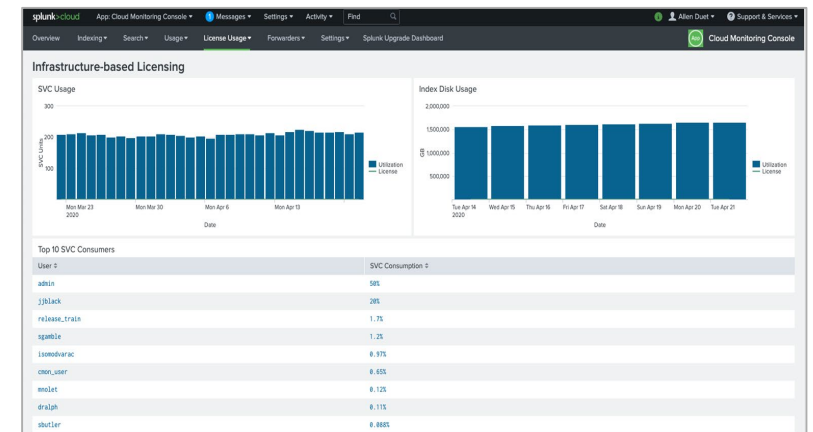
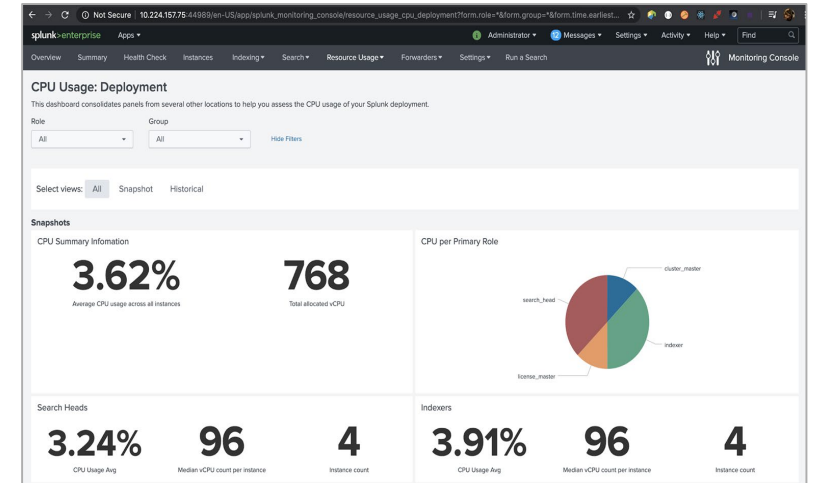
Monitor Usage in Cloud or On-Prem

Use the Cloud Monitoring Console (Cloud) or Monitoring Console (On-Prem)

Benchmark and identify peaks in SVC/vCPU use

Identify top consumers and consuming indexes

Monitor distribution across search types, apps, and workload pools to triage & adjust



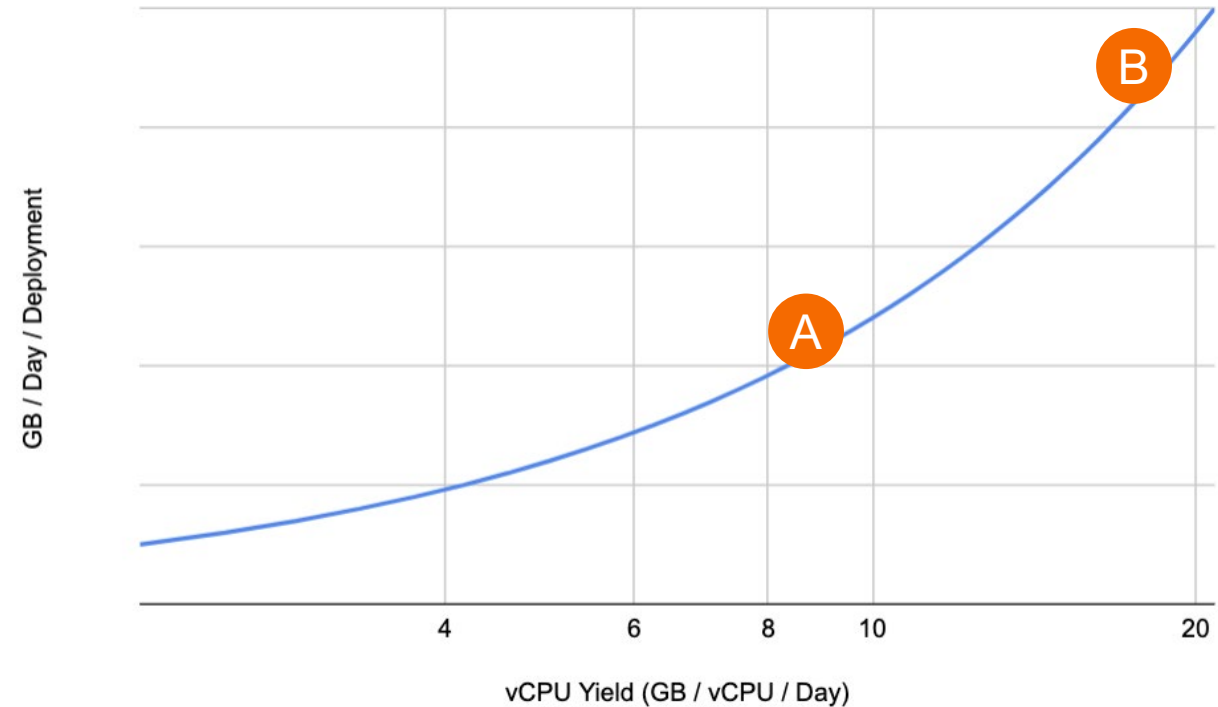
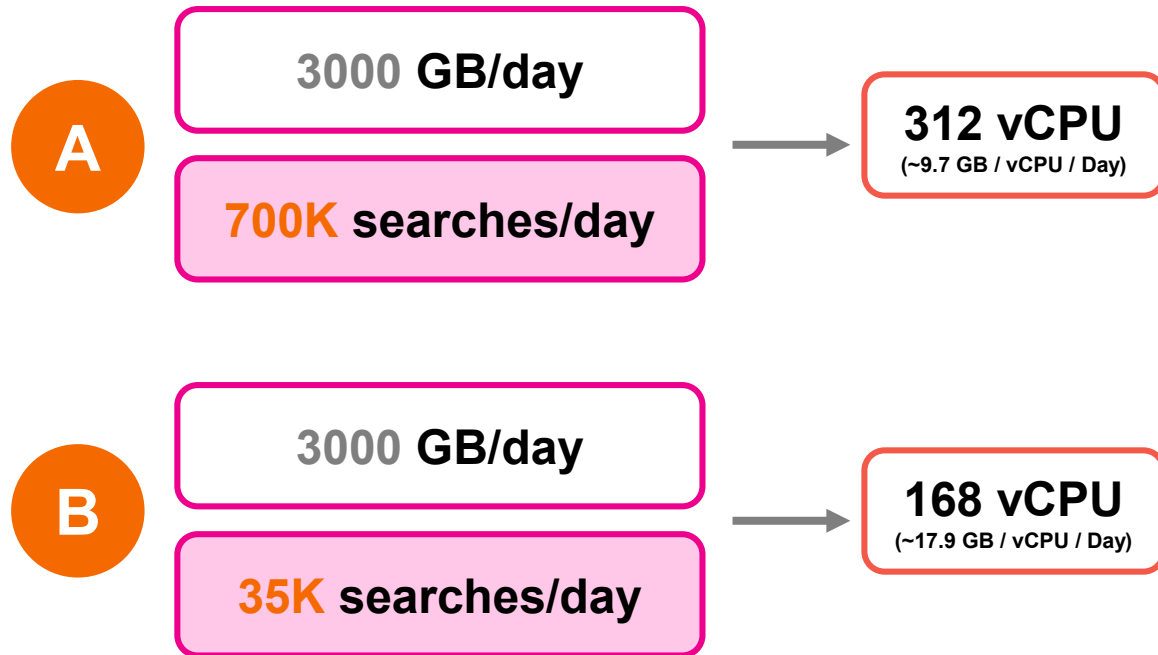


Best Practices for Infrastructure and Workload Licensing



Workload Drives Compute

The importance of vCPU Yield as a metric



For example, keeping all other factors the same, when search is not consuming CPU, more data can be ingested by the same # of vCPUs.

Oh No, is That a Dent?!




Focus on the Big Picture



Slow Performance Trap




Slow Performance Trap




Add Indexers!



Add Indexers!




Add Indexers!



Themes for the Best Practices



Themes for the Best Practices



Themes for the Best Practices

Scope



Themes for the Best Practices

Scope



Not Scope

Themes for the Best Practices



Scope



Not Scope

Themes for the Best Practices



Scope



Not Scope

Themes for the Best Practices



Scope



Not Scope

Themes for the Best Practices



Scope



Not Scope



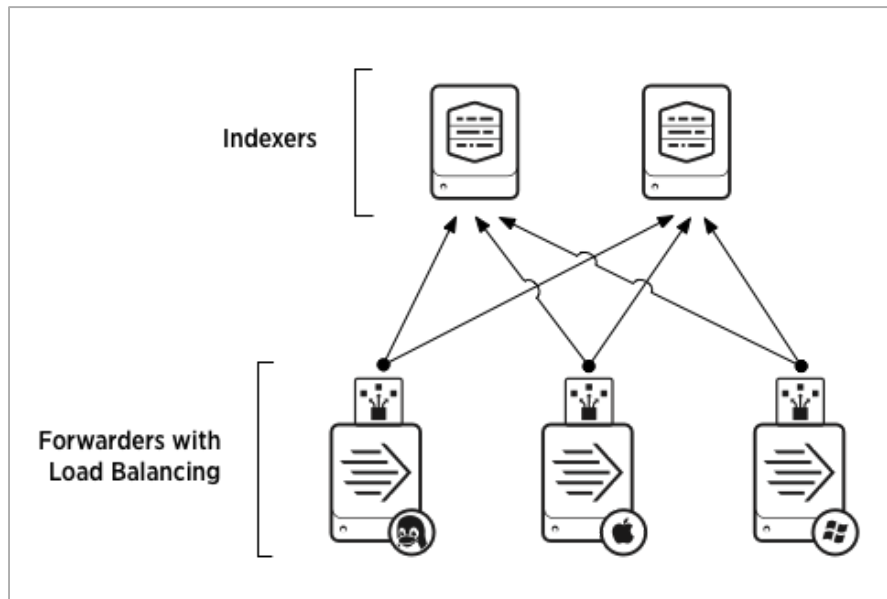
Best Practices

1. Equal Data Distribution
2. Efficient Data Onboarding
3. Restrictions on User Search
4. Well Defined Searches
5. Increasing Parallelization
6. Workload Management (WLM)

Equal Data Distribution

Common when sending from HF to Cloud; Parallelization; indexer data rebalance

“Configure load balancing for Splunk Enterprise”



“Indexing: Indexes and Volumes”

Snapshots

Index Structure Overview

3

Indexers

20.19 GB

Total Index Size

41.96 GB

Total Raw Data Size (Uncompressed)

2.08:1

Raw to Index Size Ratio*

* This is the ratio of uncompressed raw data size to index size.

Events Overview

111,321,470

Total Events

2020-04-08 19:15:33+0000

Earliest Event

2020-08-26 17:01:32+0000

Latest Event

Indexer Averages

6.73 GB

Average Index Size

37,107,156

Average Event Count

140 days

Median Data Age

144

Average Bucket Count

Instances (3)

| Indexer ⬇ | Data Age vs Frozen Age (days) ⬇ | Index Usage (GB) ⬇ | Home Path Usage (GB) ⬇ | Cold Path Usage (GB) ⬇ |
|--------------|-----------------------------------|--|---|--------------------------------------|
| yavin | <div><div></div></div> 140 / 3652 | <div><div></div></div> 5.24 / 10240.00 | <div><div></div></div> 5.24 / unlimited | <div><div></div></div> 0 / unlimited |
| hoth | <div><div></div></div> 140 / 3652 | <div><div></div></div> 5.81 / 10240.00 | <div><div></div></div> 5.81 / unlimited | <div><div></div></div> 0 / unlimited |
| endor | <div><div></div></div> 140 / 3652 | <div><div></div></div> 9.14 / 10240.00 | <div><div></div></div> 9.14 / unlimited | <div><div></div></div> 0 / unlimited |

Equal Data Distribution

“Rebalance the indexer cluster”

Only for...

- On Prem
- Indexer Clusters
- People that like rules of 3 ;)

The screenshot shows the Splunk Enterprise documentation page for 'Managing Indexers and Clusters of Indexers'. The page title is 'Rebalance the indexer cluster'. It includes a 'Download manual as PDF' button and a 'Download topic as PDF' button. The main content area is titled 'Rebalance the indexer cluster' and contains the following text:

By rebalancing the indexer cluster, you balance the distribution of bucket copies across the set of peer nodes. A balanced set of bucket copies optimizes each peer's search load and, in the case of data rebalancing, each peer's disk storage.

Types of indexer cluster rebalancing

There are two types of indexer cluster rebalancing:

- Primary rebalancing
- Data rebalancing

Primary rebalancing

The goal of primary rebalancing is to balance the search load across the peer nodes.

Primary rebalancing redistributes the primary bucket copies across the set of peer nodes. It attempts, to the degree possible, to ensure that each peer has approximately the same number of primary copies.

Primary rebalancing simply reassigns primary markers across the set of existing searchable copies. It does not move searchable copies to different peer nodes. Because of this limitation, primary rebalancing is unlikely to achieve a perfect balance of primaries.

Because primary rebalancing only reassigns markers and does not cause any bucket copies to move between peers, it completes quickly.


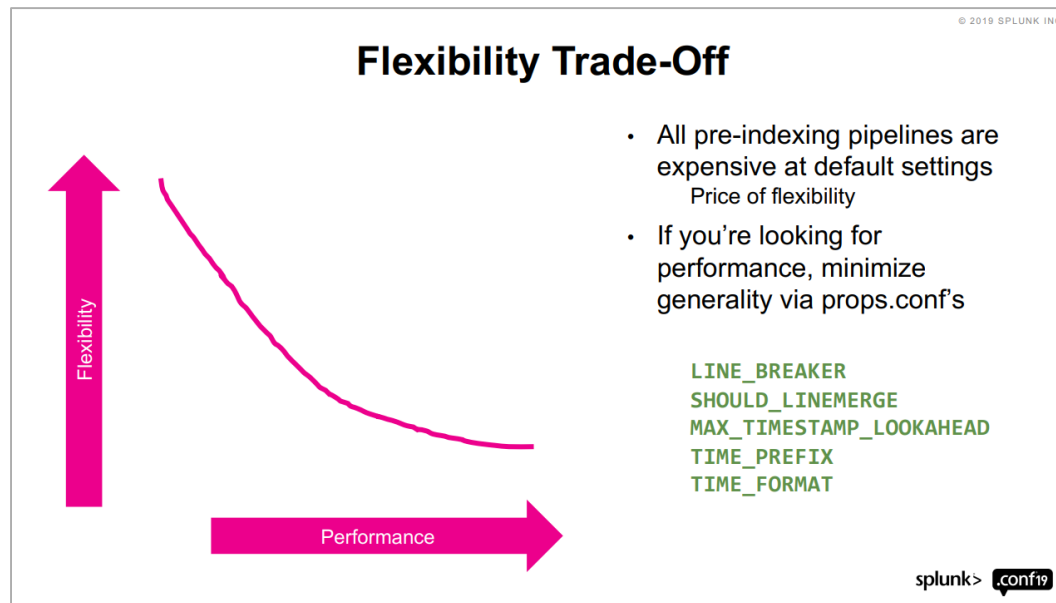
See [Rebalance indexer cluster primary bucket copies](#).

Data rebalancing

The goal of data rebalancing is to balance the storage distribution across the peer nodes.

Efficient Data Onboarding

Learn more at .conf online: “PLA1486C – Understanding Splunk Performance and Making Hardware (Virtual/Physical) Choices”



Restrictions on User Search

Search for “Create and manage roles with Splunk Web” on SplunkDocs

Role search job limit
Set a limit for how many search jobs that all users with this role can run at the same time. ?

Standard search limit

Real-time search limit

User search job limit
Set a limit for how many search jobs that a single user with this role can run at the same time. ?

Standard search limit

Real-time search limit

Role search time window limit
Select a time window for searches for this role. Inherited roles can override this setting.

Disk space limit
Set the maximum amount of disk space, in megabytes, that search jobs for a specific user with this role can use.

Standard search limit MB

Searches consume compute
-> protect against new users

Learn about incentives @:

.conf online “FN1054 - Best Practices and Better Practices for Admins”

Lantern: “Enabling users with incentives”

Well Defined Searches

Monitoring Console's "Search Usage Statistics: Deployment" dashboard

Splunk® Enterprise

Search Manual

Download manual as PDF

Hide Contents

Search Manual

Search Overview

Using the Search App

Search Primer

Optimizing Searches

About search optimization

Quick tips for optimization

Write better searches

Built-in optimization

Search normalization

Documentation / Splunk® Enterprise / Search Manual / About search optimization

Download topic as PDF

About search optimization

Search optimization is a technique for making your search run as efficiently as possible.

When not optimized, a search often runs longer, retrieves larger amounts of data from the indexes than i

Multiply these issues by hundreds or thousands of searches and the end result is a slow or sluggish syst

There are a set of basic principles that you can follow to optimize your searches.

- Retrieve only the required data
- Move as little data as possible
- Parallelize as much work as possible
- Set appropriate time windows

To implement the search optimization principles, use the following techniques.

| Long-running Searches | | | | |
|-----------------------------|------------------|---------------------------|--------------------------|--------------------------|
| Report Name/Search String ↕ | Search Runtime ↕ | Search Start ↕ | Earliest Time ↕ | Latest Time ↕ |
| search13 | 5.67s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search15 | 5.64s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search12 | 5.59s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search16 | 5.59s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search14 | 5.55s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search6 | 5.54s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search11 | 5.49s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| bundle_rep_log_base | 5.47s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search10 | 5.04s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |
| search5 | 5.01s | 08/20/2020 17:52:10 -0400 | Thu Aug 20 17:52:00 2020 | Thu Aug 20 17:52:00 2020 |

Increasing Parallelization

Search “Parallelization settings” in the SplunkDocs



Increasing Parallelization

Search “Parallelization settings” in the SplunkDocs



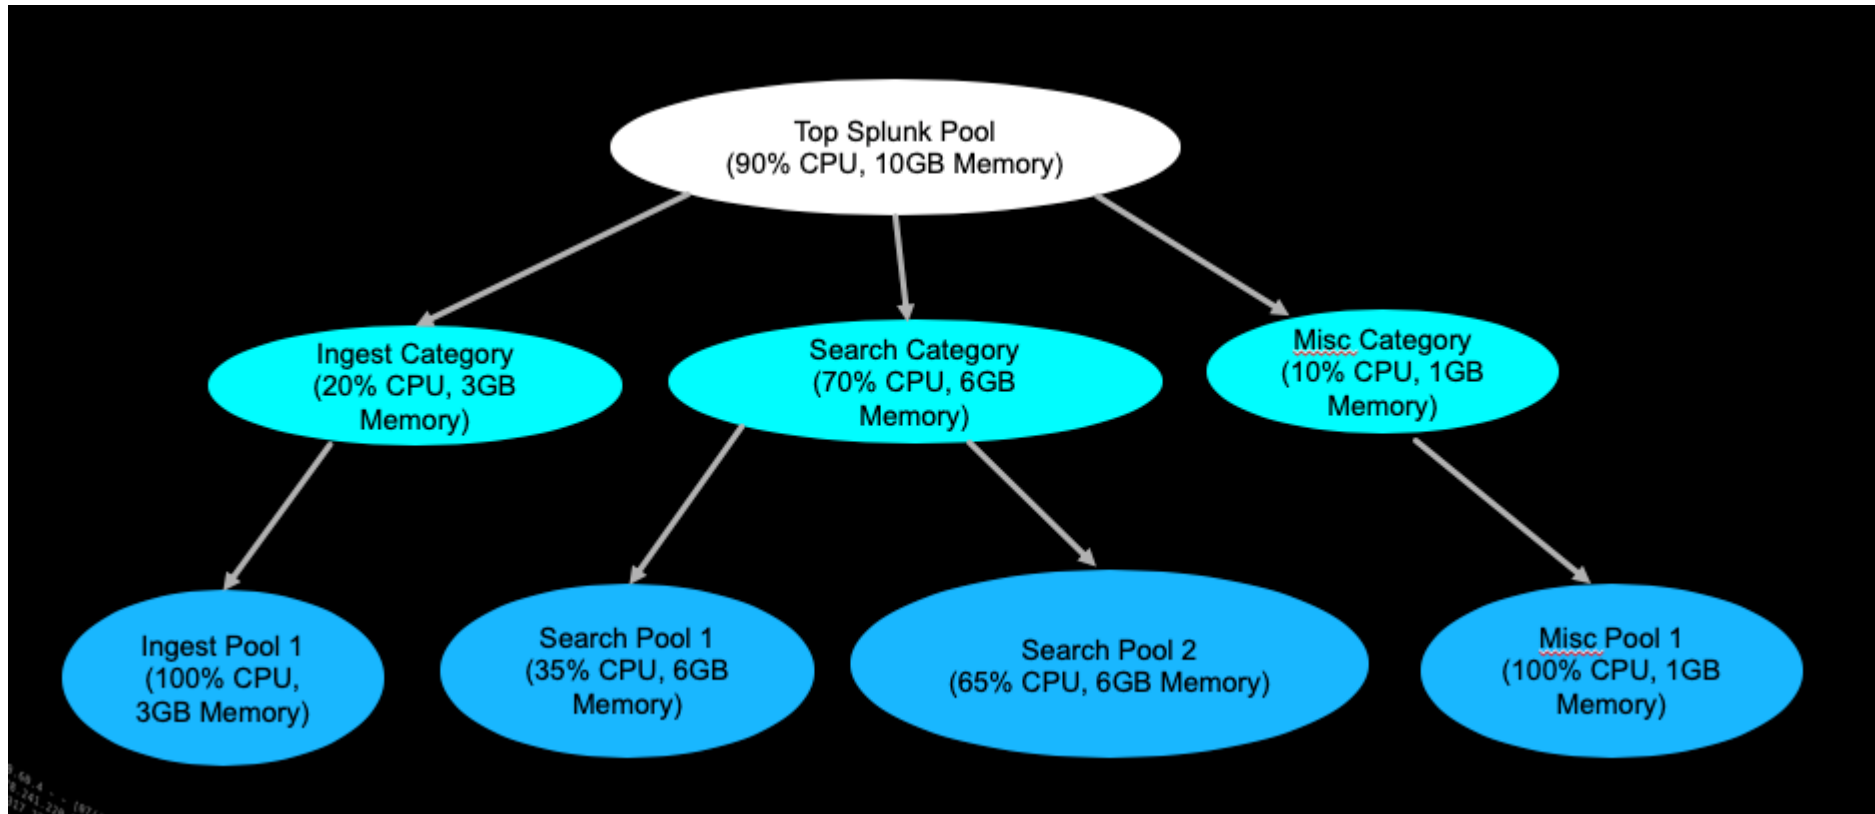
Cloud Managed. But on prem HFs.

Could drive up utilization, which is good for on prem

Could be impacting if insufficient resources available

Workload Management (WLM)

Search “About workload management” on SplunkDocs



Best Practices

1. Equal Data Distribution
2. Efficient Data Onboarding
3. Restrictions on User Search
4. Well Defined Searches
5. Increasing Parallelization
6. Workload Management (WLM)



Closing



Legal Blah

Public legal definitions of licensing can be found at *Splunk Offerings Purchase Capacity and Limitations* on splunk.com: https://www.splunk.com/en_us/legal/licensed-capacity.html

Infrastructure based licensing is described publicly on the *Data-to-Everything Pricing FAQ* page on splunk.com: https://www.splunk.com/en_us/software/pricing/faqs/data-to-everything.html#infrastructure-pricing

Success plans are outlined on the support and services page on splunk.com: https://www.splunk.com/en_us/support-and-services/vcpu-plans.html

Call to Action

1. Collaborate: #licensing
 - Sign Up @ <http://splk.it/slack>
2. Monitoring Console
3. PLA1486C – Understanding Splunk Performance and making hardware (virtual/physical) choices
4. PLA1826C – Taming wild resources. Tips to manage and remediate busy Splunk instances.
5. PLA1582C – Cloud Monitoring Console Tips and Tricks for the Splunk Cloud admin
6. On Demands Services like “Ask an Expert”
 - <https://www.splunk.com/pdfs/legal/splunk-on-demand-services-catalog.pdf>



Thank You

Please provide feedback via the
SESSION SURVEY

