Taming Wild Resources

Tips to Manage Busy Splunk Resources

Aaron Kornhauser – ak@splunk.com
Vladimir Skoryk – v@splunk.com

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.

Who Are We?

Aaron Kornhauser

- Staff Sales Engineer
- 7 years at Splunk in various roles:
  - PS Consultant
  - PS Manager
  - Sales Engineer
- Focus on Cloud, GDI, and large-scale Splunk Architectures, and building cool stuff
- Pittsburgh, PA

Vladimir Skoryk

- Principal PS Architect
- 7 years in PS
- Loves photography
- Lives at MIA

Agenda

Refactor or Scale

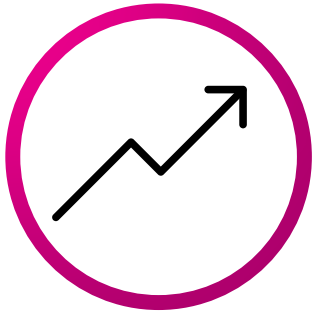

Search / Indexing Inspection

Configs to Tune

Workload Management

Cloud Intro

.conf20

splunk>

Refactor or Scale

Monitoring Console Basics

- Introduced in late 6.x to provide visibility into the health and perf of your environment
- Provides various dashboards and alerts for:
  - Indexing / Clustering / Licensing / Search
  - Resource utilization
  - Forwarder inventory and missing forwarders
  - Missing or unavailable instances
- Standardized Health Checks
  - 30 checks OOTB
  - Additional PS checks available here: https://splunkbase.splunk.com/app/4527/

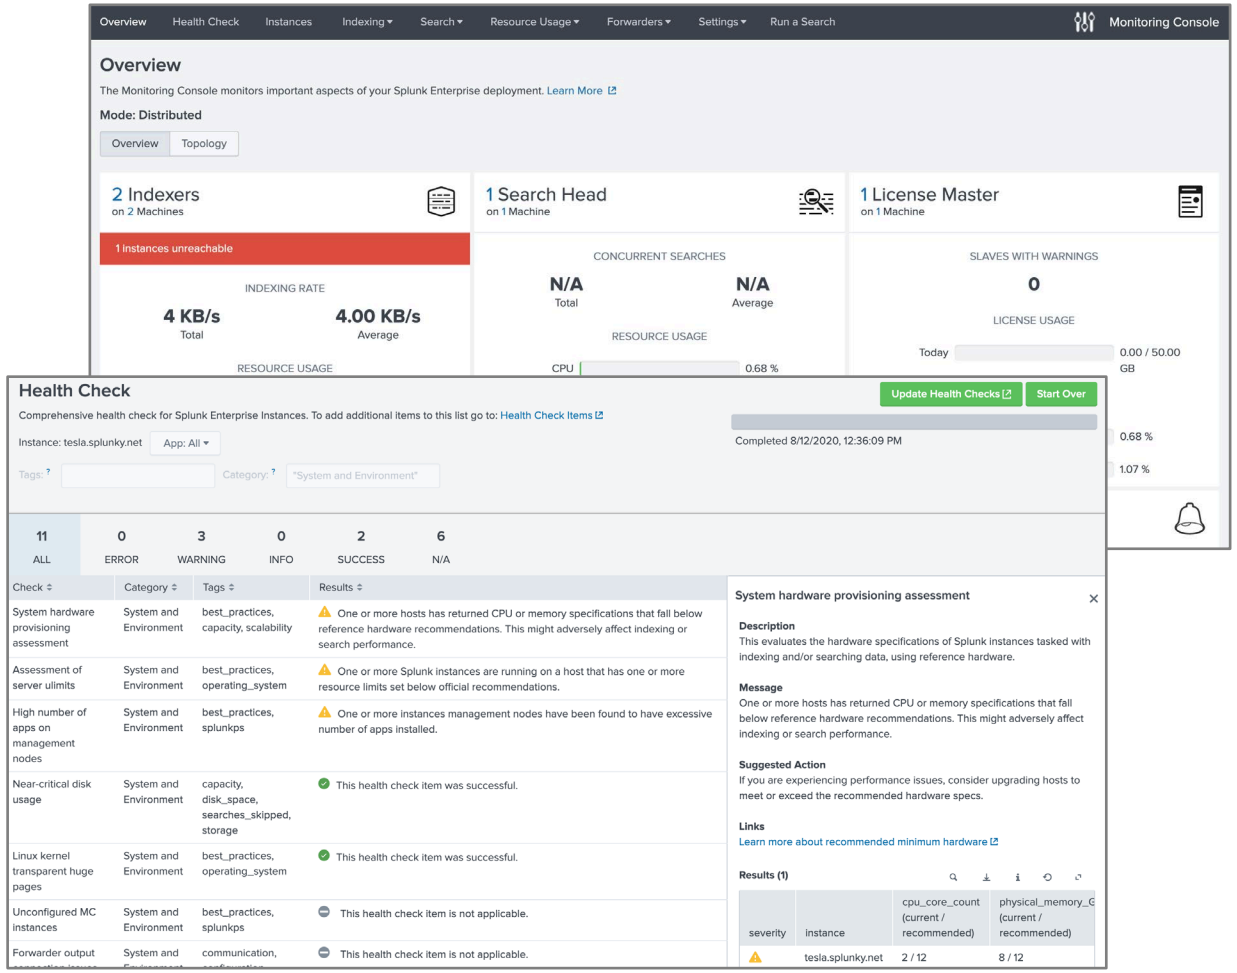

Reference Hardware Specs

Distributed Deployments

Instance	Core Splunk*	Enterprise Security†
Indexer	12 CPU cores 12GB RAM 800+ IOPS/Indexer RAID 1+0 **Target Ingest: 200-350GB/day**	16 CPU cores 32GB RAM 800+ IOPS/Indexer RAID 1+0 **Target Ingest: 100-150GB/day**
Search Head	16 CPU cores 12GB RAM 2 x 300GB, 10k RPM SAS HD in RAID1	16 CPU cores 32GB RAM 2 x 300GB, 10k RPM SAS HD in RAID1

Note: Core counts equate to the physical core count, not the vCPU count with hyperthreading

* http://docs.splunk.com/Documentation/Splunk/latest/Capacity/Referencehardware
† http://docs.splunk.com/Documentation/ES/latest/Install/DeploymentPlanning

When to Consider Scaling

- Adoption of new use cases, increase of users, and premium applications (ES, ITSI, UBA)
- Search concurrency limits are being hit
  - Can be vertically scaled or through configurations if server resources are stable
- Data retention and available storage will not meet SLA
- System load
  - System load > # of cores
- Deploying SHC/IDX clustering
  - Primarily a horizontal expansion
- Increased indexing volume has surpassed past recommended values
  - Searches begin to skip, queues begin to fill, overall performance will appear slow
- Increased number of forwarders
  - If using intermediary forwarders, want a 2:1 ratio of IF pipelines to indexers to prevent starvation

When to Consider Refactoring

- Scale has surpassed data center capacity or segmentation challenges with HA/DR
  - Review architecture to determine efficiencies or simplifications can be implemented
- Customer environment shifting to Cloud or retiring hardware
  - Opportunity to review legacy environment for shift into BYOL or Splunk Cloud
  - Introduces opportunities to review and refactor how data is being collected
- Hard-to-manage collection tier and inefficient data routing with many “hops”
  - Complex app and configuration tiering across the deployment for data onboarding
  - Review forwarding and data routing architectures for efficiency gains, especially removal of heavy forwarders
  - Consider direct push based approach with DSP or HEC

.conf20

splunk>

Search / Index Inspection

Search Introspection

- Additional capacity can be achieved by scaling and/or housekeeping
  - Look for search patterns, heavy periodicity (scheduling peaks, large lookback windows)
  - Scheduling overlap (search needs to run every 5 minutes, but takes 8 minutes to complete)
  - Skipping searches (system or user limits)
  - Look for searches with no results
  - Look for search anti-patterns (greedy, not specifying *index*, *sourcetype*, *source*, using excessive wildcards, etc)
  - Look for same searches running multiple places (multiple search heads)
  - Look for searches/data models/report accelerations running on indexers (**tip**: disable scheduler)
  - Look for searches that are similar to each other (*index=foo host=A error* … *index=foo host=B error*)
- Establish process for periodic review of content

Data Indexing

- Choices made during ingestion of data can have significant performance implication at search time
  - Ensure that *host*, *source*, *sourcetype* fields are set correctly
  - Depending on searching patterns, introducing additional indexed fields might be required
- Avoid mixing distinct data sources within the same *index* (or worse *sourcetype*, *source*)
  - Want to avoid pulling a bucket of data, only to use 1% of the data within it
- Initial data onboarding scope should include use-cases and uses for the data
- Establish periodic review of search patterns, map back to data indexing strategy

Data Indexing cont.

- Indexer data imbalance can have a significant impact on search performance
- Strive for equal data balance across all indexers across a 5-minute time interval, or better
- Watch out for
  1. Data that only makes it to select indexers
  2. Data that is prominent across subset of indexers
  3. Data that across all indexers, but favors subset
  4. No data being sent to indexers (outdated configuration, performance issues, firewalls)

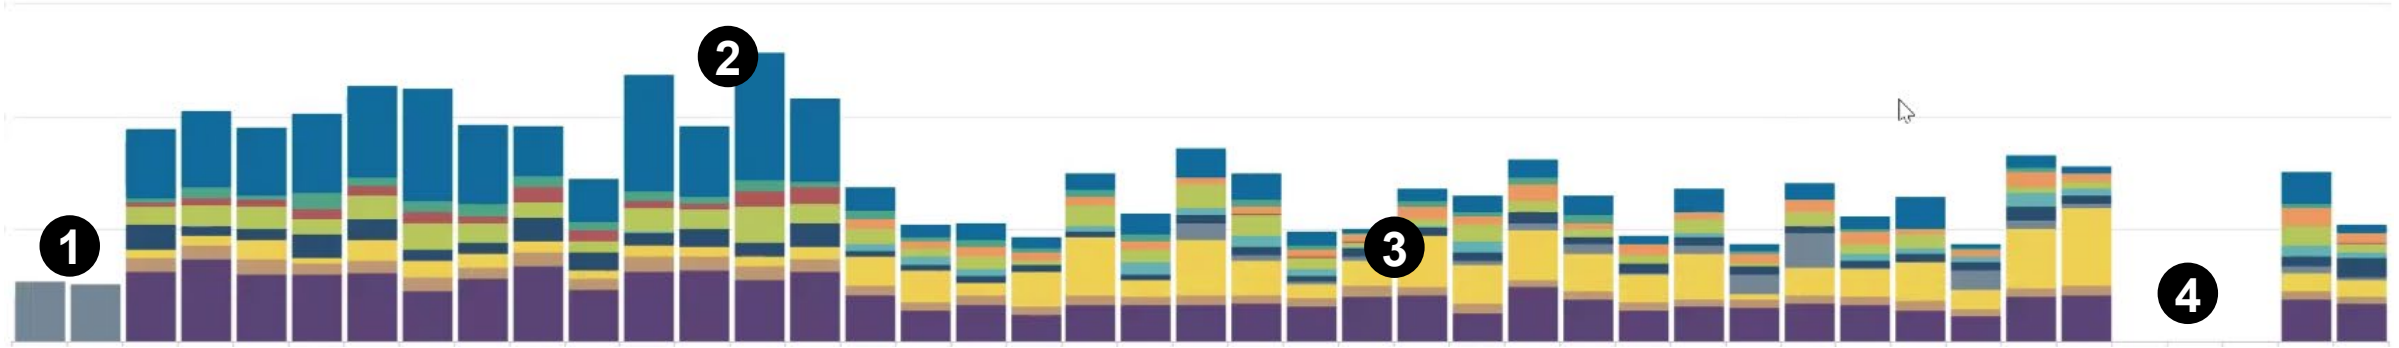

Configs to Tune

What if my server is under-utilized?

Settings to Stretch Resources

- Remember cause and effect of a distributed environment
  - If search tier is under utilized, does not mean indexing tier is
  - Review Splunk architecture, does it still meet its design goals?
  - Adjust advanced settings with caution; involve your friendly Splunk Team

- Additional search concurrency can be achieved by adjusting scheduler settings in ***limits.conf***
  - base_max_searches – constant # of searches (default 6)
  - max_searches_per_cpu - # of searches per CPU (default 1)
    concurrent_search_slots = max_searches_per_cpu x **number_of_cpus** + base_max_searches
  - max_searches_perc - % of total concurrent search slots for scheduled searches
  - auto_summary_perc - % of total concurrent search slots for RA/DM

- Data ingest throughput can be lifted by parallelizing ingestion pipelines, settings in ***server.conf***
  - parallelIngestionPipelines – number of parallel pipelines (default 1)

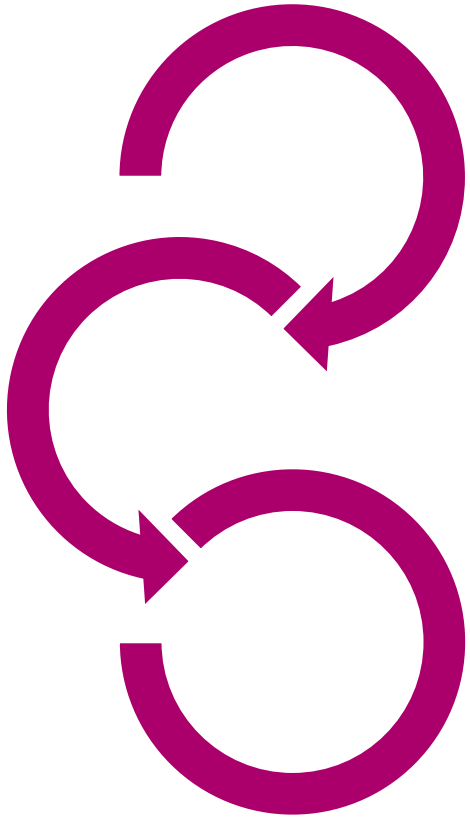

.conf20

splunk>

Workload Management

Workload Management

Allocating resources to specific tasks and users

- Introduced in 7.2 to tame inefficient users, content, and bad practices
- Allocate resources based on business priorities using the enhanced rules framework now including type of search (ad hoc, scheduled, all-time etc.)
- Allows for:
  - Prioritization of resources to critical search workloads, such as security investigations or alerts
  - Restriction of resources to prevent data ingest delay due to heavy search
  - Reduction of resource allocation to low-priority or inefficient content i.e. real-time searches
- Only available on specific *nix distros – no Windows support

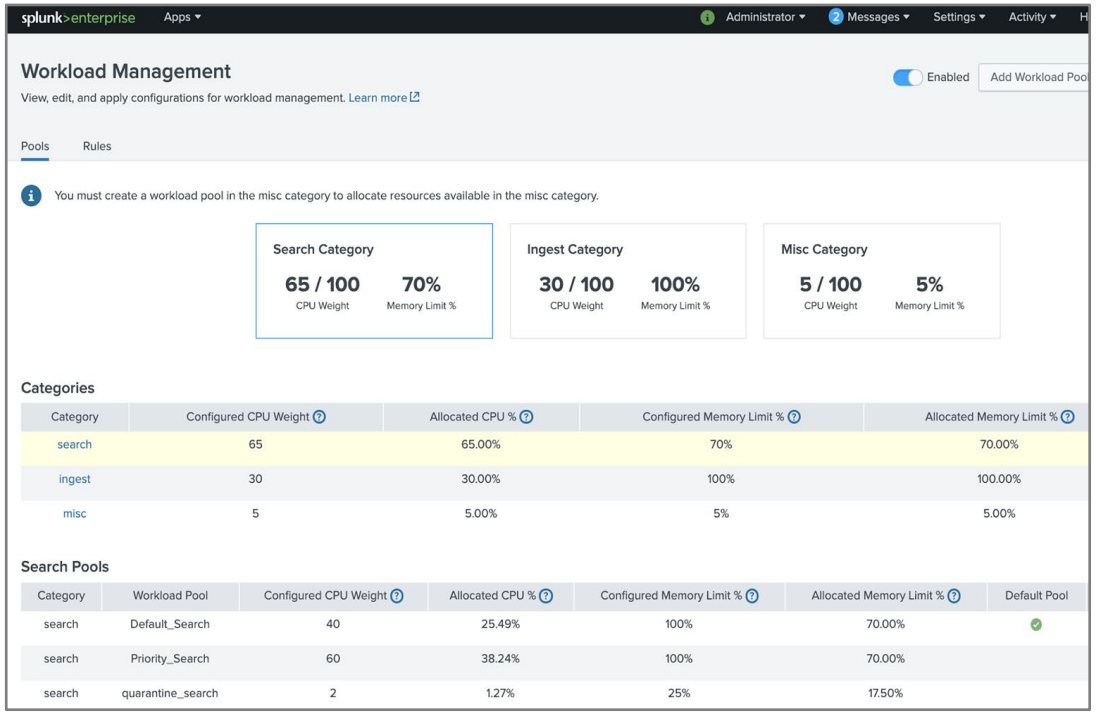

https://docs.splunk.com/Documentation/Splunk/latest/Workloads/Keyconcepts

Using Workload Management

Settings -> Workload Management

splunk>enterprise Apps ▾ | Administrator ▾ | Messages ▾ | Settings ▾ | Activity ▾ | Help ▾ | Find

Workload Management

View, edit and apply configurations of workload management. Learn m

Enabled | Add Workload Pool | Add Workload Rule | Apply Changes

Define various pools and resource allocation

Workload Pool	CPU (%)	Memory (%)	Default Search Pool	Default Ingest Pool	Actions
pool_1	35	35			Edit Delete
pool_2	20	20			Edit Delete
pool_4	15	15	✓		Edit Delete
pool_5	30	30		✓	Edit Delete

Filter workload pools to various predicates

Order	Workload Rule	Predicate	Workload Pool	Actions
1	rule_1	app=search	pool_1	Edit Delete
2	rule_2	role=admin	pool_2	Edit Delete
3	rule_3	role=analyst	pool_2	Edit Delete
4	rule_5	app=splunk_instrumentation	pool_4	Edit Delete

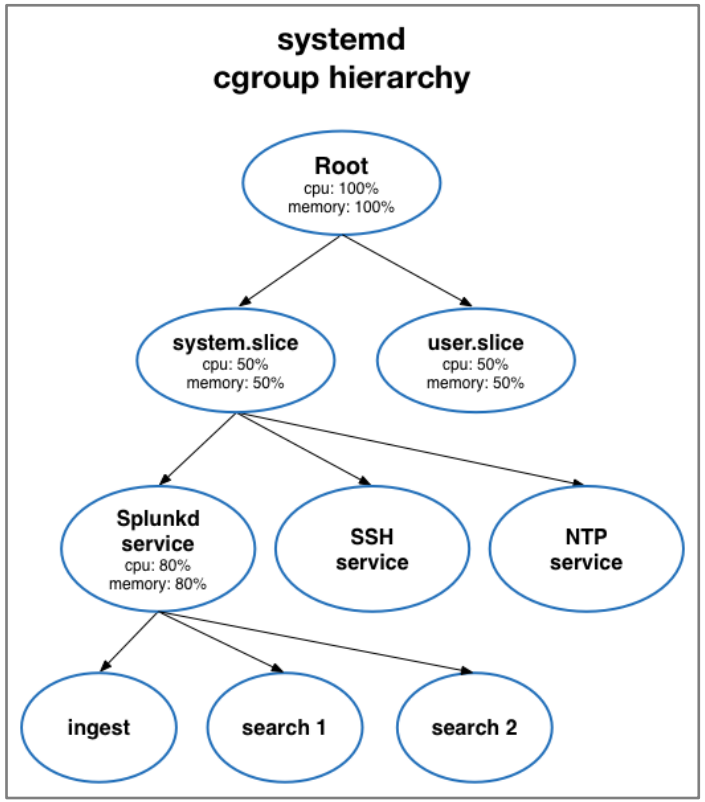

splunk> .conf20

Using Workload Management Continued

Using / Auditing WLM

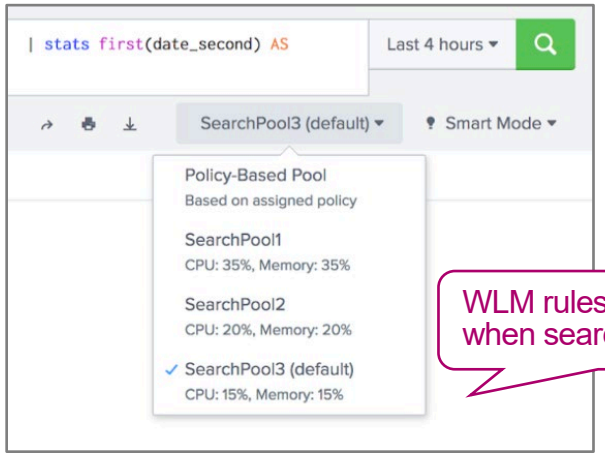

WLM rules and default when searching (adhoc)

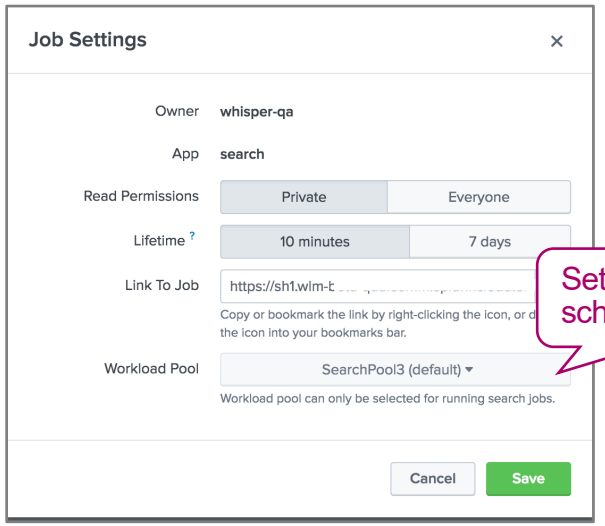

Setting pool for scheduled jobs

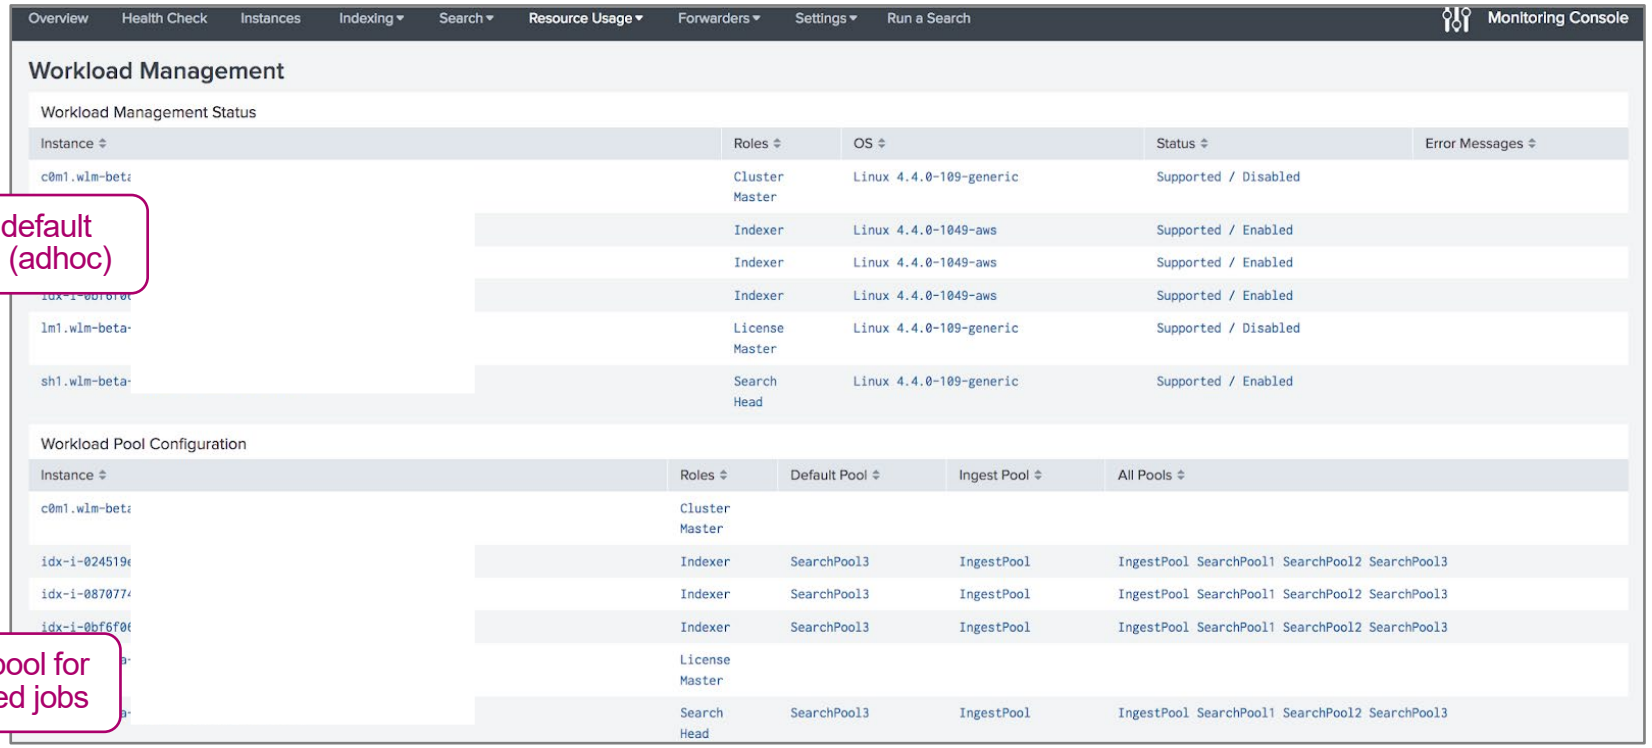

splunk> .conf20

.conf20

splunk>

Tired of managing on-prem Splunk?

Splunk Cloud

The benefits of Splunk as a service

Fastest time to value

Eliminates infrastructure requirements

Maximizes value from limited resources

splunk> cloud™

Managing Splunk Is More Than Just Software

	Responsibility	Splunk Ent Deployed On-Premises	Splunk Cloud
Admin Tasks: One-time Setup	Purchase/rent HW	Customer	**Splunk**
	Rack and stack, cable, network all HW	Customer	**Splunk**
	Install Splunk	Customer	**Splunk**
	Install OS	Customer	**Splunk**
	Configure Splunk (create users, load apps, configure)	Customer	**Splunk**
	Configure indexes	Customer	**Splunk**
	Setup HA/clustering	Customer	**Splunk**
	Setup disaster and recovery	Customer	**Splunk**
	Configure forwarders	Customer	**Joint**
	Onboard data	Customer	**Joint**
	Integrate with LDAP/AD	Customer	**Joint**
Admin Tasks: Ongoing	Scale up HW	Customer	**Splunk**
	Install Splunk patches / upgrades	Customer	**Splunk**
	Install OS patches / upgrades	Customer	**Splunk**
	Monitor deployment / health checks	Customer	**Splunk**
	Manage forwarders	Customer	Customer
	Create users / roles	Customer	Customer
	Manage indexes	Customer	Customer
	Onboard additional data	Customer	Customer
	Load search head only apps	Customer	**Both***
	Load distributed apps	Customer	**Both***
	Load premium apps	Customer	**Splunk**
	Export data	Customer	**Splunk**
User Tasks	Search, alerts, reports, dashboards	Customer	Customer

Managing a Splunk deployment involves 12 on-going admin tasks, 8 of which are conducted by Splunk for a Cloud based deployment

~80% reduction in management tasks

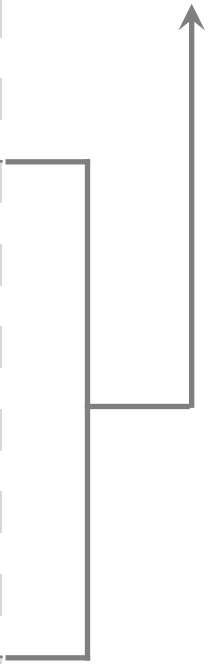

splunk> .conf20

Thank You

Please provide feedback via the

SESSION SURVEY