.conf2013

YOUR DATA
NO LIMITS

Architecting and Sizing your Splunk Deployment

Simeon Yep
Sales Engineering Manager, Splunk

#splunkconf

splunk>

Legal Notices

During the course of this presentation, we may make forward-looking statements regarding future events or the expected performance 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 could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. The forward-looking statements made in this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements we may make. 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 functionality described or to include any such feature or functionality in a future release.


.conf2013
YOUR DATA
NO LIMITS
Introduction
splunk>

About Me

- 5+ years @ Splunk
- Experience:

- Based in HQ (San Francisco office)
- Currently: Business Development, Technical Synergies

Agenda

- Sizing Fundamentals
- Architecting Fundamentals
- Deployment Topologies


.conf2013
YOUR DATA
NO LIMITS
Sizing Fundamentals
splunk>

Sizing Fundamentals

- Understand the sizing factors
- Data volume
- Search volume
- Sizing sheet

Sizing Factors

- How much data (raw sizes)?
  - Daily volume
  - Peak volume
  - Retained volume (archive size)
  - Future volume?
- How much searching?
  - Use cases
  - How many people? How often?
- Jobs
  - Summarization, alerting, reporting

Data Volumes

- Estimate input volume
  - Verify raw log sizes
  - Leverage _internal metrics to get actual input volumes
- Confirm estimates with actual data
  - Create a baseline with real or simulated data
  - Find compression rates (range from 30%-120%, typically 50%)
  - Determine retention needs
- Document use cases
  - Use case determines search needs
  - Plan for expansion as adoption grows (search and volume)

Data Sizing Exercise

- Via Filesystem
- Use the Splunk log files: metrics.log or license_usage.log
- Optionally:
  - License report view in Splunk Enterprise 6
  - S.o.S app in 5.x

Search Volumes

- Gather use case information
  - How much ad-hoc searching?
  - How much background searching?
- Ad-hoc searching
  - Evaluate the data being searched
  - Evaluate the time duration (real-time vs historic)
  - Real-time searches are typically less overhead
- Background searching
  - Alerting and monitoring
  - General reports
  - Summary indexing

Final Sizing Numbers

- Data capacity
  - Daily and peak
- User capacity
  - Concurrent and total
- Search capacity
  - Concurrent and total

*Document the use cases!!


.conf2013
YOUR DATA
NO LIMITS
Architecture
splunk>

Architecture

- Splunk server roles: distributed/clustered deployments
- Reference server
- Rules of thumb
- Hardware factors

Splunk Distributed Roles

Search Head (regular and job server)

search head

Indexer

indexer

Forwarder (universal)

forwarder

Splunk Distributed Roles

Cluster Master (clustering/replication requirement)

License Master

Deployment Server

Recommended Configurations

	Stand-alone	Indexer (distributed)	Search head (distributed)	Indexer (clustered)	Search head (clustered)	Cluster master (clustered)
Forwarding	*	*	*	*	*	
Searching	√		√	*	√	
Indexing	√	√	*	√		
Deployment server		*	*			
License master		*	√		*	*
Cluster master						√

√ - common * - uncommon

What's a "Reference" Server?

- Sizing based on commodity x86 servers
- Dual quad-core CPUs at 3.0 GHz (dual six core is common)
- 8 GB of RAM – (16 GB is common)
- 64-bit OS
- **4x10k RPM local SAS drives in RAID 1+0 (800+ IOPs)**
- Variations cause corresponding changes in performance/ requirements

Rules of Thumb

- These all have exceptions and qualifications
- 1 reference indexer per 100 GB/day
- 1 reference search head per 8 to 12 users
- 1 reference job server per 20 concurrent jobs
- 1 deployment server per 3000 polls/min
- Replication later…

How Many Indexers?

- Rule of thumb says: 1 per 100 GB/day
- Leaves room for:
  - Daily peaks
  - Light searching and reporting for about 3 concurrent users
- Need more indexers for:
  - Heavy reporting
  - More users
  - Slower disks, slower CPUs, fewer CPUs

How Many Search Heads?

- Rule of thumb says: 1 per 8 to 12 concurrent users
- Limit is concurrent queries
- 30-50 web sessions
- 1:1 ratio of search query to CPU core
- Only add first search head if ≥3 indexers
- Don't add search heads; add indexers: indexers do most work
- But you need more if:
  - Running a lot of scheduled jobs on the search head

Search Head vs. Job Server

- Search Head Pooling (SHP): uses NFS to manage user profiles/ configurations and job queue
- Search head and job server are equivalent with SHP
- Use job servers for scheduled searches (summaries, alerts, and reports)
- Use search heads for ad-hoc searching

How Many Deployment Servers?

- Rule of thumb says: 1 per 3000 polls/minute
- Just use one deployment server, and adjust the polling period
- Small deployments can share the same splunkd
- Low requirement for disk performance (good candidate for virtualization)
- Windows OS – 1 per 500 polls/minute
- Or use something other than deployment server
  - Puppet, SCCM, cfengine, chef…

More is Better?

- CPUs
  - Search process utilizes up to 1 CPU core (1:1)
  - Indexers still need to do the heavy lifting (search exists on indexer AND search head)
  - Limited benefit for indexing (up to 2 CPU cores for indexing)
- Memory
  - Good for search heads and indexers (16+ GB)
- Disks
  - Faster is better (15k rpm)
  - More disks in RAID 1+0 = faster

Performance and Sizing Tips

System change	Search Speed	Indexing Speed
Faster disks	++	++
Add an indexer	++	++
Add a search head	+	
Report acceleration/ summaries	++	

Performance and Sizing Tips

System change	Search Speed	Indexing Speed
Optimize searches	+++	
Optimize field extraction	+	
Optimize input parsing		+
Faster CPU	+	+

Capacity → Architecture

- Sizing recipe
  - Capacity
  - Rules of thumb determines number of servers
- Building blocks for architecture

Architecture Factors

- What are my sizing requirements?
- Where is the data?
- Where are the users?
- What is the security policy?
- What are the retention and compliance policies?
- What is the availability requirement?
- What about the cloud?

Architecture Factors

- What are my sizing requirements?
  - Data capacity
  - Search capacity
  - User capacity
- Obtained from the sizing process

Architecture Factors

- Where is the data?
  - Local or remote to the indexing machine
  - If remote – use forwarders when possible
  - Index in local data center (zone) or index centrally
  - Persist network data to disk as a best practice
  - Use intermediate forwarders to distribute data
- Where are the users?
  - User experience affected by search head location
    - Time zone tuning (5.x +)
    - Distributed search over LAN vs WAN

Architecture Factors

- What is the Security Policy?
  - Apply user security policies
    - Auth method
    - Roles
    - Filters
  - Apply physical security policies
    - Index location

Architecture Factors

- Retention, compliance, governance
  - Where is the data allowed to be?
  - Where is the data *not* allowed to go?
  - Where *must* the data go?
- Availability
  - Local failover, fault-tolerance, clustering
  - Geographic disaster recovery/fault-tolerance
  - Index replication!

Architecture Factors

- Same old story
- Cloud considerations
  - Authentication restrictions
  - Data transfer costs
  - Security – SSL tunnel
  - Zones


.conf2013
YOUR DATA
NO LIMITS
Topologies
splunk>

Architecture Factors → Topology

Topology Examples		
Centralized	Decentralized	Hybrid

Centralized Topology

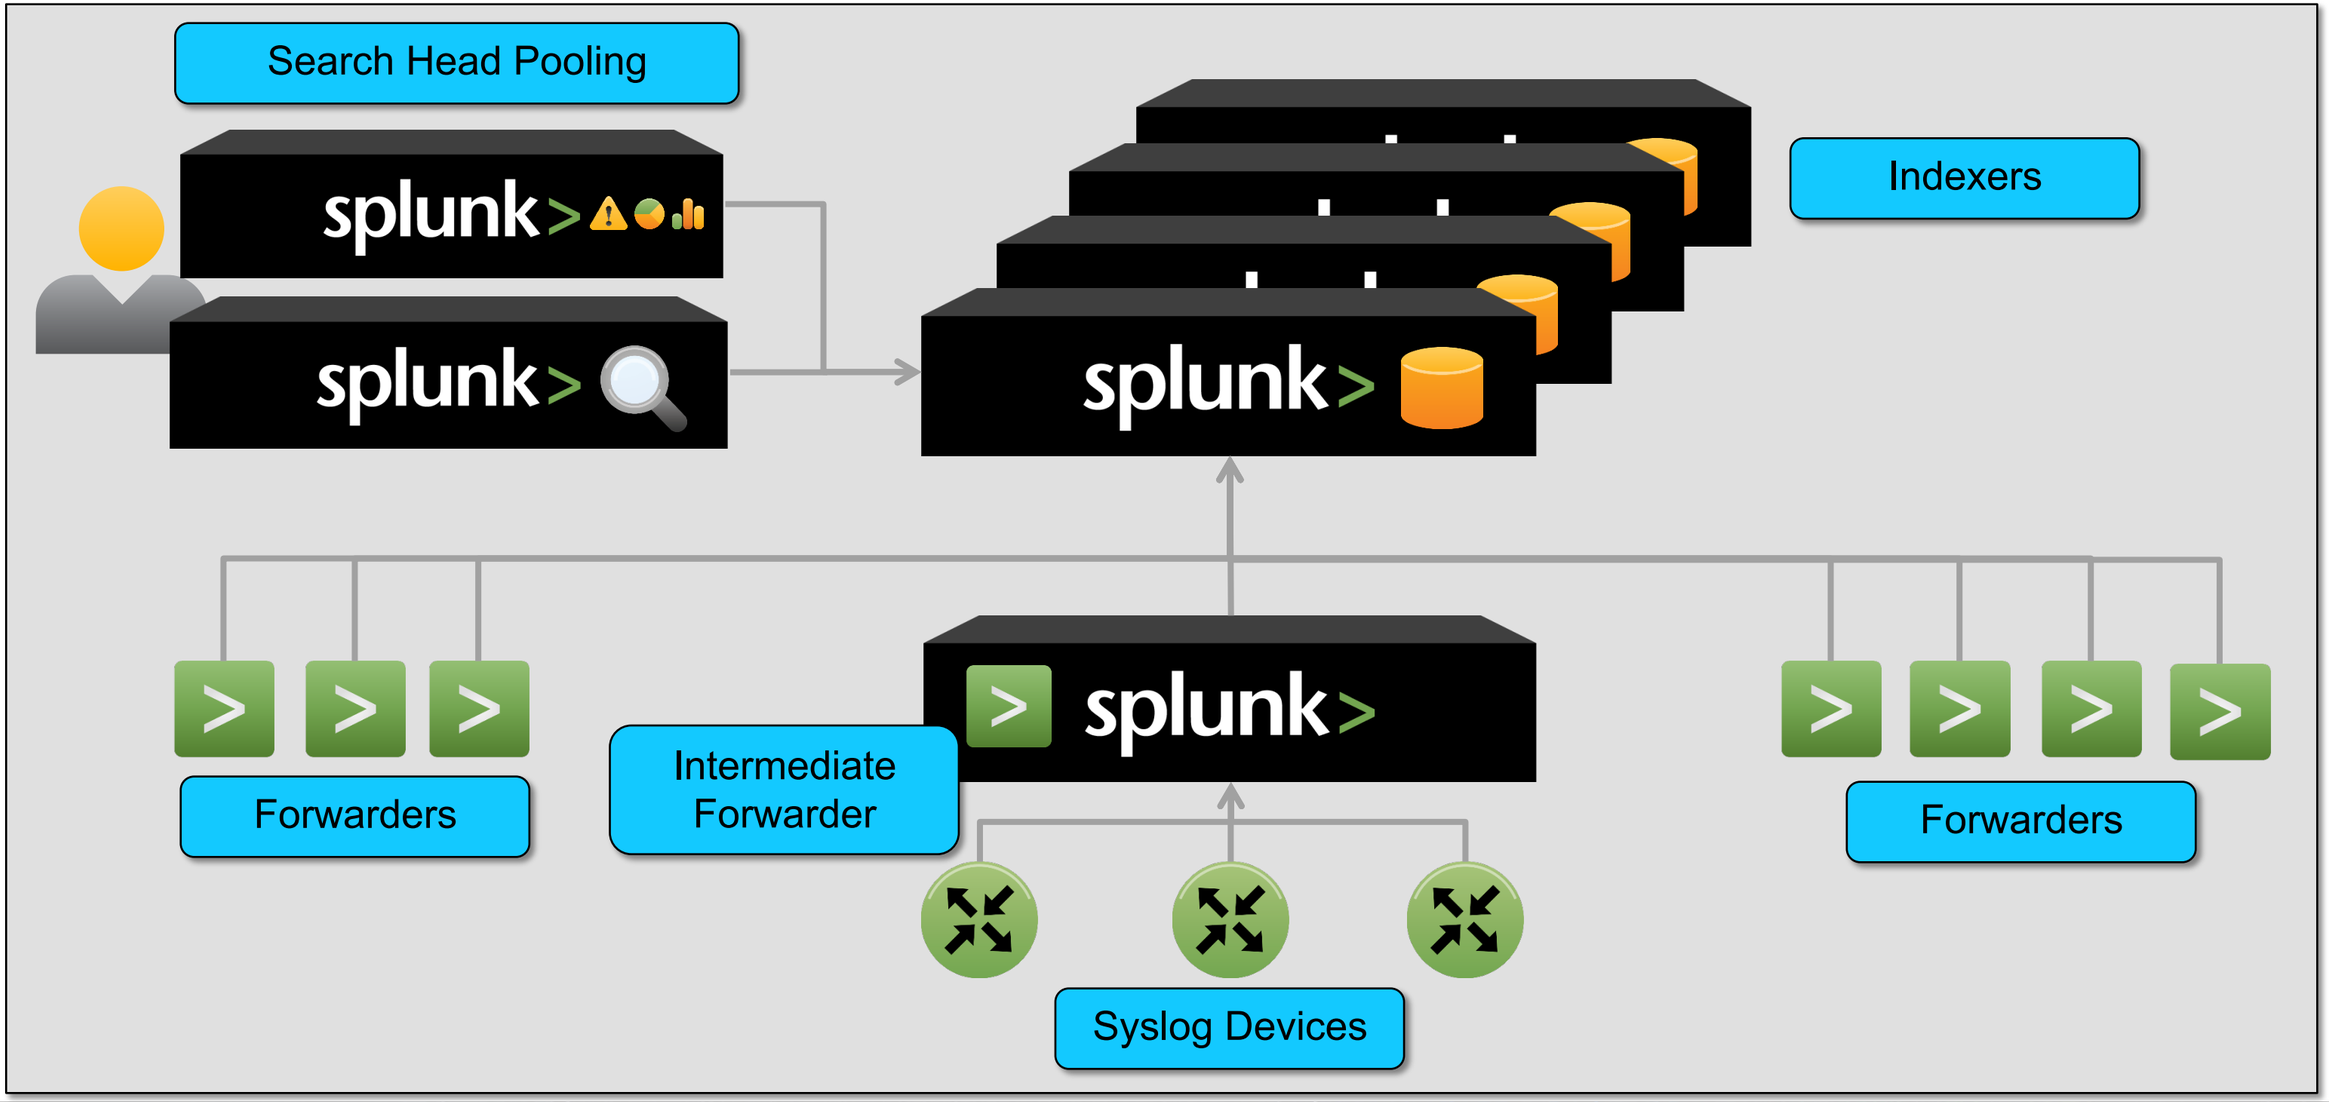

Decentralized Topology

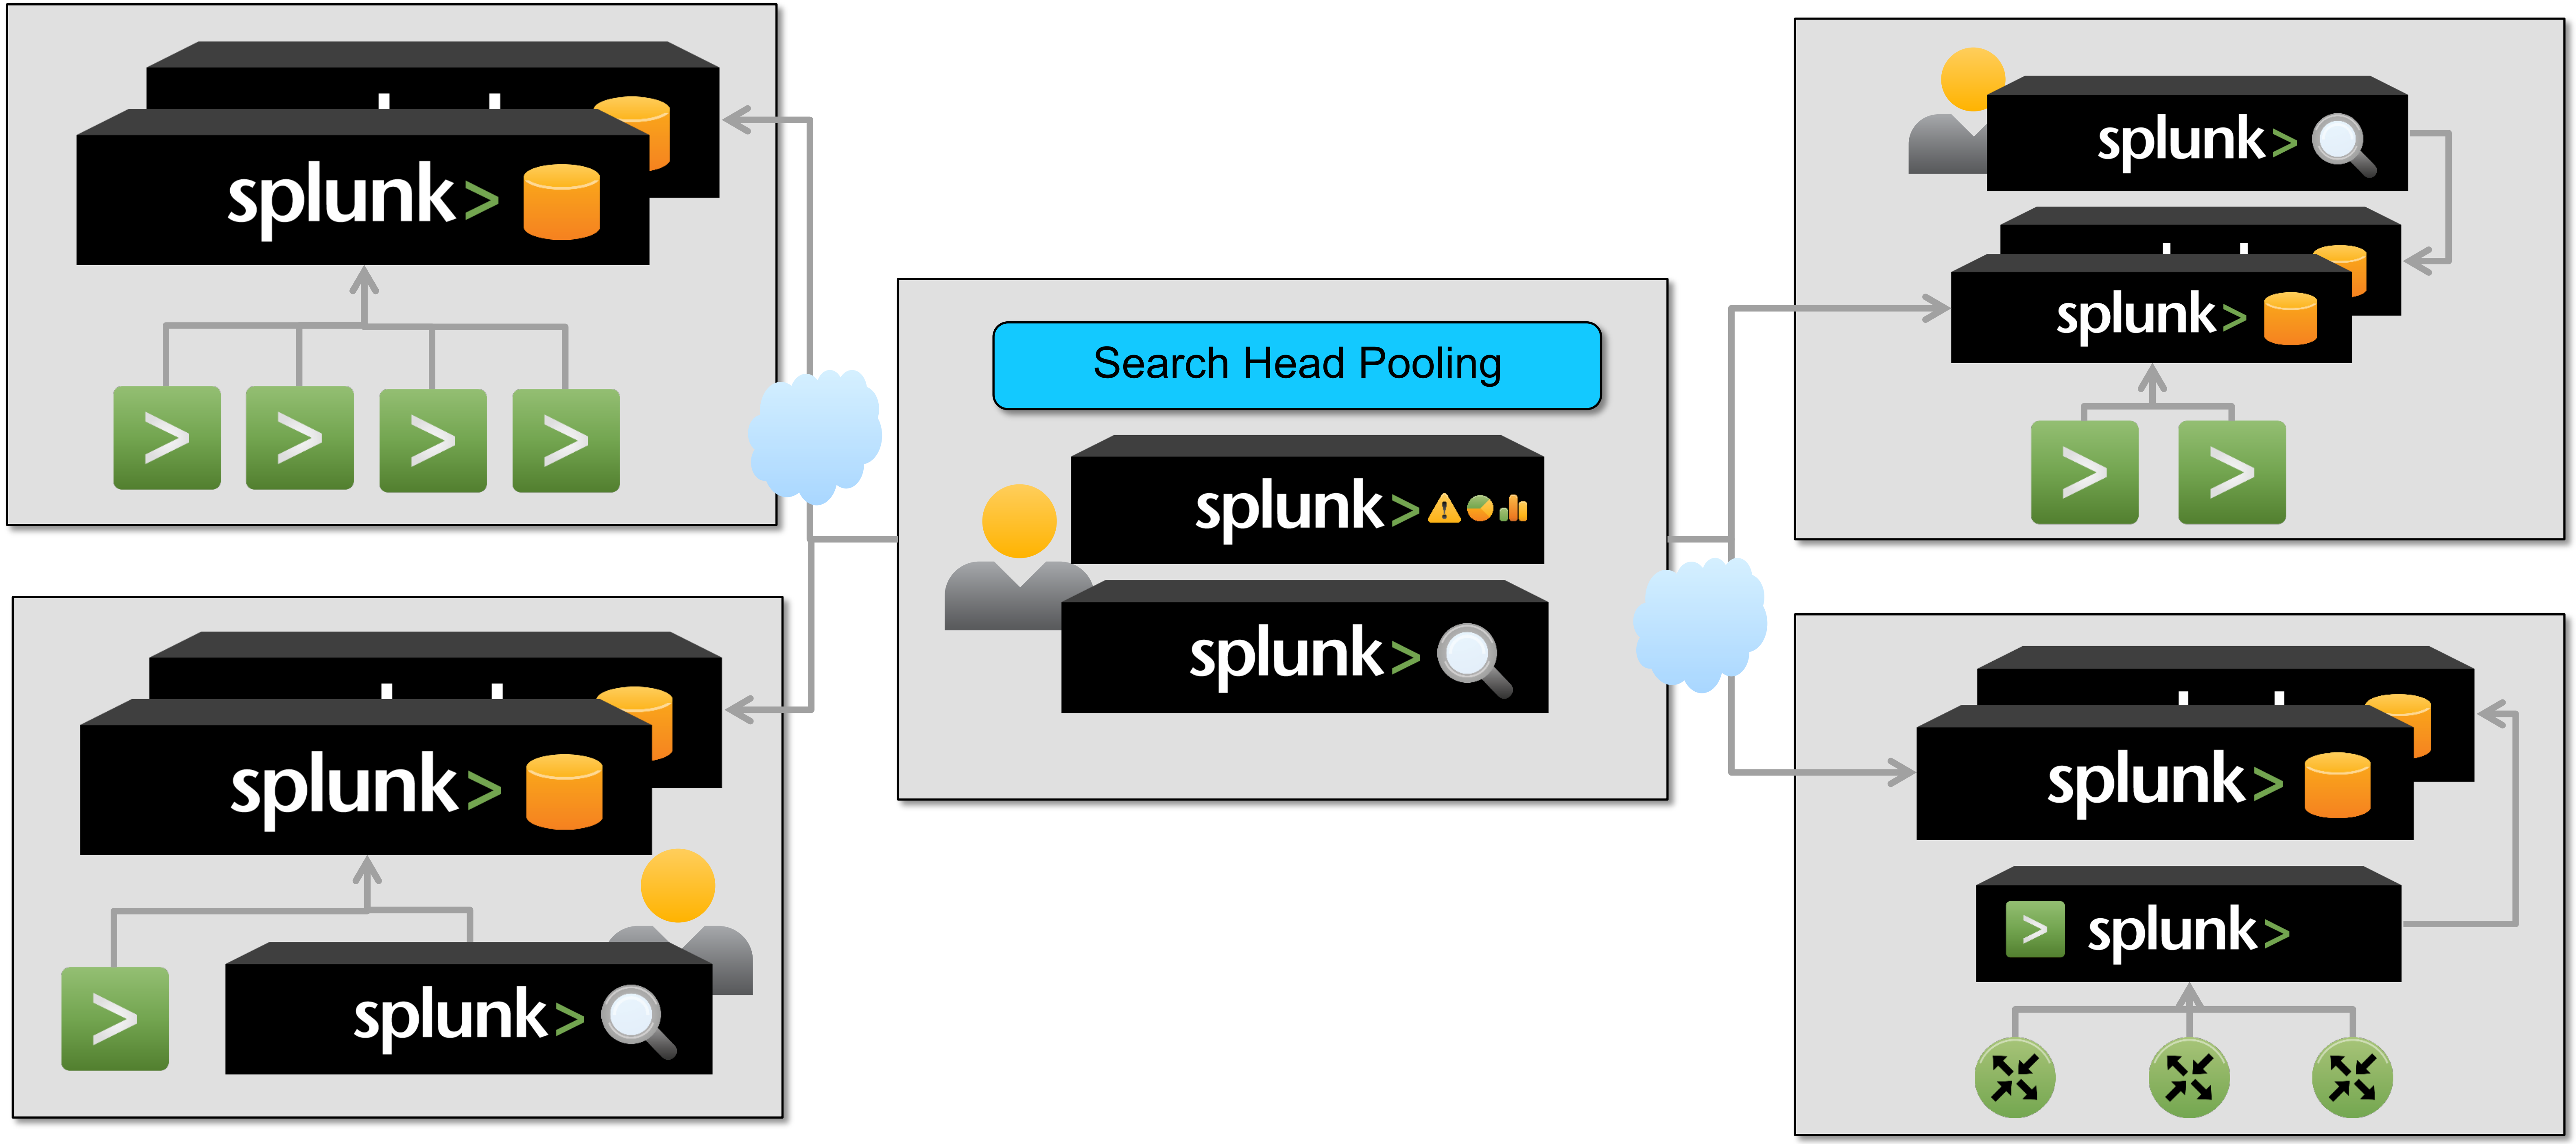

Hybrid Topology

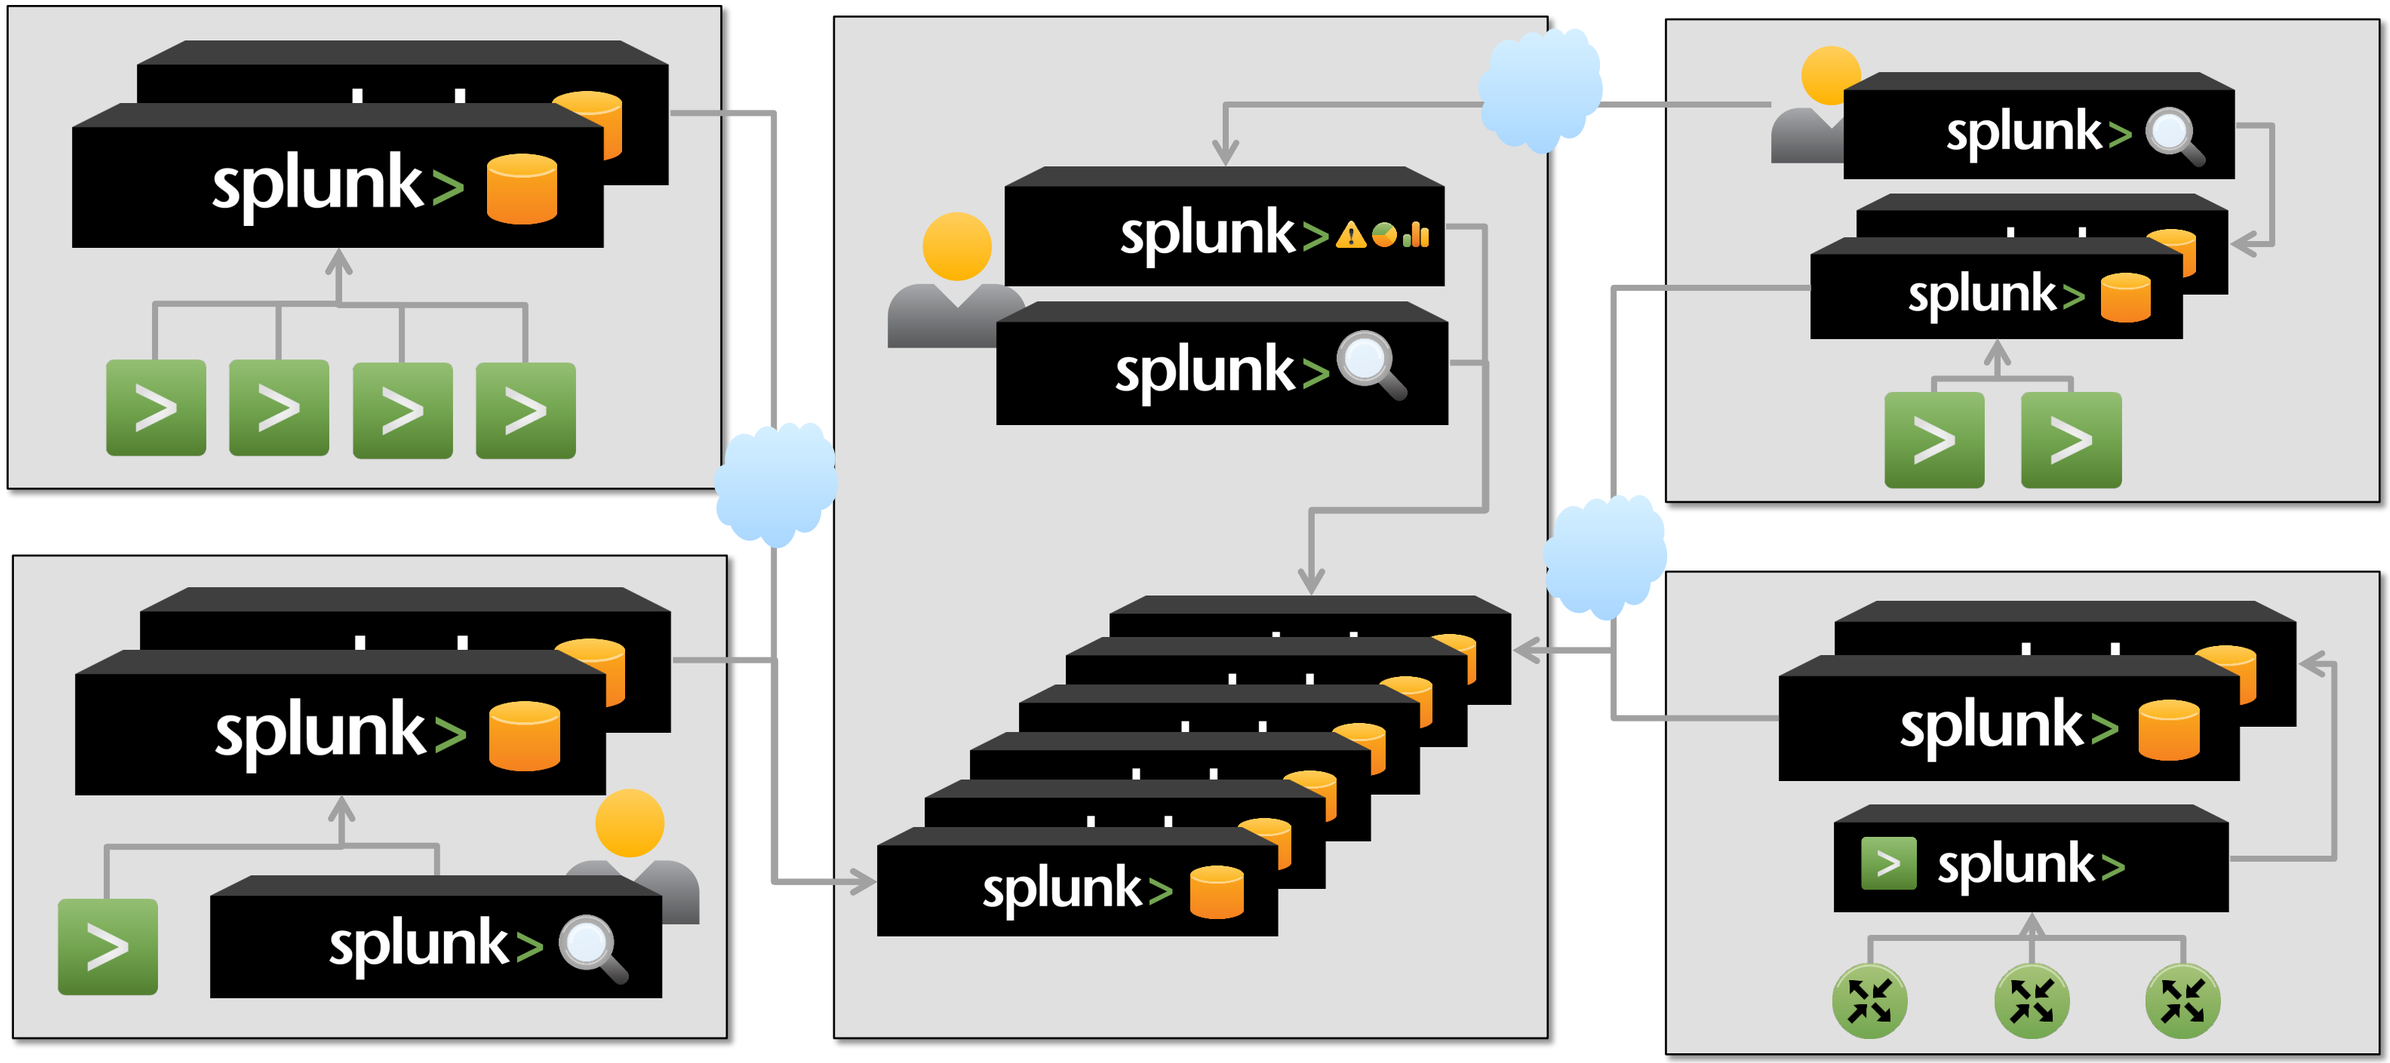

Scaling and Expansion

- Add to your indexer pool for more performance or capacity
  - Mixed platform and hardware is okay
- Use search head pooling for more UI capacity
  - Requires NFS
- Create new indexes for new data types
  - Follows best practices

Index Replication (aka Clustering)

- What is it?
  - Indexes are replicated to 1 or more indexers (tunable)
  - Splunk controlled
- Basics
  - Master node (manages indexing and searching location)
  - Distributed deployment
  - NOT = “index and forward “
- HA license VS. index replication
  - HAL – Separate fully functioning Splunk deployments
  - IR – Data is made available on 1 or more indexers

Index Replication

- Rule of thumb: 1 per 50 GB/day
  - Assume simple replication (2 in existence)
  - Increase in I/O, CPU, and disk requirement
- Need more indexers if:
  - Increase in replication factor
  - Performance or capacity needs (search and index)

Index Replication (aka – Clustering)

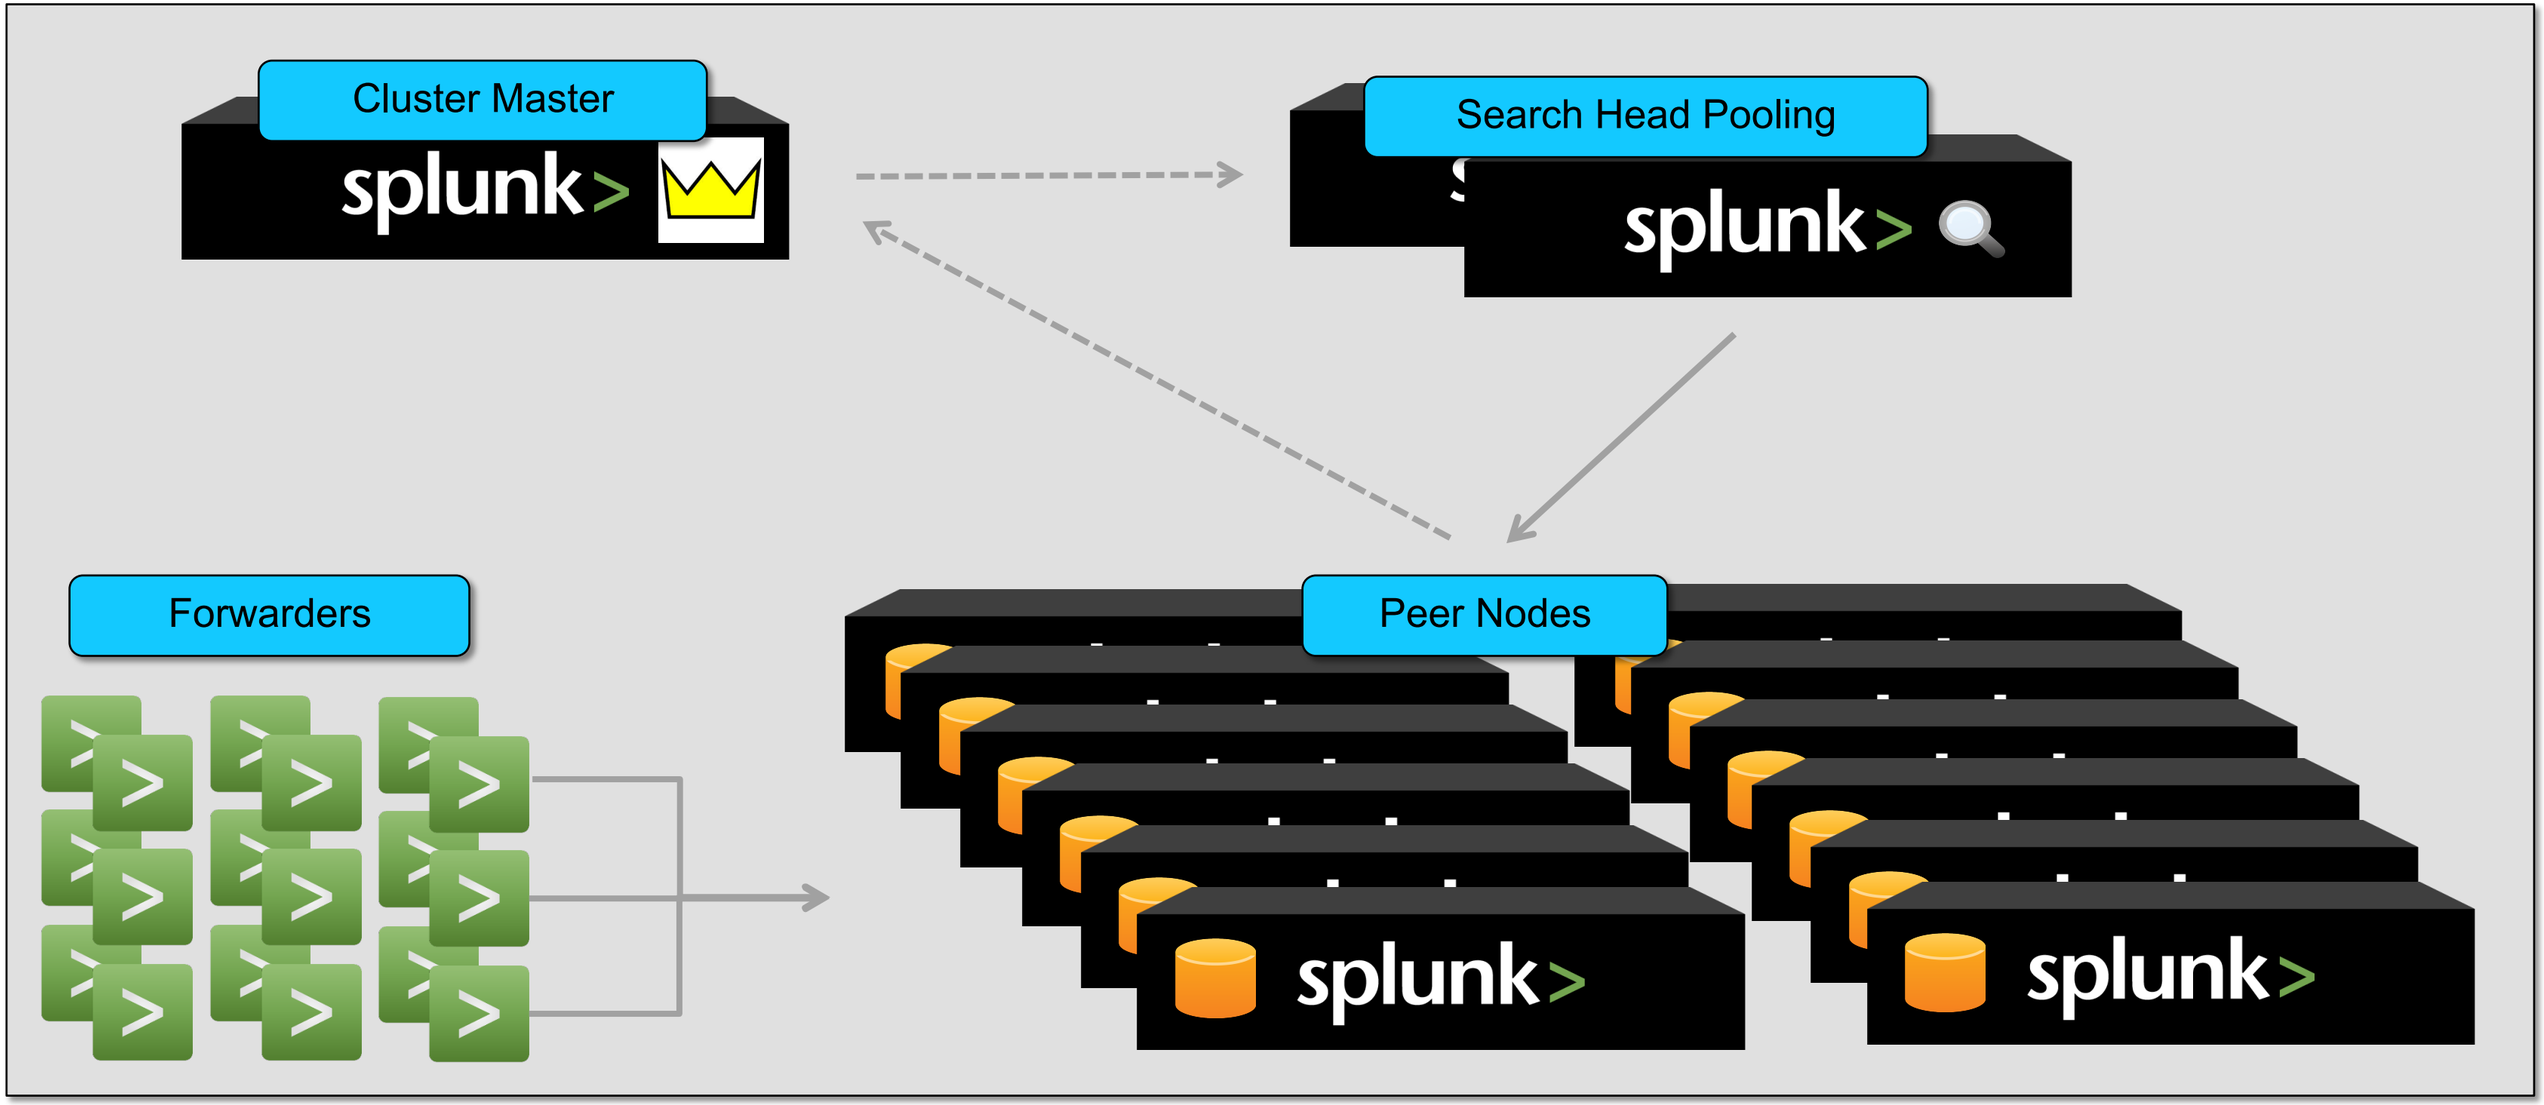

Index Replication

- Data is replicated and made available
- WAN configuration is not recommended
- Careful consideration when inserted into standard topologies
- Increases
  - Storage requirement
  - I/O requirement (disk, network)
  - Total indexer requirement
- Disaster recovery and high availability .conf session

Final Thoughts

- Sizing is more than data volume – it's also search load
- Centralized architecture is the baseline
- Variations on architecture are driven by
  - Sizing
  - Data location
  - User location
  - Retention/access/governance
  - Availability requirements

Next Steps

1. **Download the .conf2013 Mobile App**
   If not iPhone, iPad or Android, use the Web App

2. Take the survey & **WIN A PASS FOR .CONF2014**… Or one of these bags!

3. **View the sessions listed on the next slide**
   Available on the Mobile App

More Information

- Contact: syep@splunk.com
- Documenation: http://docs.splunk.com
- Answers: http://answers.splunk.com
- Other presentations
  - Best Practices: Deploying Splunk on Physical, Virtual, and Cloud Environments
  - Architecting Splunk for High Availability and Disaster Recovery


.conf2013
YOUR DATA
NO LIMITS
THANK YOU
splunk>