Buckets Full of Happy Tiers

Cory Minton – Principal SE @ EMC

Jenny Hollfelder – Global Alliances SE @ Splunk

.conf2016

splunk>

Disclaimer

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 the 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.

J. Cory Minton

Principal SE and Global Data Fabrics Leader

- 2015 Principal SE Class
- EMC Elect ‘15-16
- Lead Global Ninja Team
- Certified Splunk SE
- Oracle and SAP Consulting Background
- BS Engineering and MBA
- Professional Chef on side

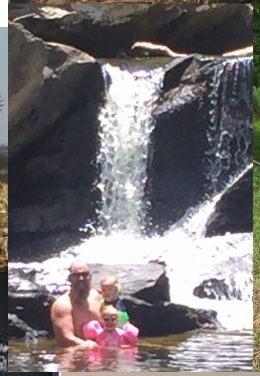

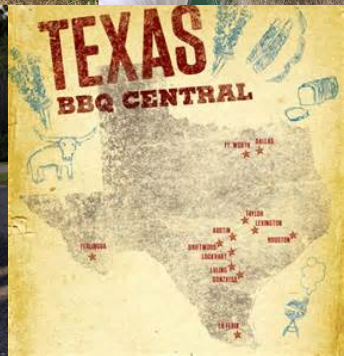

Jenny Hollfelder

Sales Engineer, Global Strategic Alliances

- Splunk Certifications
  - Architect
  - Consultant I, II
  - SE I, II and III
- Splunk Customer since 2007
  - Spoke at first SplunkLive! in 2009
  - Wrote Splunk app for Bacula
- Linux Administration Background
- BS Computer Science
- Professional Hooper on the side (not really... but I do hoop a lot!)

Splunk Trends we are seeing…

Splunk is now a business critical application:

- ✓ Size of Splunk deployments is increasing rapidly
- ✓ Splunk search performance is important
- ✓ Availability/Reliability is becoming crucial
- ✓ Underlying infrastructure is aligning to enterprise IT
- ✓ Efficiency at scale is becoming a top concern

EMC provides a scalable and efficient enterprise solution for deploying Splunk.

Why Does Splunk Grow?

Performance

- ✓ Volume Of Ingest
- ✓ Search Performance
- ✓ More Users
- ✓ Big Apps

Capacity

- ✓ Volume Of Ingest
- ✓ Index Retention Periods
- ✓ Indexer Clustering
- ✓ Big Apps

EMC provides a scalable and efficient enterprise solution for deploying Splunk.

Scaling

Then comes the real world…

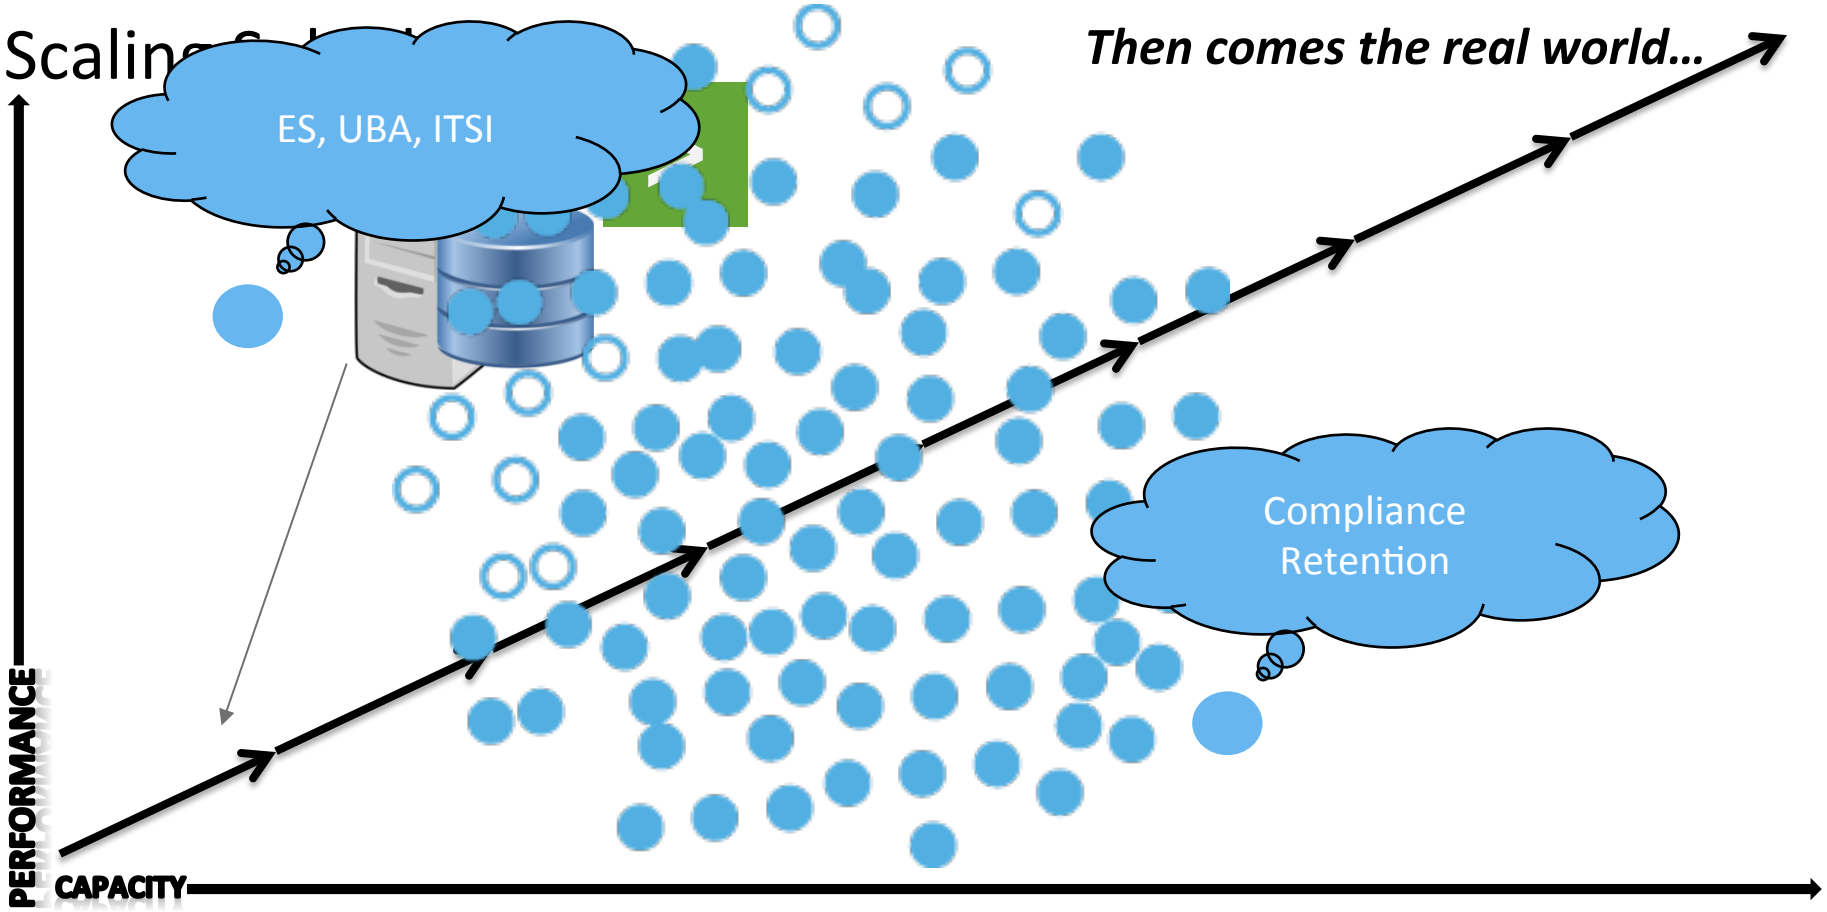

Splunk: Decouple Compute From Storage

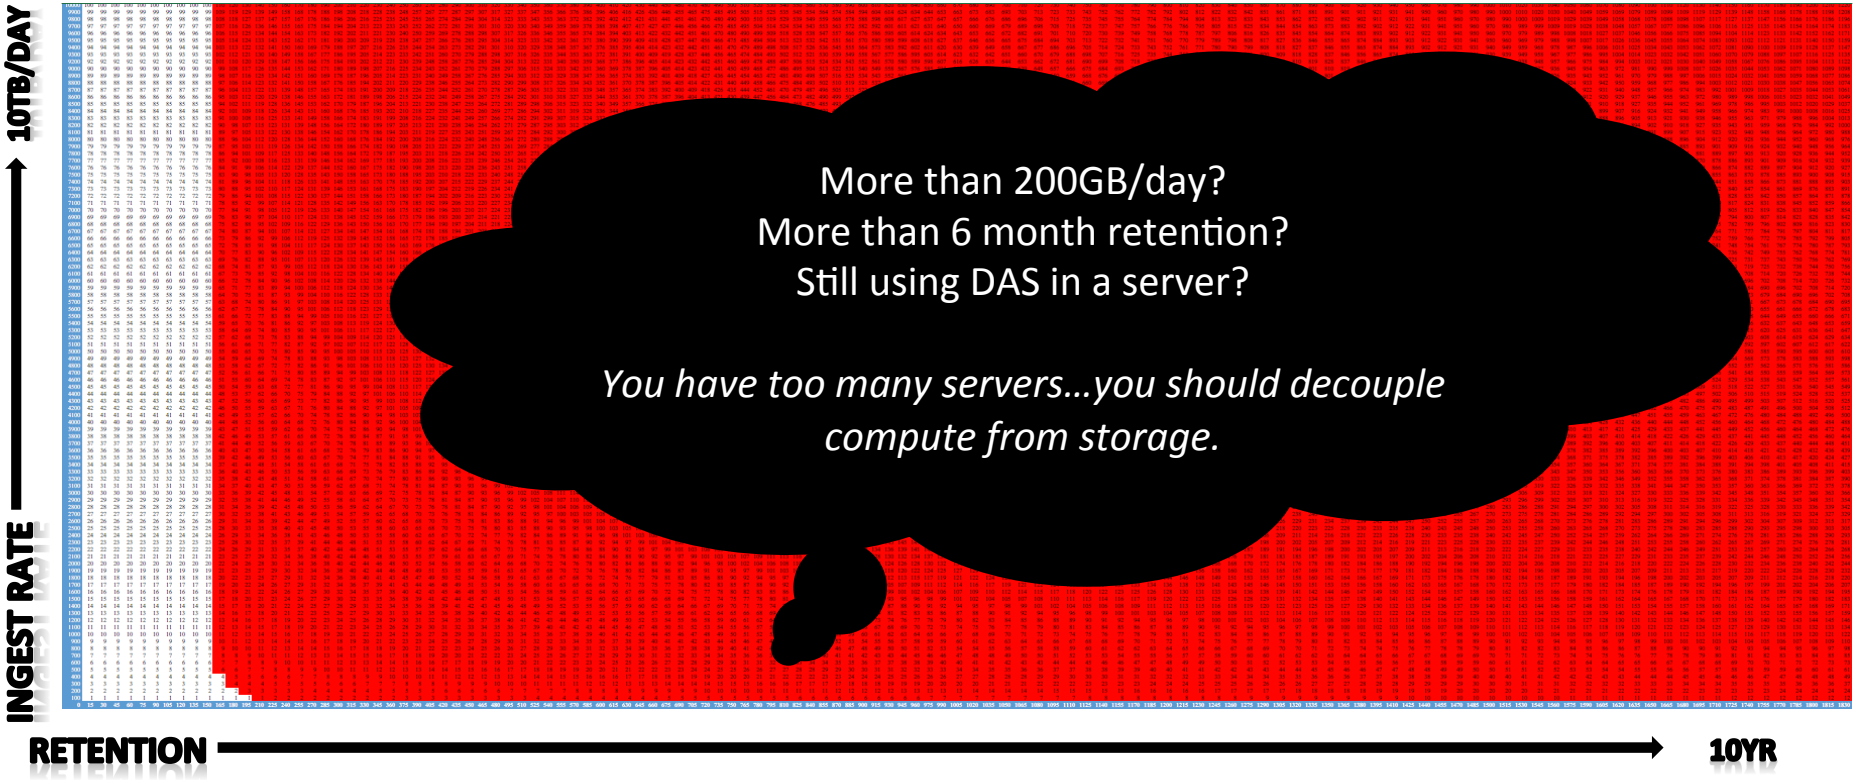

Splunk Index Architecture

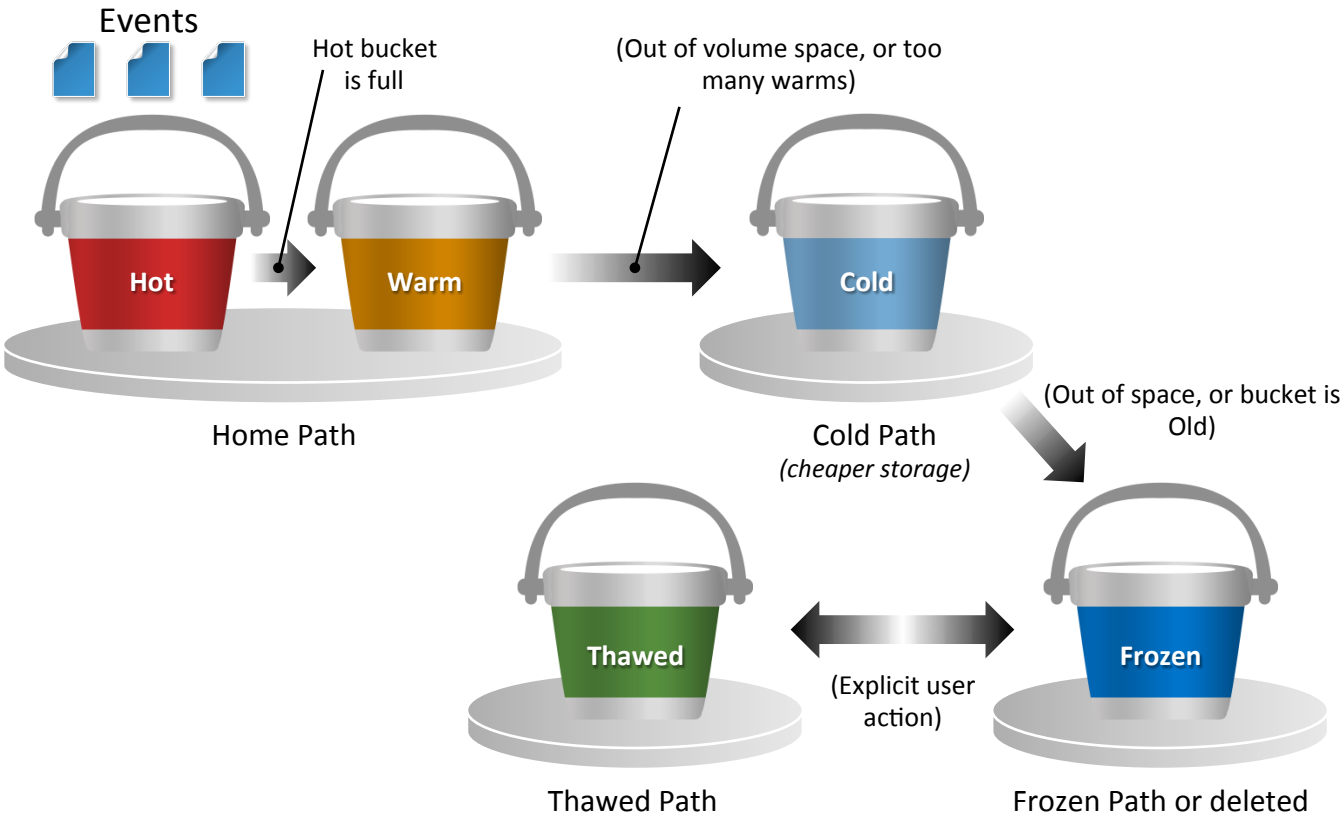

Why EMC For Splunk

Optimized infrastructure for big & fast data

Optimized Shared Storage & Tiering

Hot & Warm Data Deployed On XtremIO or ScaleIO

Cold & Frozen Data Deployed On Isilon

Cost-Effective & Flexible Scale-Out

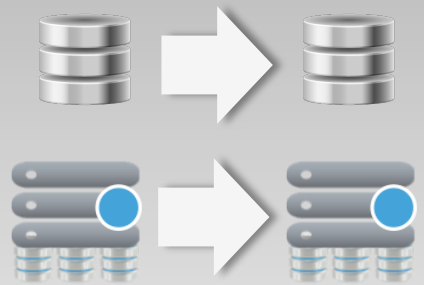

Scale-Out Capacity & Compute Independently Or As Converged Platform

Powerful Data Services

Encyption & Security

Index File Compression

Snapshots For Backups

Deduplication Of Clustered Indexes

EMC Xtremio & Splunk

All-flash high speed infrastructure for hot & warm data

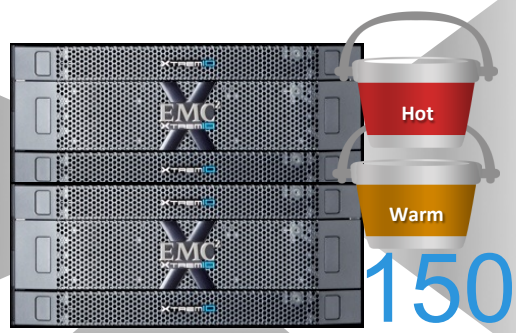

High-Speed Search

Accelerate SuperSparse & Rare Searches

Scale-Out Flash For I/O-Bound Data

>1M IOPS & <1ms Latencies

5TB → 320TB

Flash

150K IOPS/Brick

Data Services For Hot & Warm Data

Self-Encrypting Flash Drives

Index File Compression

In-Memory Data Copy Services

Dedupe Clustered Index Copies

EMC Scaleio & Splunk

Existing server + disk capacity leveraged to hold hot and warm data

SUPREME ELASTICITY	Hyper-converged infrastructure for Splunk
UNPARALLELED FLEXIBILITY	Leverage under-utilized HD/SSD storage from existing Servers
STRONG PERFORMANCE	10M+ IOPS
MASSIVE SCALABILITY	1000+ Servers
COMPELLING ECONOMICS	Reduce CAPEX 50%

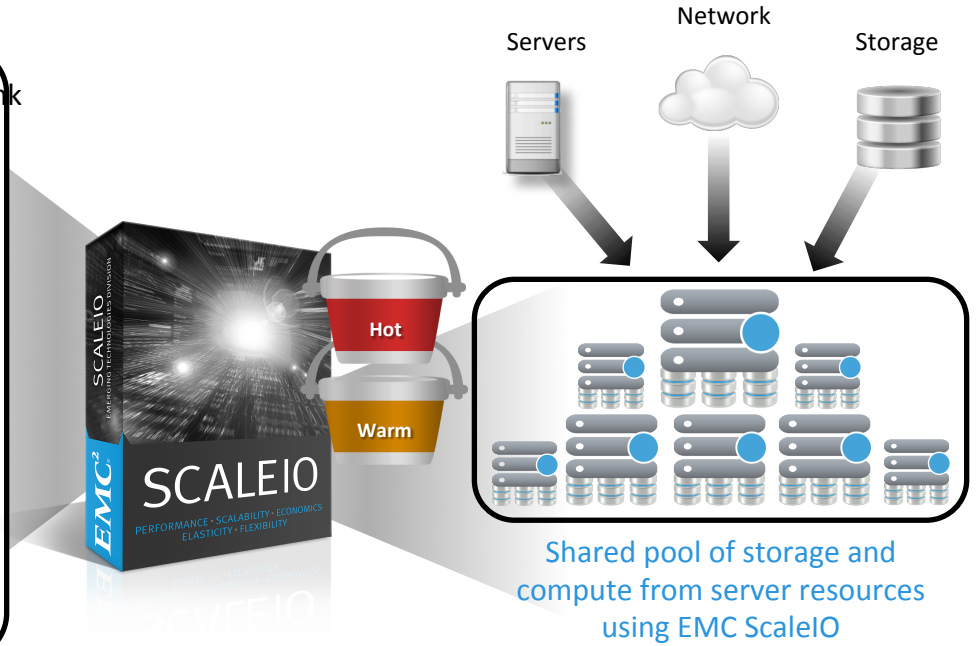

SOFTWARE-DEFINED SCALE-OUT SAN

EMC Isilon & Splunk

Low-cost & secure scale-out for cold and frozen data

Hadoop and Splunk

Single Common Data Lake for Hadoop and Splunk

Bottomless Cold Bucket

No need to freeze or thaw.
Single Volume scaling to

68 PB

Powerful Data Services

Encyption & Security

Index File Compression

Snapshots For Backups

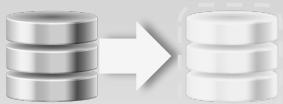

Deduplication Of Clustered Indexes

EMC Reference Architectures For Splunk Enterprise

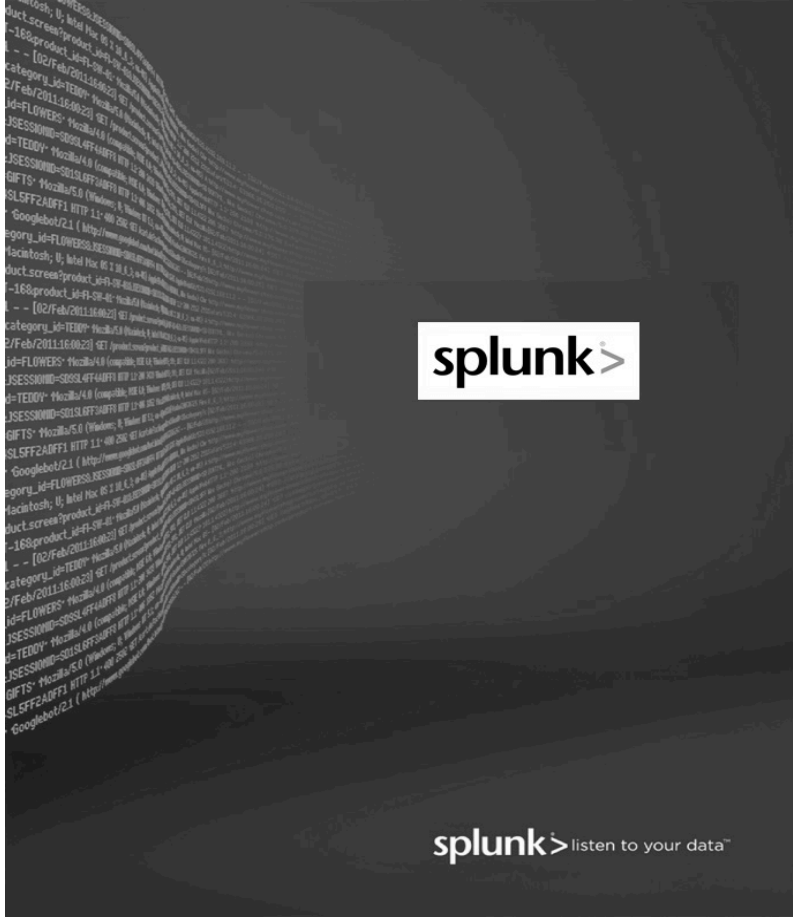

[XtremIO and Isilon Reference Architecture]

[ScaleIO and Isilon Reference Architecture]

Why Vce For Splunk?

Cloud-like experience, on-premise solution

SPEED TIME-TO-DEPLOYMENT

- Factory Physical and Logical Build
- Compliance-Ready
- Performance and Availability

SIMPLIFY ONGOING OPERATIONS

- Roadmap and New Feature Planning
- Configuration and Patch Management
- Single Support Through VCE

Scaling Splunk

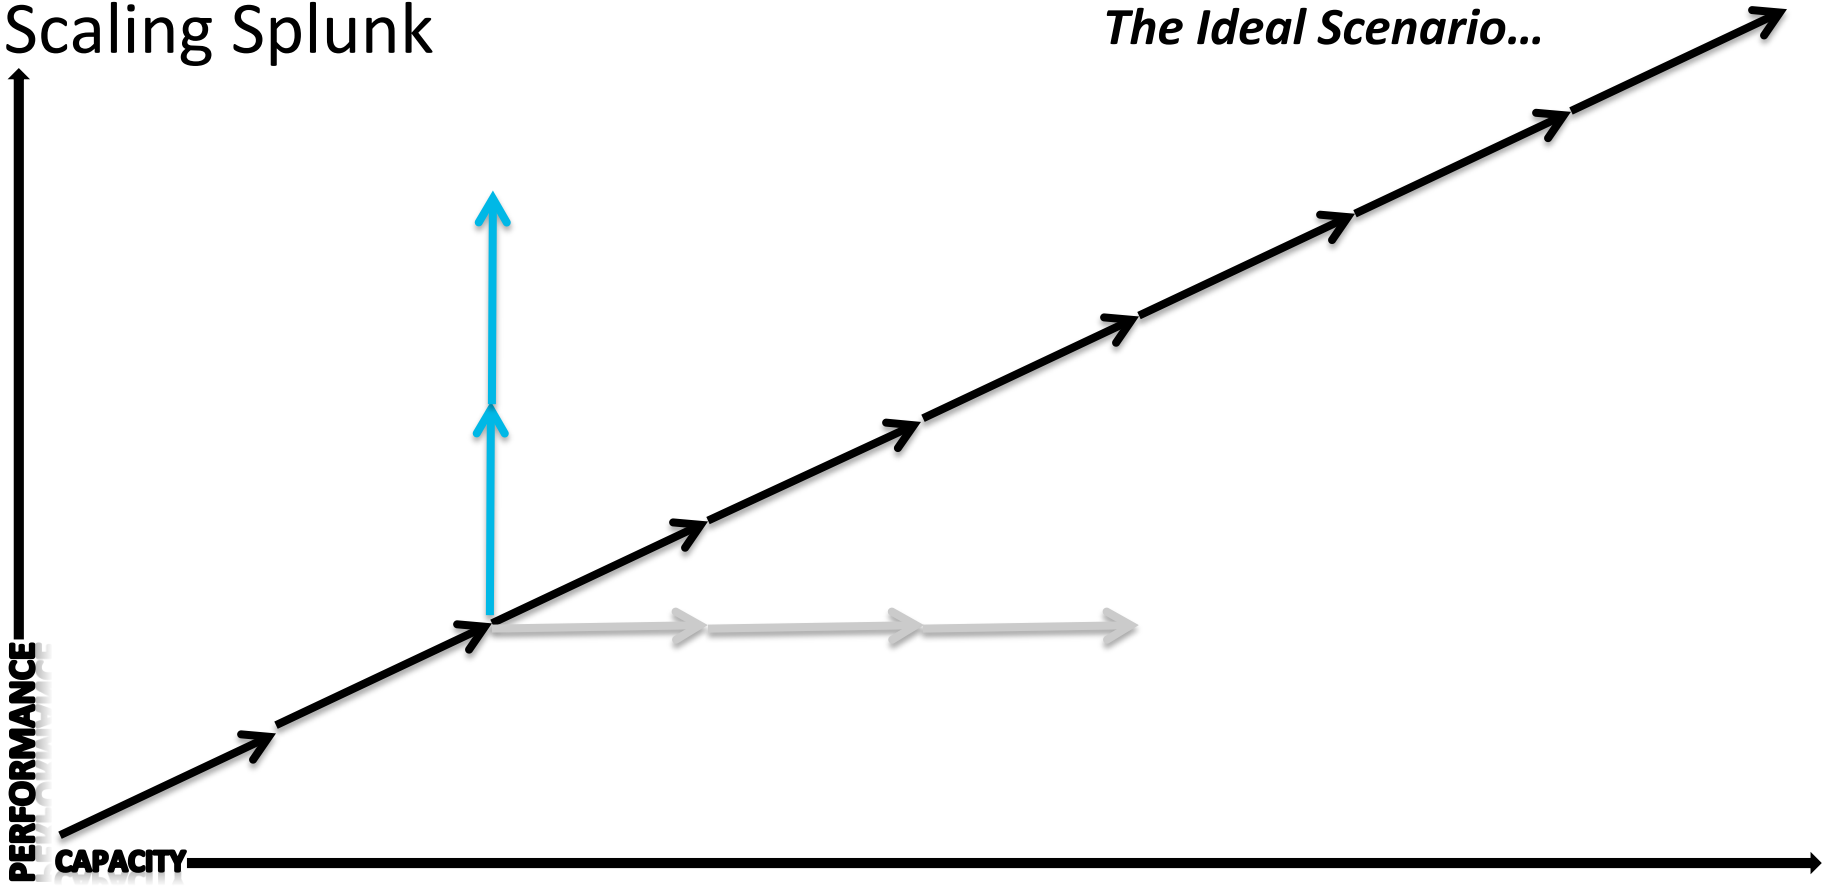

Splunk Apps From EMC

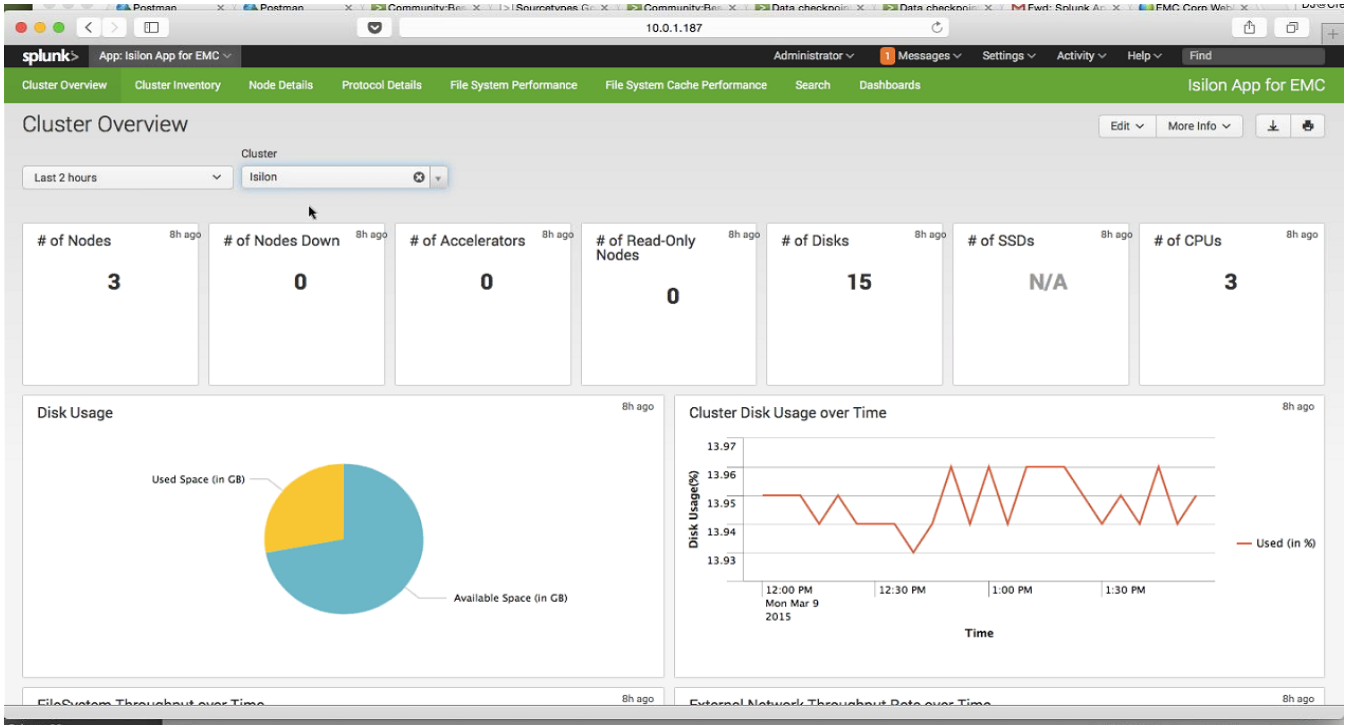

[XtremIO App]

[Isilon App]

[Isilon FS Audit App]

[VNX App]

Demo Time!!!

Jenny Hollfelder – Splunk SE Global Alliances

Cory and Jenny's Top 10

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.

#1 - Follow published EMC best practices

- Do it across EMC Platforms
  - VMware
  - XtremIO
  - VNX
  - VMAX
  - ScaleIO
  - VSAN
  - Isilon
  - ECS

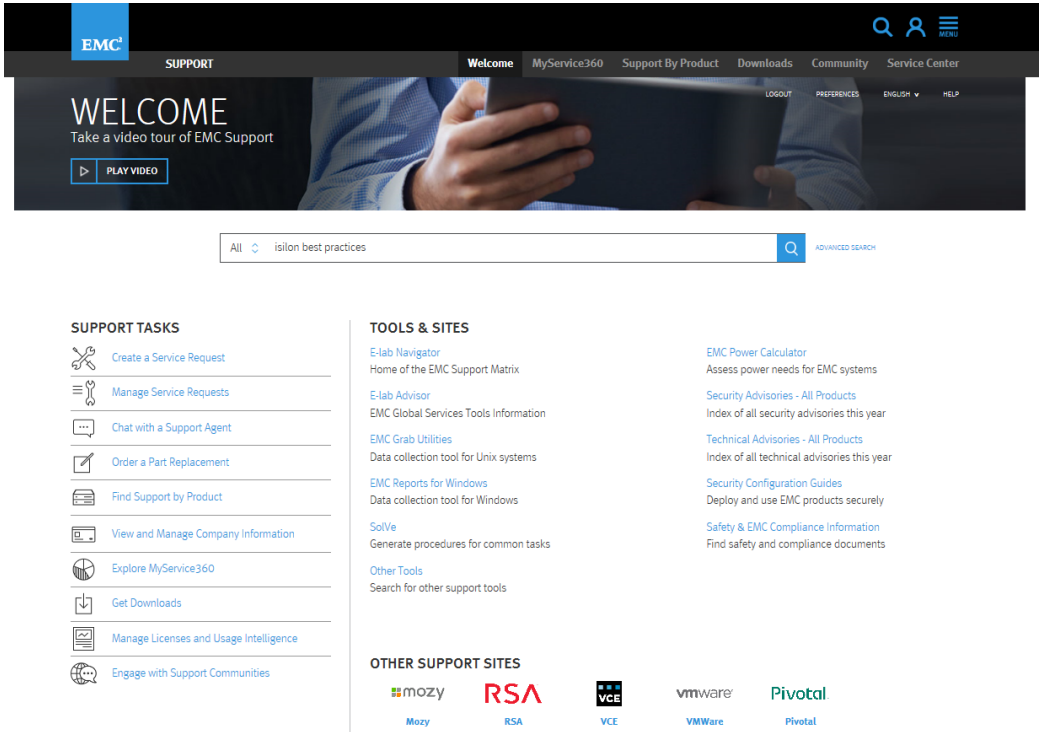

In the absence of Splunk-specific Best Practice, use base.

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!

#2 – Use Splunk Best Practices

- Disable Transparent Huge Pages (THP)
- Increase ulimit for number of open files AND number of processes

```
splunk          hard    nofile          65535
splunk          soft    nofile          32000
splunk           soft    nproc           16384
splunk           hard    nproc           16384
```

- Enable Hyper-Threading Technology (HHT or HT)
  - Provides ~30% performance benefit
  - ALWAYS allocate the equivalent of the recommended number of physical cores when provisioning resources in a virtualized environment

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure

#3 – Watch The Levers You Pull…

Performance

- ✓ **Volume Of Ingest**
- ✓ Search Performance
- ✓ More Users
- ✓ Big Apps

Capacity

- ✓ Volume Of Ingest
- ✓ **Index Retention Periods**
- ✓ Indexer Clustering
- ✓ Big Apps

Scaling is easy, if we know what is scaling!

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.

#4 – Networking Matters

- Dedicated Fabrics
  - H/W is FC SAN
  - Cold/Frozen is IP LAN
- Use 10GB...trust us.
- Not a fan of converged fabrics for Splunk…today.

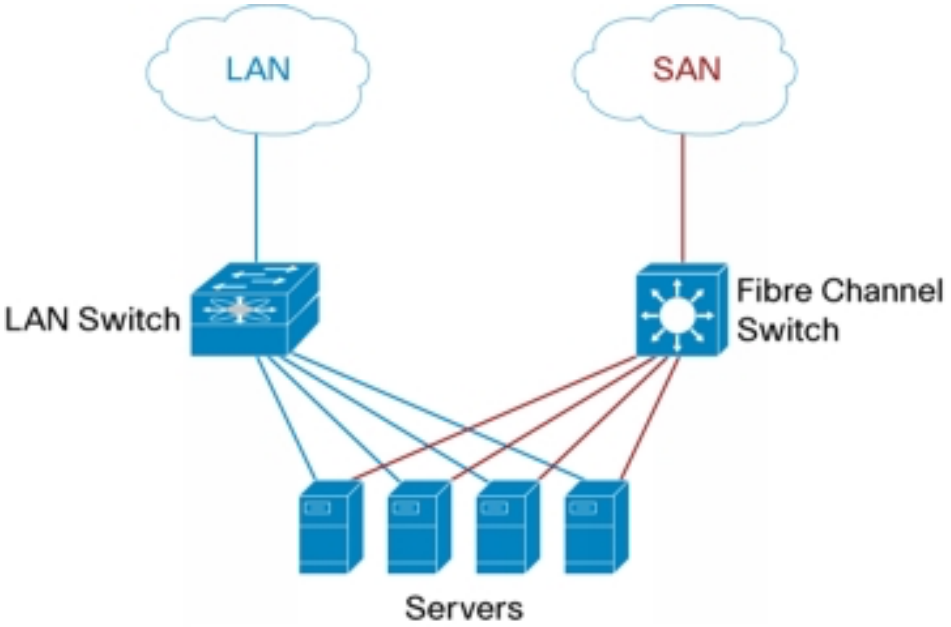

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.
5. If using Isilon for Cold, configure SmartConnect Advanced PROPERLY!

#5 – Do Isilon SmartConnect Advanced RIGHT!!!

- SmartConnect = Load Balancer for Isilon Cluster
- Use it with FQDN for NFS or…

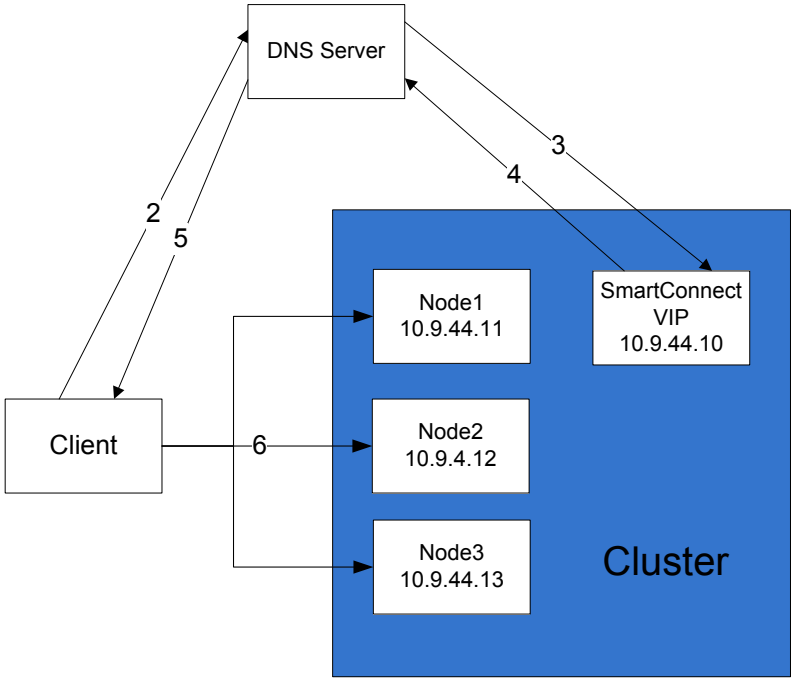

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.
5. If using Isilon for Cold, configure SmartConnect Advanced PROPERLY!
6. If you are going virtual and using XtremIO, us vSphere 5.5+.

#6 – Use VMware 5.5+ with XtremIO

- It's all about automating space reclamation.
- Check out this blog.
- XtremIO Host Configuration guide - key settings described on pages 36-38

```
# mount -o discard /dev/vg_ext4_test/lv_ext4_test /mnt/ext4
```

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.
5. If using Isilon for Cold, configure SmartConnect Advanced PROPERLY!
6. If you are going virtual and using XtremIO, us vSphere 5.5+.
7. Traditional Benchmarks are not to be trusted…

#7 - Existing Benchmark Tools Not Accurate

	Advantages	Drawbacks
Bonnie++ (http://www.coker.com.au/bonnie++/)	• Simple to install and test • Current Splunk standard for benchmarking IOPS (random seeks per second) • Splunk App to guide usage of running and collecting bonnie++ results (https://splunkbase.splunk.com/app/3002/)	• Differences in results depending on bonnie++ version (1.03e vs. 1.96) • Latency numbers are questionable • **Poor at accurately indicating performance of Splunk with shared storage environments** (block size, read/write workload, etc)
SplunkIT (https://splunkbase.splunk.com/app/749/)	• Simple to setup • Measures indexing throughput • Measures search response time	• No multi-user scenarios • No concurrent/streaming workloads • No support for distributed environments (single instance only)

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.
5. If using Isilon for Cold, configure SmartConnect Advanced PROPERLY!
6. If you are going virtual and using XtremIO, us vSphere 5.5+.
7. Traditional Benchmarks are not to be trusted…
8. Trust in joint EMC and Splunk validations.

#8 - Trust in Splunk Validation

- EMC Builds It
- Splunk Tests It
- We publish
  - ”Meets or exceeds Splunk’s reference hardware.”
- You get our commitment to your success.

Single Server

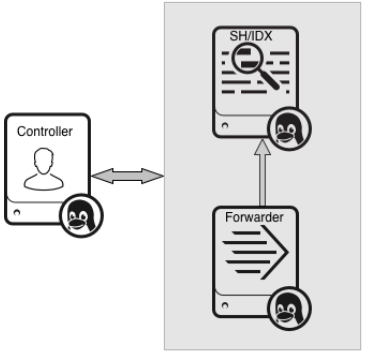

Multi Server

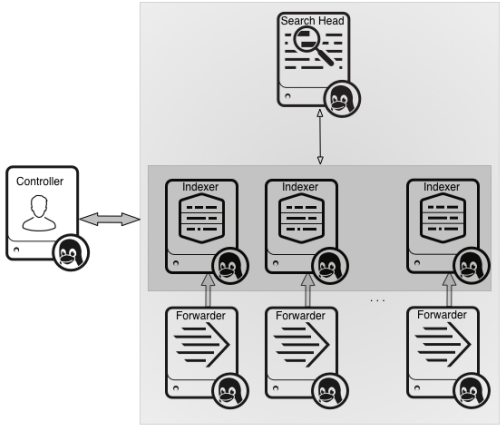

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.
5. If using Isilon for Cold, configure SmartConnect Advanced PROPERLY!
6. If you are going virtual and using XtremIO, us vSphere 5.5+.
7. Traditional Benchmarks are not to be trusted…
8. Trust in joint EMC and Splunk validations.
9. The Splunk Apps for EMC are handy, but need your creativity!

#9 – Install the Splunk Apps from EMC

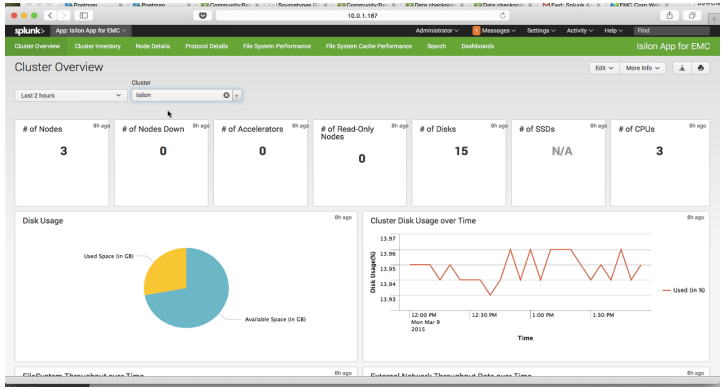

[XtremIO App]

[Isilon App]

[Isilon FS Audit App]

[VNX App]

- Go Install them!
  - Apps = The Pretty Face…
  - TAs are slick…
- Use cases make them really awesome!
  - What would you do with them?
  - Share it with us!

Cory and Jenny's Top 10

1. When in doubt, follow published EMC best practices.
2. Do the same for Splunk!
3. Slight changes in Splunk = Massive impacts on infrastructure
4. When decoupling, networks are REALLY important.
5. If using Isilon for Cold, configure SmartConnect Advanced PROPERLY!
6. If you are going virtual and using XtremIO, us vSphere 5.5+.
7. Traditional Benchmarks are not to be trusted…
8. Trust in joint EMC and Splunk validations.
9. The Splunk Apps for EMC are handy, but need your creativity!
10. Don't hesitate, call an EMC Splunk Ninja TODAY!!!

#10 – Call an EMC Splunk Ninja

- We are Splunk SE trained
- We know our kit and Splunk
- Global representation of 40+
- SME resources…use us
- Who doesn't love a Ninja?

THANK YOU

splunk>