

Finding Your Faults Before Mom... Deploying Splunk For It Troubleshooting And Capacity Planning On Large Scale Integrated Data Center Infrastructure

Wissam Ali-Ahmad

Splunk

Karthik Karupasamy

Cisco

.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.



Karthik Karupasamy

- Technical Marketing Engineer with Cisco UCS focusing on Big Data solutions and UCS Director Express for Big Data.
- Co-authored several Cisco Validated Designs with ISV and Hardware partners.
- Main focus areas: architecture, solutions, and emerging trends in big data and infrastructure in the Data Center.

About...



Wissam Ali-Ahmad



- Solution Architect, Technical Alliance Lead in Global Strategic Alliances team at Splunk.
- 15 years of technical experience in big data, security, cloud infrastructure and enterprise software.
- Prior to Splunk, held engineering leadership roles at AppSense, Infoblox, Qualys, Vernier Networks, PSS Systems and Verizon Labs.

Agenda

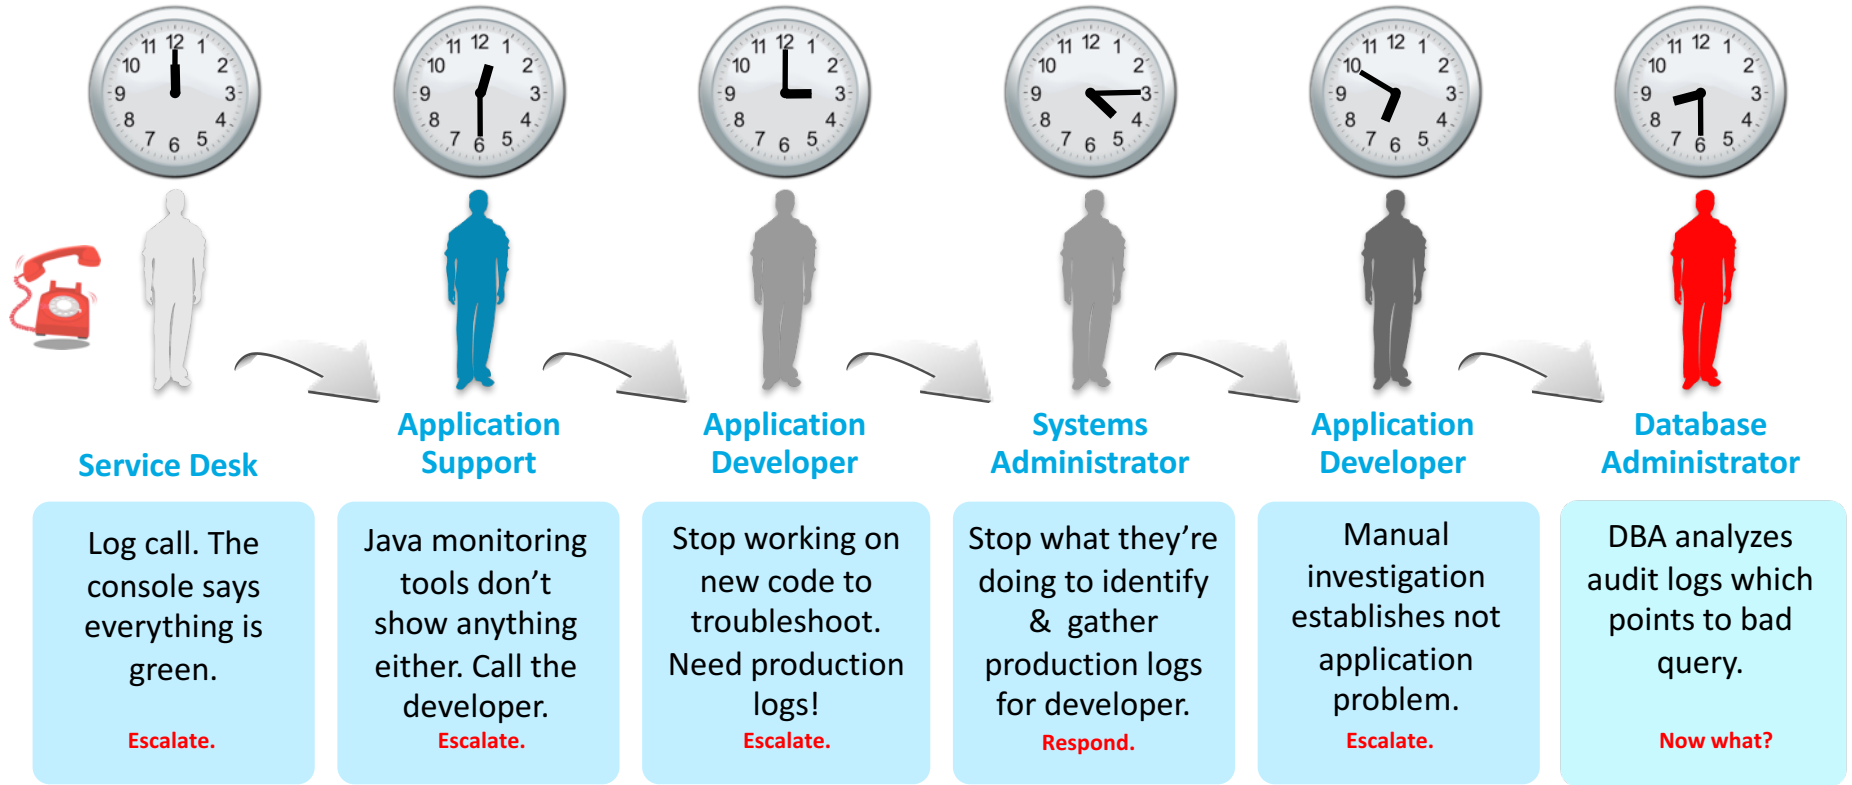
- IT Operational Analytics at enterprise scale
 - Current State
- Deployment of a large scale analytics platform
 - Challenges
 - How can Splunk and Cisco help?
- Splunk and Cisco UCS for IT Troubleshooting
 - Architecture
 - Data Collection and Mapping
 - Use Cases
 - Demo

IT Operations Analytics at Enterprise Scale



.conf2016

What It's Like in the IT Trenches



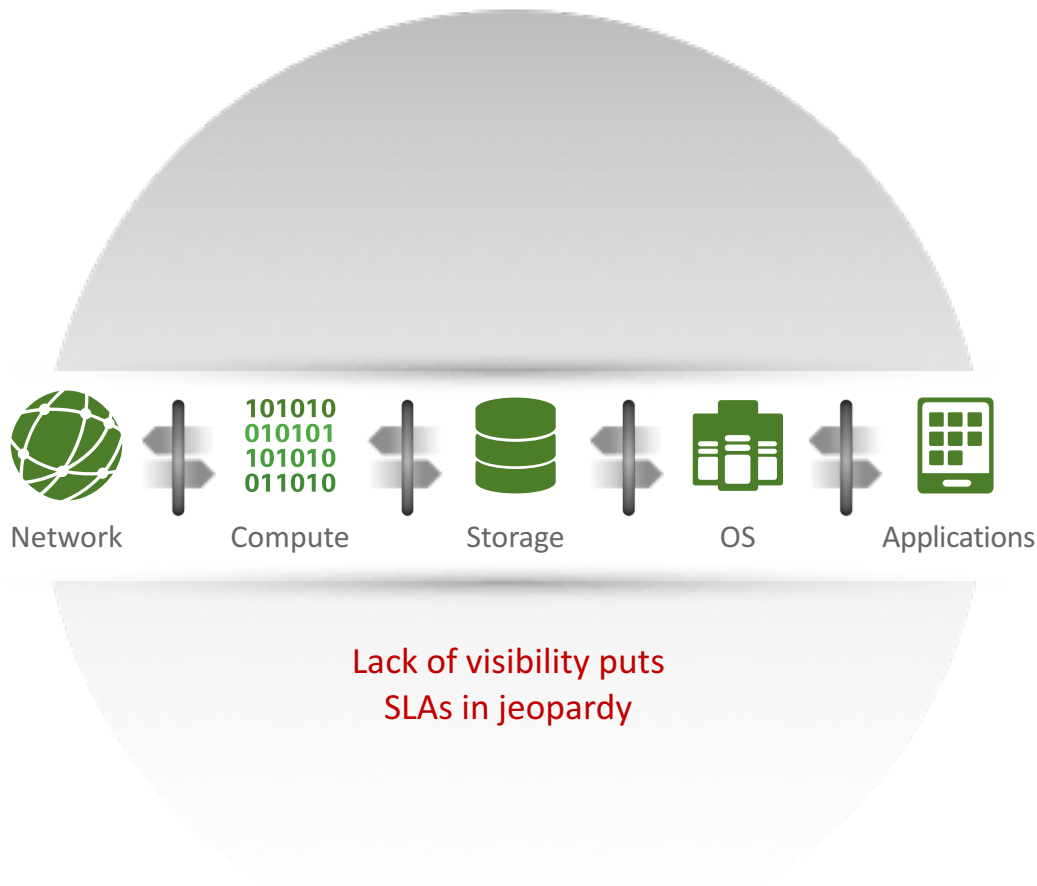
IT silos create obstacles to consistent application, infrastructure performance

Obstacles to End-to-End Visibility, High Availability/Performance

Organizational silos

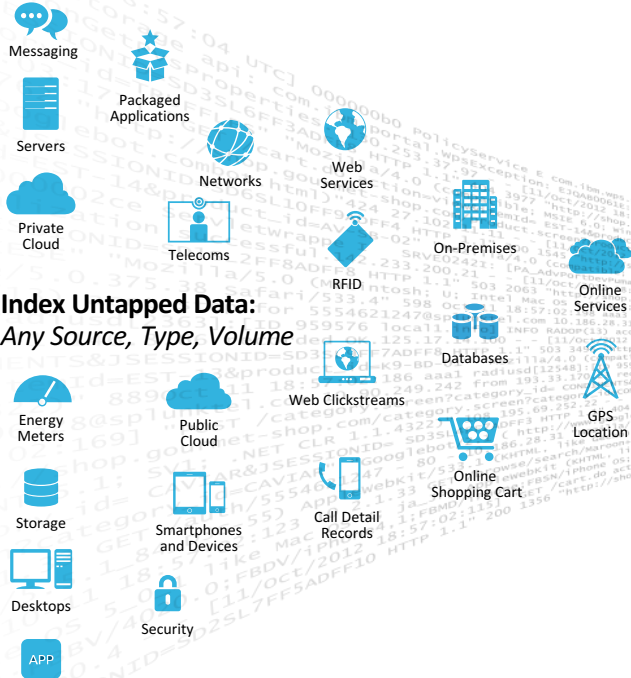
Different management and monitoring tools

Prevent IT from delivering predictive services, meeting SLAs



Required: Operational Intelligence at Scale

Data-Driven Insights To Enhance IT Performance with Splunk



Index Untapped Data:
Any Source, Type, Volume



Any Data
Any Format
Any Volume
Any Location

Real-time visibility across stack

Scalable platform

One software, many use cases

Quick root cause and issue resolution

Splunk and Cisco UCS

High Performance High Scale Infrastructure Platform for IT Operations Analytics

Analytics Software

- **Single software platform** integrates across infrastructure silos, enabling visibility to data anywhere
- **Flexibility** to identify, analyze new data sources
- **Fast time to value**
- **Comprehensive IT management functionality** to improve IT productivity



Monitoring



Incident
Mgmt



Problem
Mgmt



Capacity
Mgmt

Analytics Infrastructure



Highly
scalable



Consistent, split-
second
response times



Low TCO

Deployment of large scale analytics platform



.conf2016

Cisco Unified Computing System

A differentiated, revolutionary approach



Simplified Architecture

- Networking with fewer components
- Lower cost and easier scaling
- Fewer management touch points
- Stateless: any resource, any time
- Better TCO/ROI



Unified Management

- Faster deploy/provision
- Unification leads to reduced complexity
- Management via a single interface



Higher Performance

- Brings out the best of x86 architecture
- Optimized resource utilization for compute, networking, and management



Scale

- Ultimate Scalability
- Enhanced design capability
- Designed for the future, today

Cisco UCS: Industry-Leading Performance

110 World-record Benchmark Results

26

CPU

17

Virtualization/
Cloud

9

Database

18

Enterprise
Application

15

Enterprise
Middleware

19

HPC


6

Big
Data

Cisco UCS Benchmarks that held world record performance records as of date of publication For details, please see source document "Cisco Unified Computing System and Intel Xeon Processors: 100 World-Record Performance Results" at http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/te_32801_pb_ucs_worldrecords.pdf



UCS Enables Operational Intelligence at Scale

- Seamless Scalability Facilitates Rapid Growth
 - Ability to scale from a single instance to distributed deployment
 - Runs on the same UCS C-Series as other big data platforms
- Split Second Response Times
 - Exceptional performance for “needle-in-a-haystack” searches
 - Consistent performance as simultaneous users increase exponentially
 - Up to 6x search performance gains upgrading 6.2 to 6.3 on UCS
- Enterprise Class Analytics Infrastructure
 - Reliable, manageable server platform
 - 1/2 the datacenter footprint of reference hardware
- Simplified, Repeatable Deployments
 - Four pre-tested UCS Integrated Infrastructures, CVD



Cisco UCS Integrated Infrastructure for Big Data

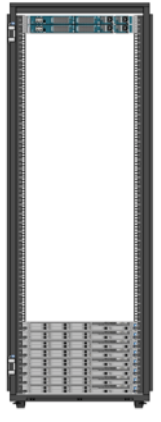
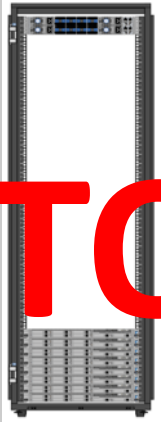




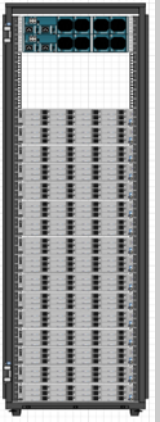

Optimized for Big Data Workloads Today and Tomorrow



- Pre-tested and pre-validated configuration
- Fabric-based infrastructure integrates computing, networking, and storage resources
- Designed for high performance and availability

Cisco UCS Integrated Infrastructure for Big Data

4th Generation of Reference Architectures and Bundles

UCS-SL-CPA4-S	UCS-SL-CPA4-H	UCS-SL-CPA4-P1	UCS-SL-CPA4-P2	UCS-SL-CPA4-P3	UCS-SL-CPA4-C1	UCS-SL-CPA4-C2	<TBD>
							
<h1 style="color: red; font-size: 100px; opacity: 0.5;">TO UPDATE</h1>							
<p>Network: 2x 6248 Servers: 8 X UCS-BD-C220M4-S1 Server Type: C220 M4 SFF CPU: 2x 2620v4 Memory: 128GB DDR4 Drives: 8 x 1.2TB 10K SAS HDD VIC: VIC 1227 RAID: 12Gps SAS, 2GB UCSD: no Cores: 128 Memory: 1024 Raw Storage: 76.8 I/O Bandwidth: 7.5 Gbytes/sec</p>	<p>Network: 2x 6332 Servers: 8 X UCS-BD-C220M4-H1 Server Type: C220 M4 SFF CPU: 2x 2680v4 Memory: 256GB DDR4 Drives: 8 x 960GB SSD VIC: VIC 1387 RAID: 12Gps SAS, 2GB UCSD: no Cores: 224 Memory: 2048 Raw Storage: 60 I/O Bandwidth: 20 Gbytes/sec</p>	<p>Network: 2x 6296 Servers: 16 X UCS-BD-C240M4-P1 Server Type: C240 M4 SFF CPU: 2x 2680v4 Memory: 256GB DDR4 OS: 2 x 240GB SSD Drives: 24 x 1.2TB 10K SAS HDD VIC: VIC 1227 RAID: 12Gps SAS, 2GB UCSD: yes Cores: 448 Memory: 4096 Raw Storage: 460.8 I/O Bandwidth: 45 Gbytes/sec</p>	<p>Network: 2x 6296 Servers: 16 X UCS-BD-C240M4-P2 Server Type: C240 M4 SFF CPU: 2x 2680v4 Memory: 256GB DDR4 OS: 2 x 240GB SSD Drives: 24 x 1.8TB 10K SAS HDD VIC: VIC 1227 RAID: 12Gps SAS, 2GB UCSD: yes Cores: 448 Memory: 4096 Raw Storage: 691.2 I/O Bandwidth: 48.75 Gbytes/sec</p>	<p>Network: 2x 6332 Servers: 16 X UCS-BD-C240M4-P3 Server Type: C240 M4 SFF CPU: 2x 2680v4 Memory: 256GB DDR4 OS: 2 x 240GB SSD Drives: 24 x 1.8TB 10K SAS HDD VIC: VIC 1387 RAID: 12Gps SAS, 2GB UCSD: yes Cores: 448 Memory: 4096 Raw Storage: 691.2 I/O Bandwidth: 48.75 Gbytes/sec</p>	<p>Network: 2x 6296 Servers: 16 X UCS-BD-C240M4-C1 Server Type: C240 M4 LFF CPU: 2x 2620v4 Memory: 128GB DDR4 OS: 2 x 240GB SSD Drives: 12 x 6TB 7.2K SAS HDD VIC: VIC 1227 RAID: 12Gps SAS, 2GB UCSD: yes Cores: 256 Memory: 2048 Raw Storage: 1152 I/O Bandwidth: 26.25 Gbytes/sec</p>	<p>Network: 2x 6296 Servers: 16 X UCS-BD-C240M4-C2 Server Type: C240 M4 LFF CPU: 2x 2620v4 Memory: 256GB DDR4 OS: 2 x 240GB SSD Drives: 12 x 8TB 7.2K SAS HDD VIC: VIC 1227 RAID: 12Gps SAS, 2GB UCSD: yes Cores: 256 Memory: 4096 Raw Storage: 1536 I/O Bandwidth: 26.25 Gbytes/sec</p>	<p>Network: 2x 6332 Servers: 9 X UCS-BD-C3220-HC1 Server Type: C3260 (2 x servers) CPU: 2 x 2680v4 Memory: 256GB DDR4 OS: 2 x 240GB SSD Drives: 424 6TB 7.2K SAS HDD VIC: VIC 1387 RAID: 12Gps SAS, 2GB UCSD: yes Cores: 504 Memory: 4608 Raw Storage: 2544 I/O Bandwidth: 57.97 Gbytes/sec</p>

Splunk Reference Architecture

Cisco Validated Design (CVD) for Splunk on UCS

From Pallet to Production

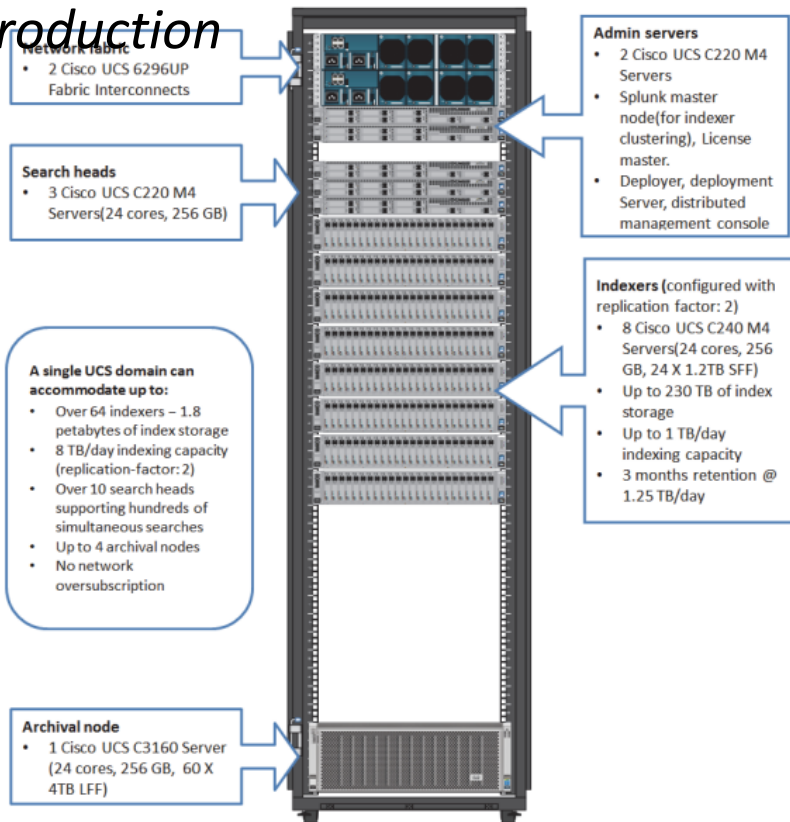


Cisco UCS Integrated Infrastructure for Big Data with Splunk Enterprise

With Cluster Mode for High Availability and Optional Data Archival

Last Updated: June 8, 2015

<http://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/big-data/splunk-enterprise.pdf>



Indexer Reference Configurations

	Configuration 1		Configuration 2		Configuration 3	
Server Spec	C240 M4 with <ul style="list-style-type: none"> RAID10: 24 1.2TB 10K SAS HDDs 		C240 M4 with <ul style="list-style-type: none"> RAID10: 24 1.8TB 10K SAS HDDs 		C240 M4 with <ul style="list-style-type: none"> RAID5: 8 x 960 GB SSDs (HOT/WARM) RAID50: 16 x 1.8TB 10K SAS HDDs (COLD) RAID50: 4 HDDs per span RAID50: 5 HDDs per span + 1 spare	
Retention/ Storage	Duration (days)	Storage (TB)	Duration (days)	Storage(TB)	Duration (days)	Storage (TB)
HOT/WARM	30	14 TB	45	21 TB	30	6
COLD	60		90		120	21
Total Retention	90	14 TB	135	21 TB	150	27

C240 M4 servers with 2 E5-2680 v4 CPUs, 256GB RAM and 24 SEF HDD or SSD drives -- 250 GB per day


Splunk on UCS sizing

Daily Indexing Capacity (per day)	Number of Indexers Required (per above configuration and retention)		
	No Replication	Replication Factor(RF)=2 Search Factor(SF)=2	RF=3 SF=3
250 GB	1	2	3
500 GB	2	4	6
1 TB	4	8	12
2 TB	8	16	24
4 TB	16	32	48
8 TB	32	64	96

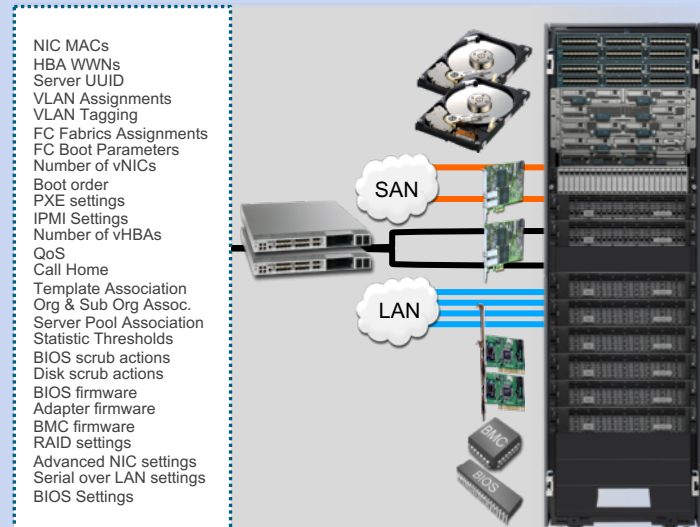
Cisco UCS Advantage - Unified Management

Robust management delivers superior programmability, scalability, and automation for Big Data deployments

Traditional

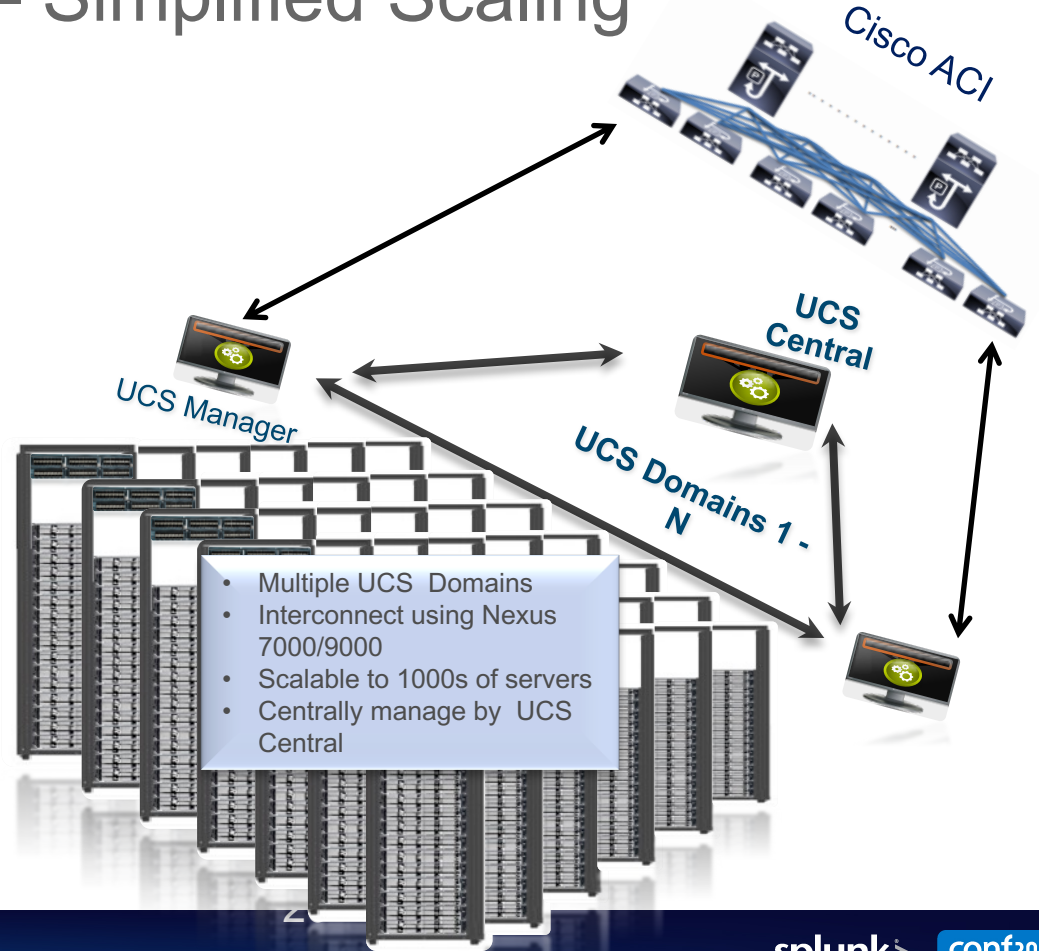



UCS Manager Service Profile



Abstraction of all configuration and identity information into a service profile speeds deployment, reduces errors, lowers costs

Cisco UCS Advantage – Simplified Scaling



UCS Director Express for Big Data

Deploy your Splunk Enterprise Cluster in hours – not in days or weeks

Features:

- Indexer clustering – customizable Replication and Search Factors
- Search Head clustering
- License Master, Deployer for SHC
- Ability to grow the Search Head, Indexer clusters.
- DMC
- Automates more than 90% of the CVD process.
- (Manual configuration required for completing the remaining 10%).

Unified Management with UCS Director Express for Big Data

Programmability, Scalability and Automation


UCSD Express



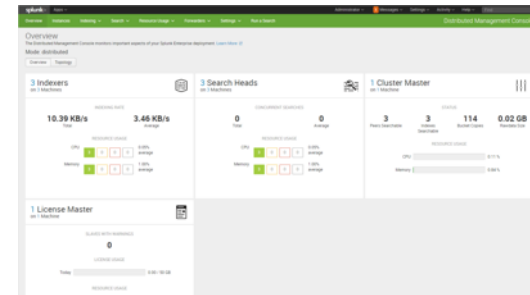

UCS 6200 Series Fabric Interconnect



UCS C240 M4 Series Rack Server



UCS C3160 Rack Server



Creating a Splunk cluster

- Cluster Name
- OS (RHEL)
- Splunk version
- UCS Manager
- Organization

Instant Splunk Cluster Creation

Big Data Account Name *
Enter Big Data Account Name with atmost 10 alphanumeric characters

UCSM Policy Name Prefix *
Enter UCSM Policy Name Prefix with atmost 5 alphanumeric characters

SSH (root) Password *

Confirm SSH Password *

Splunk Manager Password *

Confirm Splunk Manager Password *

OS Version *
Choose RHEL 6.5 for M4 Servers

Splunk Distribution Version *

Multi UCSM

UCS Manager Account *

Organization *

UCS SP Template

	Account Nam	DN	Name	Associated St	Assigned Ser
<input type="checkbox"/>	UCSM	org-root/lis-ucs	ucs	<input type="radio"/> unassociate	
<input type="checkbox"/>	UCSM	org-root/lis-bd10	bd103cloudera	<input type="radio"/> unassociate	

Optional →

Submit Close

Creating a Splunk Cluster

- Server-pools (per role)
- Map vNIC to IP-Pools.
 - Mgmt,
 - Data1(Ingest),
 - Data2(Replication)

Replication
Factor,
Search
Factor

Server
Pools

Networking

NOTE: eth0 → MGMT pool
binding shown.

- Click **Submit**

Instant Splunk Cluster Creation

PXE VLAN ID: 102

Replication Factor: 2

Search Factor: 1

Splunk Server Roles

Node Type	Node Count	Host Name Pr	SSD Boot Dri	Server Pool
Indexer	4	uIdx	true	UCSM;org-root/c
Search Head	3	uSearch	true	UCSM;org-root/c
Admin	1	uAdmin	false	UCSM;org-root/c

Total 3 Items

vNIC Name	IP Pool	MAC Address Pool	VLAN ID
eth0	Mgmt:50.1.1.1	Mgmt	101
eth1	Data1:0.0.0.0	Data1	201
eth2	Data2:0.0.0.0	Data2	202

Submit Close

PXE VLAN

Creating a Splunk Cluster -- Server Pool Selection

Server Pools

Server Pools

Hostname Prefix

Edit Splunk Server Roles Entry

Node Type: Indexer

Node Count:

Host Name Prefix:

SSD Boot Drives Available for OS

Server Pool

	ID	Server Pool	Server Pool F	Assigned	Size
<input type="checkbox"/>	UCSM;org-root/	default		0	1
<input checked="" type="checkbox"/>	UCSM;org-root/	test_indexers		0	4
<input type="checkbox"/>	UCSM;org-root/	test_searchhead		0	2
<input type="checkbox"/>	UCSM;org-root/	test_admin		0	1
<input type="checkbox"/>	UCSM;org-root/	test_forwarder		1	1
<input type="checkbox"/>	UCSM;org-root/	ucs		1	13
<input type="checkbox"/>	UCSM;org-root/	Hadoop		8	8
<input type="checkbox"/>	UCSM;org-root/	test_searchhead		0	3
<input type="checkbox"/>	UCSM;org-root/	Splunk		10	21

Total 9 items

Creating a Splunk Cluster -- VNIC configuration

- Map vNIC to IP-Pools.

NOTE: eth0 → MGMT pool binding shown.

- Click **Submit**

Edit vNIC Template Entry

vNIC Name: eth0

IP Pool: Mgmt(50.1.1.31 - 50.1.1.60)

MAC Address Pool: Mgmt (503)

VLAN ID: 101
[4048-4093],[1-3967]
(MGMT VLAN)

Submit Close

Instant Splunk Cluster Creation

PXE VLAN ID: 102
[4048-4093],[1-3967]

Replication Factor: 2

Search Factor: 1

Splunk Server Roles

Node Type	Node Count	Host Name Pri	SSD Boot Dri	Server Pool
Indexer	4	uIdx	true	UCSM;org-root/c
Search Head	3	uSearch	true	UCSM;org-root/c
Admin	1	uAdmin	false	UCSM;org-root/c

Total 3 items

vNIC Name	IP Pool	MAC Address Pool	VLAN ID
eth0	Mgmt:50.1.1.1	Mgmt	101
eth1	Data1:0.0.0.0	Data1	201
eth2	Data2:0.0.0.0	Data2	202

Submit Close

The Secret Trick is Hidden in UCS Template

- Splunk Cluster is powered by Underlying UCS HW Template
- Splunk's UCS HW Template comes with Flexible RAID Policy
- RAID Policy Supported:
 - RAID1, RAID0
 - RAID5, RAID6
 - RAID10 (default)
 - Future (RAID50, RAID60)

Splunk UCS HW Template – RAID Policy

The screenshot shows the 'Modify UCS SP Template for Big Data' interface. On the left, a sidebar lists configuration tasks, with 'Local Disk Configuration Policy' selected. The main area is divided into two sections:

- Local Disk Configuration Details:** Includes a checkbox for 'Use LVM For Disk Configuration' (unchecked) and a 'Partition Configuration' table.
- Configure Splunk RAID Policy:** A table defining RAID levels for different components.

Two blue callout bubbles highlight specific features: 'RAID Policy' points to the RAID configuration table, and 'Custom Partitions' points to the '1 with grow' entry in the partition table.

Custom Partitions

RAID Policy

Splunk UCS HW Template – Inside the RAID Policy

Edit Entry

Write Mode: Write back *

Read Mode: Read ahead *

Use Cache

Use Cache if Bad BBU

Strip Size(MB): 128 *

Disks for Hot Data

RAID Level[Hot Data]: RAID10 *

Disks Per Group: 12 *

Write Mode: Write back *

Read Mode: Read ahead *

Use Cache

Use Cache if Bad BBU

Strip Size(MB): 128 *

Data Disks for Cold Data

Same Disk As Hot Data

Data Disks for Frozen Data

Same Disk As Cold Data

Submit Close

RAID10 for HOT/WARM

Cold data on the same RAID group

Splunk UCS HW Template – Inside the RAID Policy

RAID10 for
HOT/WARM

Edit Entry

Disks for Hot Data

RAID Level[Hot Data] RAID10 *

Disks Per Group 10 *

Write Mode Write back *

Read Mode Read ahead *

Use Cache

Use Cache if Bad BBU

Strip Size(MB) 128 *

Data Disks for Cold Data

Same Disk As Hot Data

RAID Level[Cold Data] RAID5 *

Disks Per Group 4 *

Write Mode Write back *

Read Mode Read ahead *

Use Cache

Use Cache if Bad BBU

Strip Size(MB) 128 *

Submit Close

RAID5 for
COLD

Typical Big Data Deployment Challenges

- Paralysis by HW analysis
- Inconsistent configurations
- Repeatable results
- Justifiable costs/TCO/footprint
- Scalability and sustainability



Cisco UCS Delivers

- ✓ Accelerated Sales cycle/time to production
- ✓ Reduced architectural planning and calculation for the customer
- ✓ Consistent, repeatable results
- ✓ Comprehensive automated deployment
- ✓ Facilitates Splunk expansion at a reduced footprint

Towards a UCS “Appliance” Bundle*

(Splunk Enterprise Security)

Architecture (up to ~1.5TB)

- Based on CVD
 - Search Head Cluster (3 x UCS C220-M4)
 - Index Cluster (16 x UCS C240-M4)
 - License Master
 - DMC + Deployment Server
- Sizing
 - For every 100GB/day ingest per indexer, you need 300GB extra storage


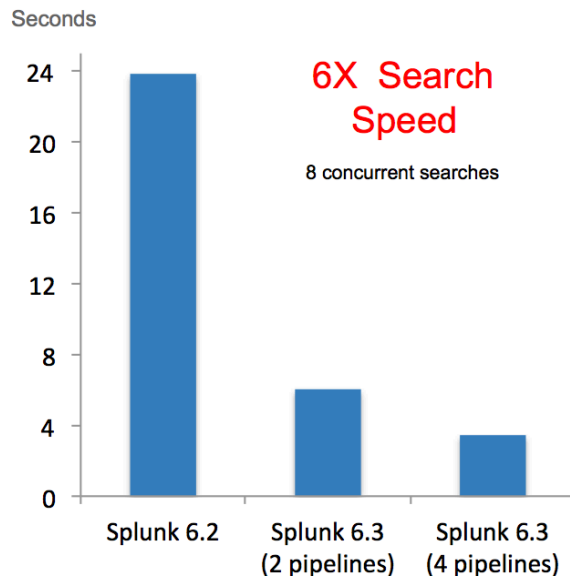
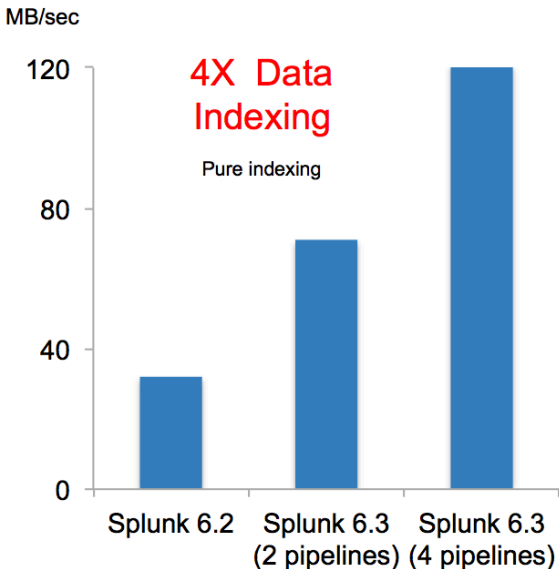
Performance Tips

- Use parallelization of data model acceleration
- Performance affected by
 - # installed/enabled Add-ons
 - # enabled correlation searches
 - Types of data sources

* OnX



Cisco UCS Benchmark Results (Splunk Enterprise 6.2 vs 6.3)



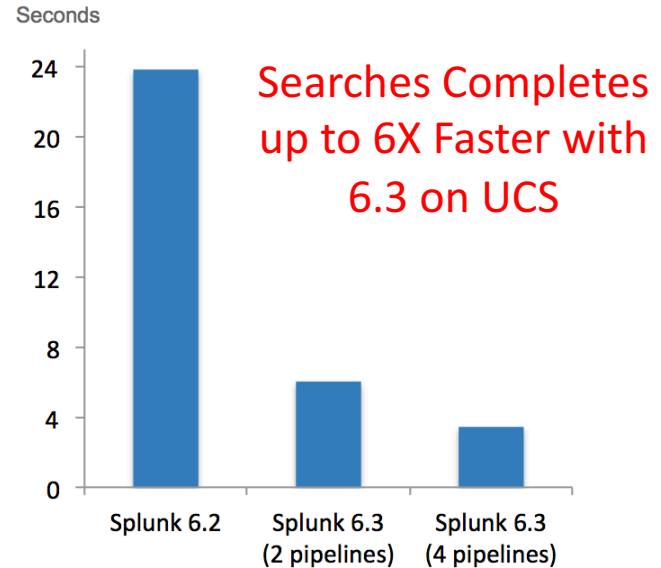
**Data for newer releases in time for .conf16*

Splunk and Cisco UCS: Better Together

- Proven at Enterprise Scale
- Exceptional Performance
- Seamless Scalability, TCO
- Splunk Supported Integration
- Faster, More Predictable Deployments



Splunk Enterprise on Cisco UCS
Performance Benchmark

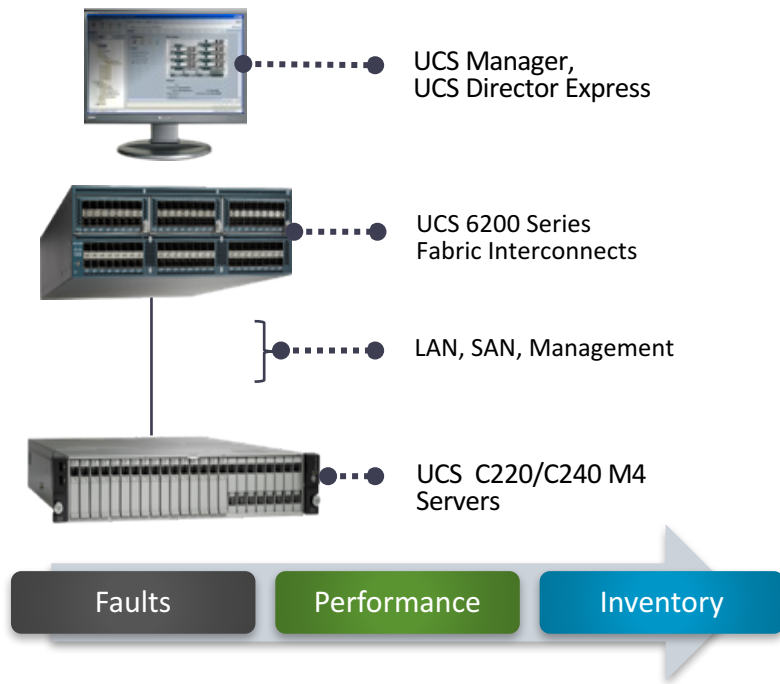


Splunk and Cisco UCS for IT Troubleshooting

.conf2016

splunk >

Splunk Add-on for Cisco UCS

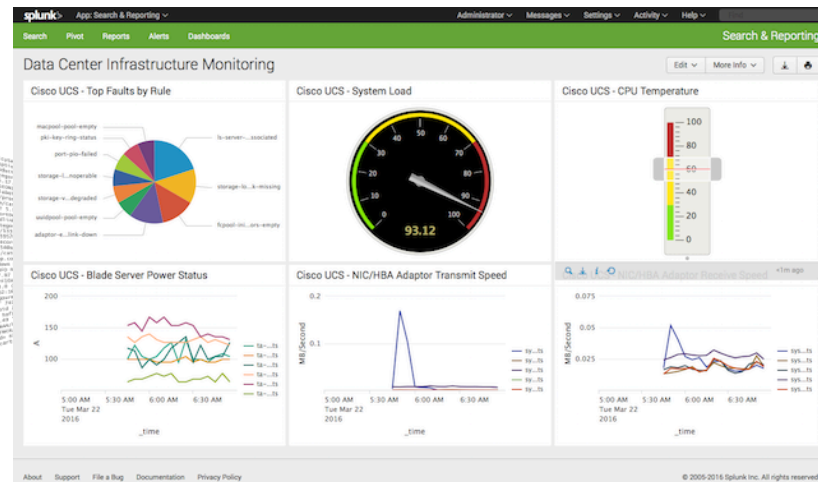


- Collects, normalizes and analyzes events from Cisco UCS Managers
- Used for troubleshooting, monitoring and capacity planning across the data center infrastructure stack
- Scales to the largest environments
- Certified for UCS platform (Cisco IVT)

* IVT: Interoperability Validation Testing

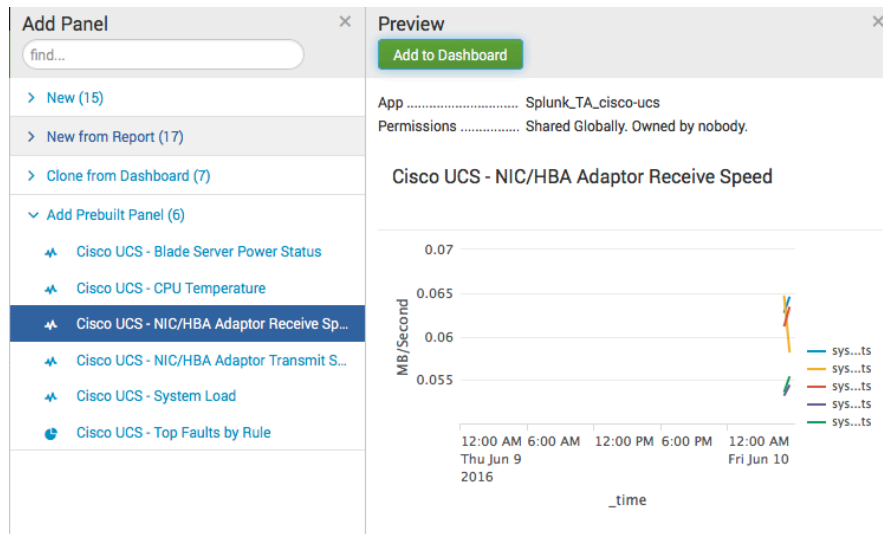


IT Troubleshooting of Application Stack with Cisco UCS



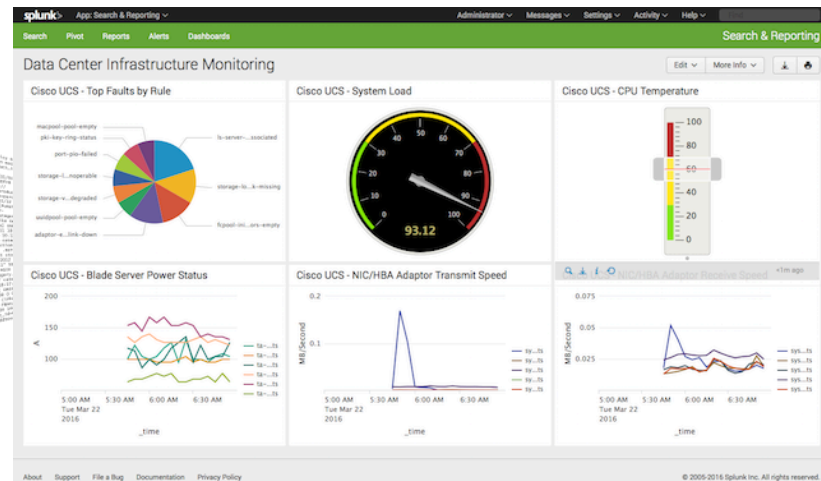
Analyze UCS Data in few minutes!

1. Download Splunk Add-on for <https://splunkbase.splunk.com/app/2731/>
2. Configure UCS Manager connectivity
3. Configure data collection task(s) for pre-defined event templates
4. Events should start showing up within 5 minutes - verify
 - `sourcetype=cisco:ucs source=*fault*`
5. Create a new dashboard
6. Add one or more pre-built panels



DEMO


IT Troubleshooting of Application Stack with Cisco UCS




Environment Health

Edit More Info Download Print


New Faults & Errors




System Load




Chassis Temperature



Network Transmission Performance

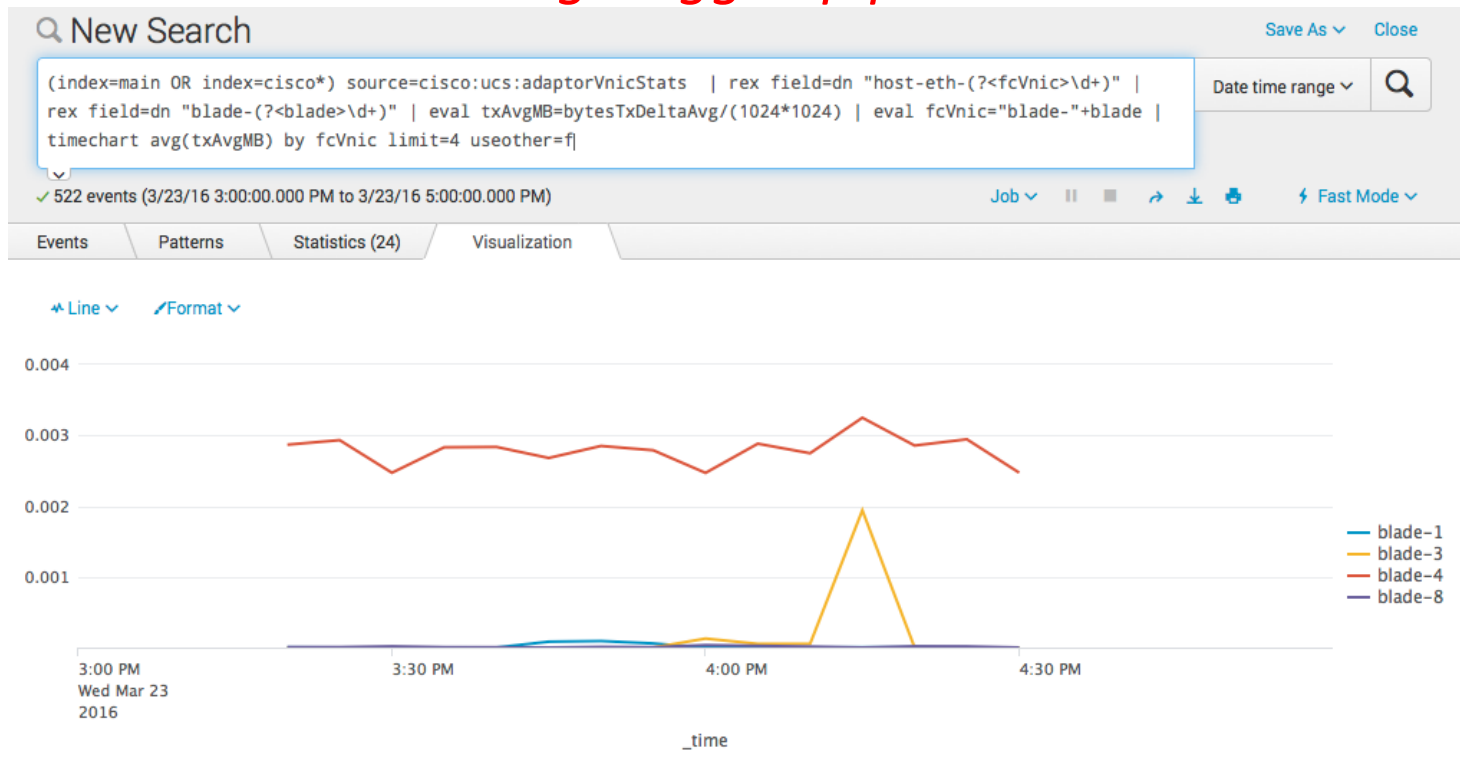


Network Reception Performance



Network Capacity & Monitoring

Finding clogged pipes...



Next Steps

Related breakout sessions and activities...


- Splunk and Cisco Architecture
<http://blogs.cisco.com/tag/splunk>
<http://www.cisco.com/go/bigdata>

- Troubleshoot your data center today

[Download Splunk for Free](#)

[Install the Cisco Add-ons](#)

See your data shine!



THANK YOU

.conf2016