goto; conference aarhus

SPLUNK IN OPERATIONS

Karsten Thygesen

CTO, Netic A/S

Who is Karsten?

- CTO, Netic A/S
- Masters, CS from Aalborg University
- In operations for 25+ years
- Splunk deployment Architect
- Evangelist of new technologies

Netic A/S

- Netic
  - Established in 2002
  - Private funded
  - Located in Aalborg
  - 16 people
  - Multiple datacenters

- Business Areas
  - Hosting
  - **Operations, Application Management, ITILv3, 24x7**
  - Consultancy and infrastructure, Stingray, VMWare etc
  - SW Development, hotspot, network provisioning etc
  - **Splunk license and consultancy**

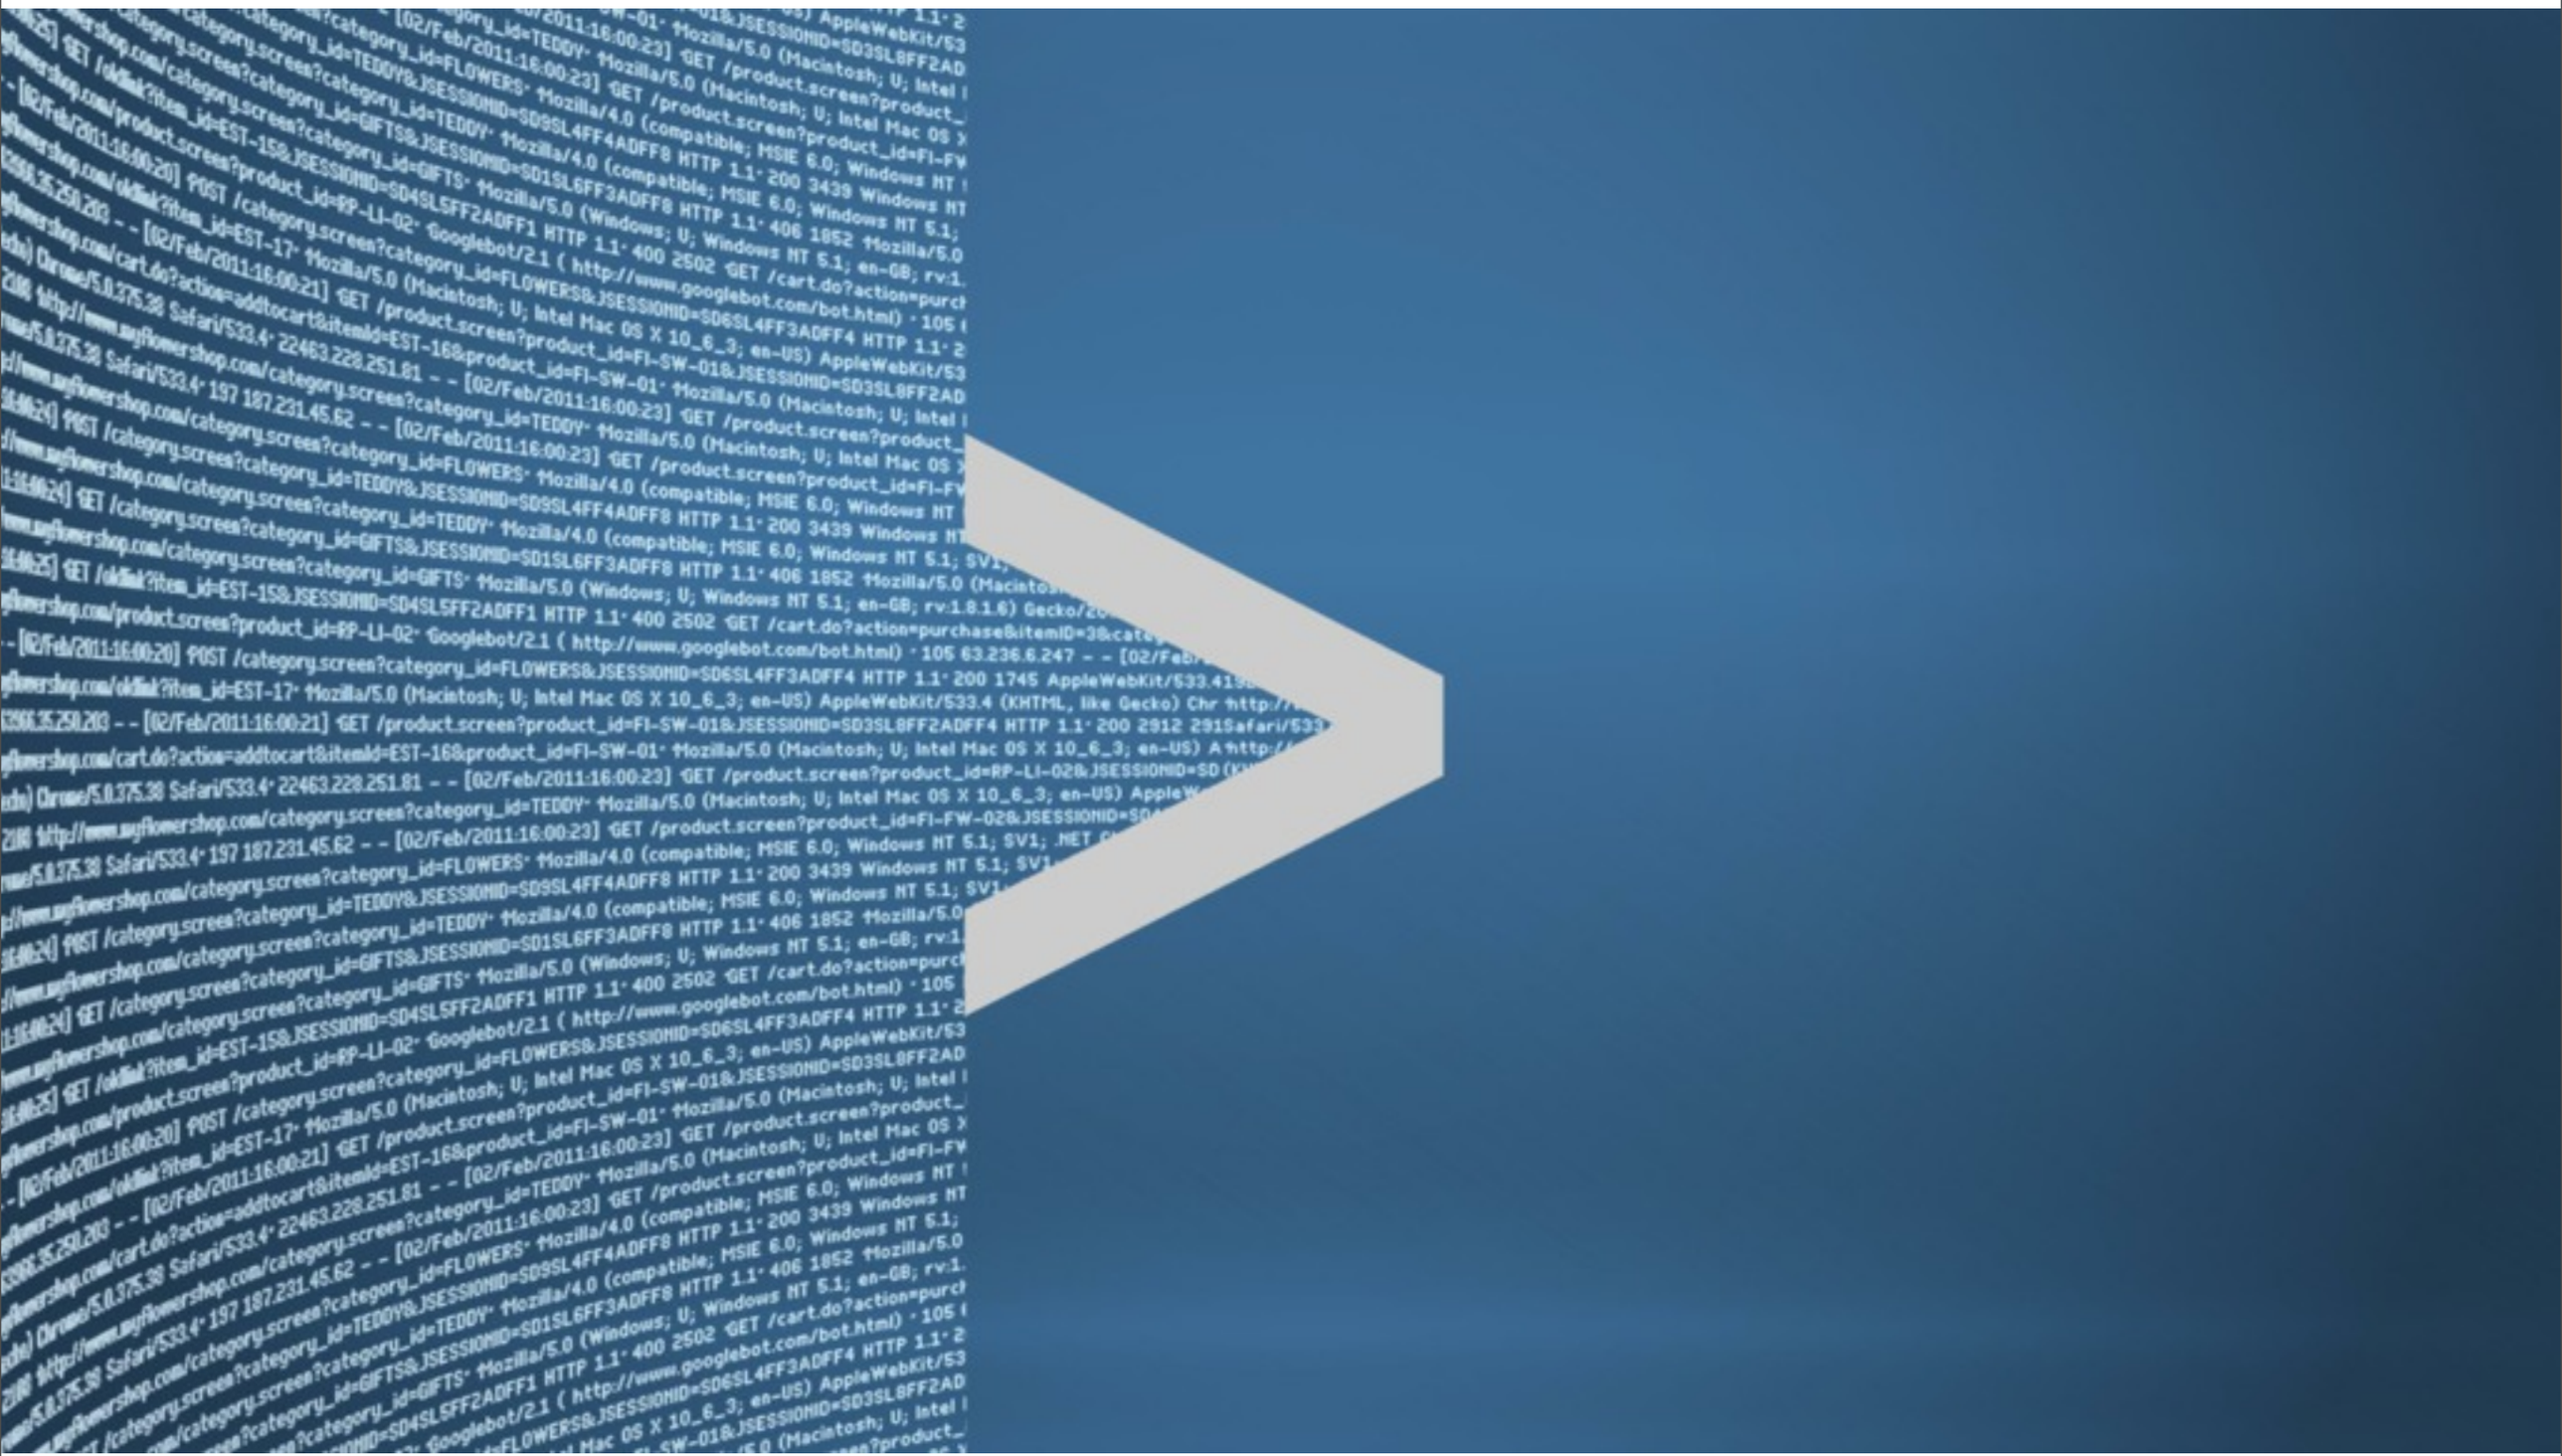

netic

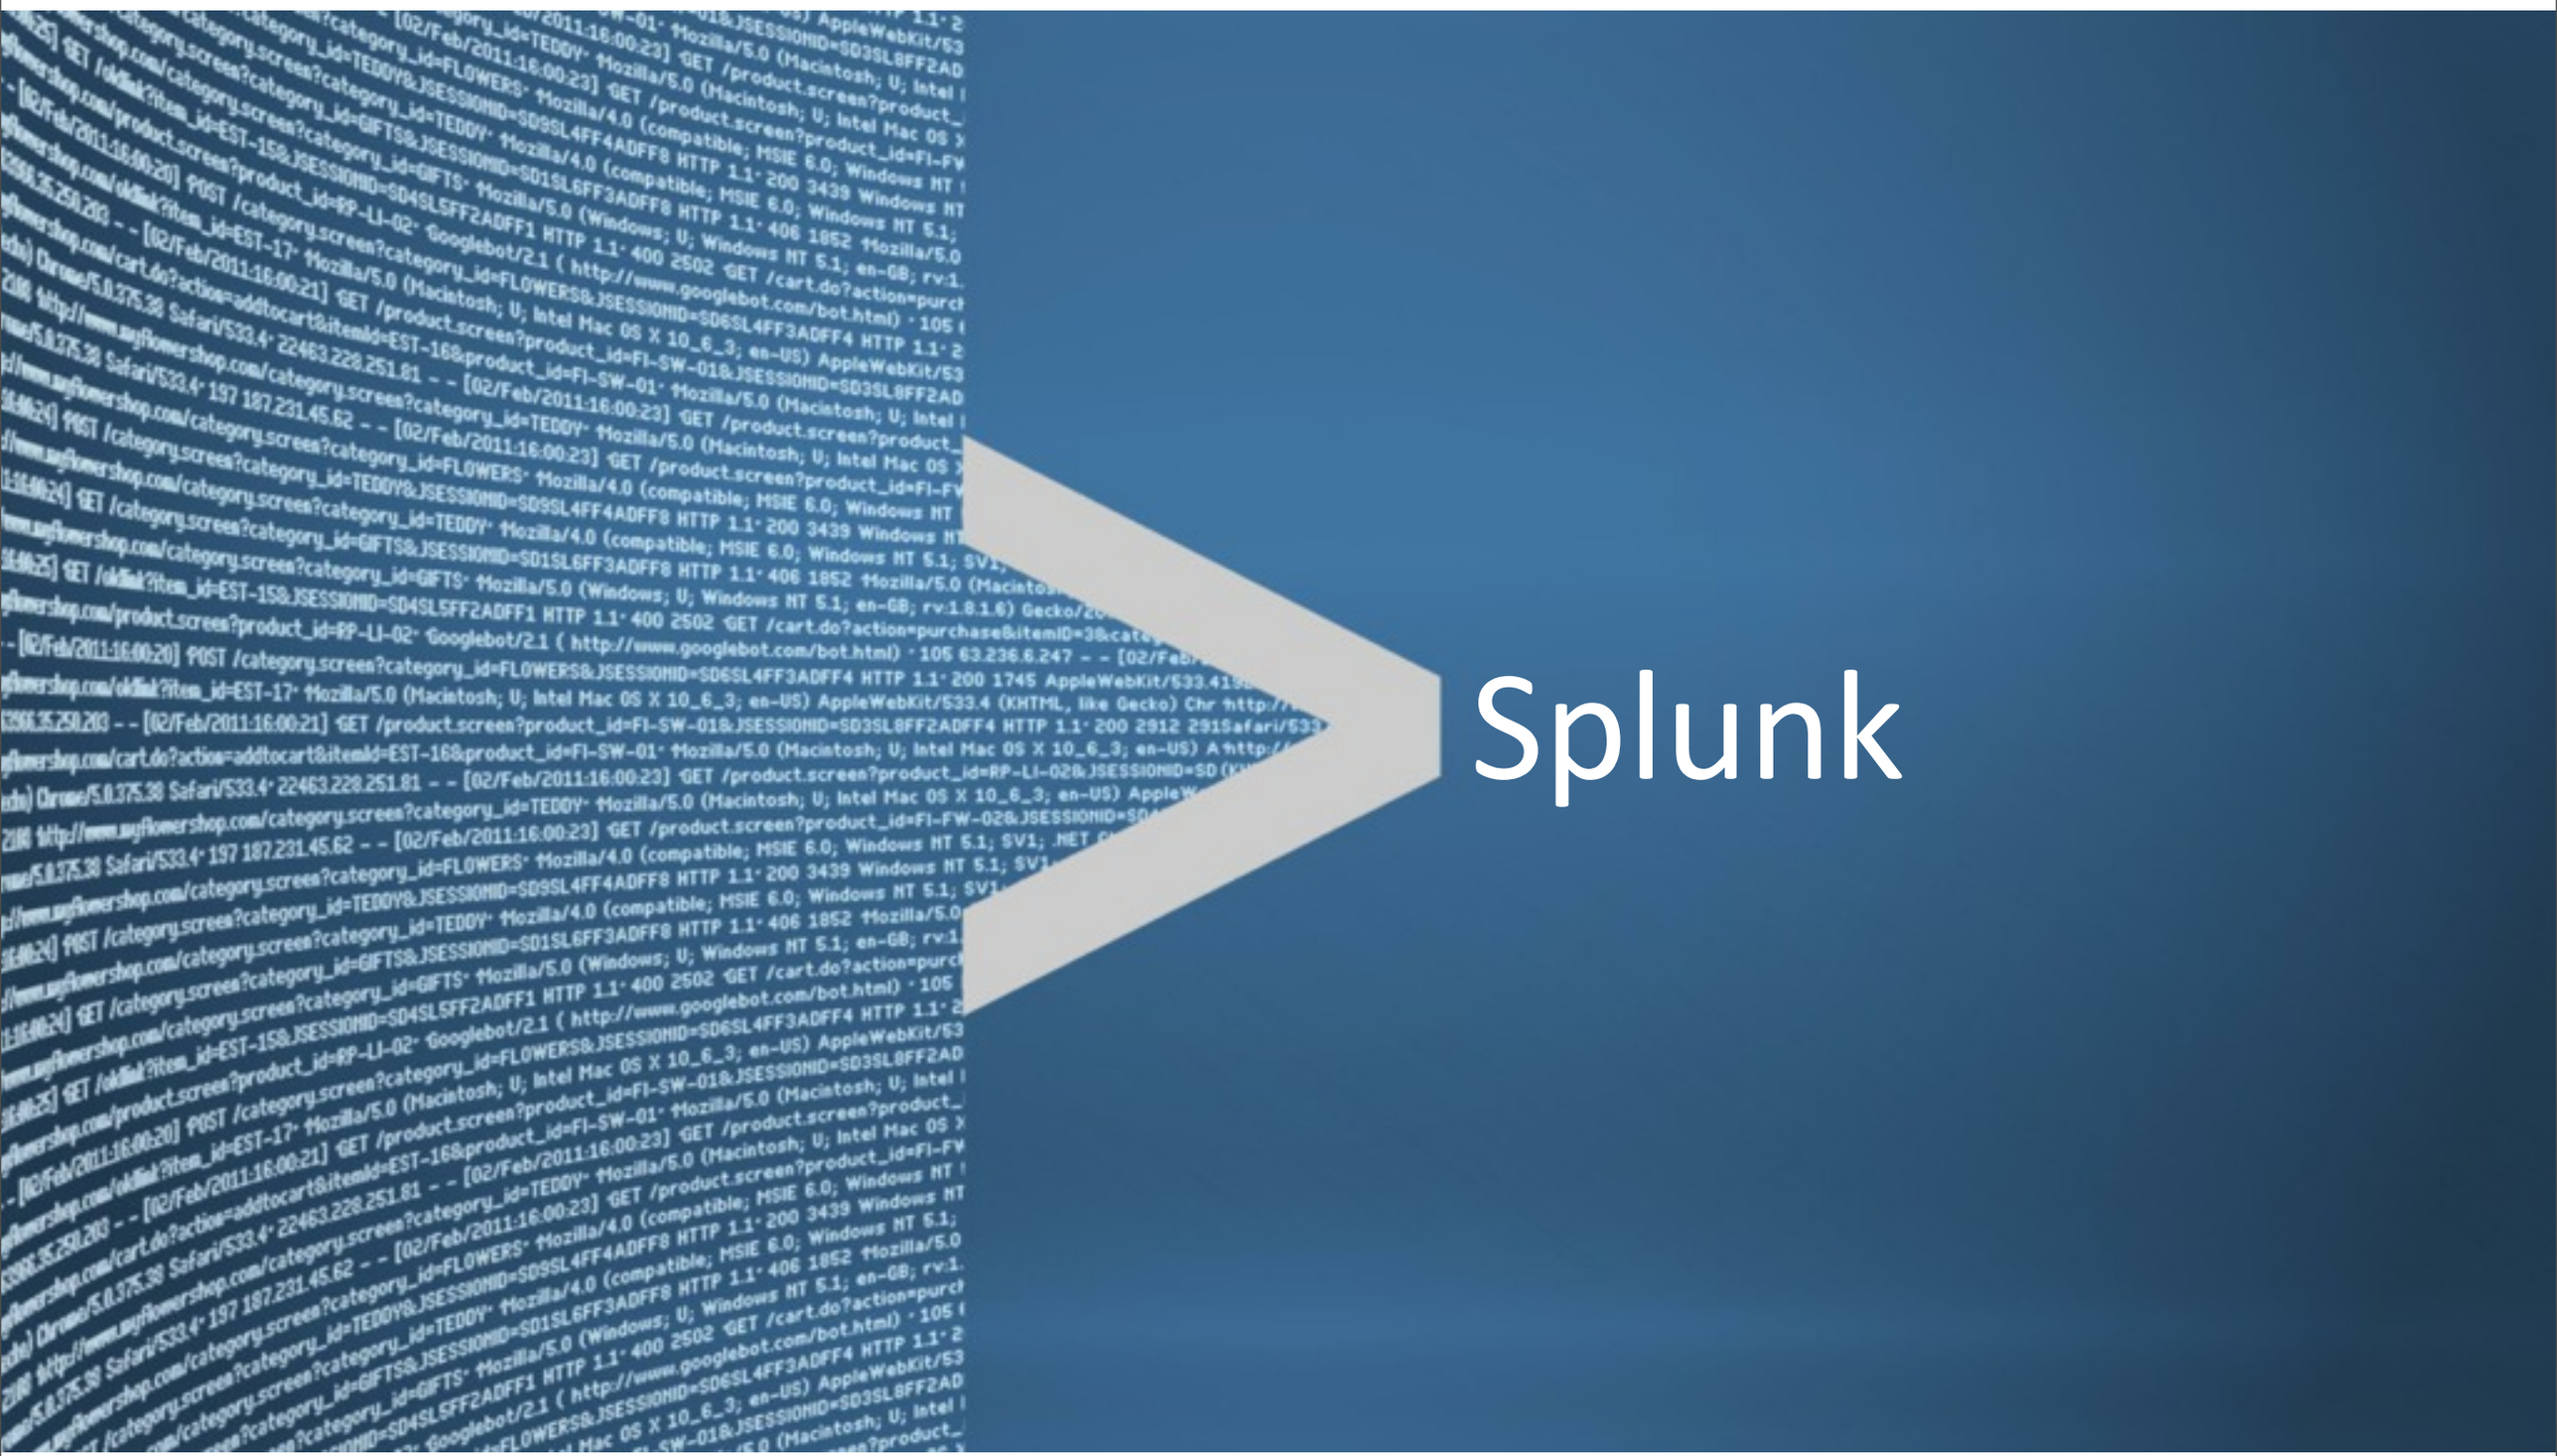

“Two-thirds of all IT spending is just to sustain the business, not to grow or transform the business”

“Two-thirds of all IT spending is just to sustain the business, not to grow or transform the business”

Gartner

Most Enterprise Data is Machine-generated

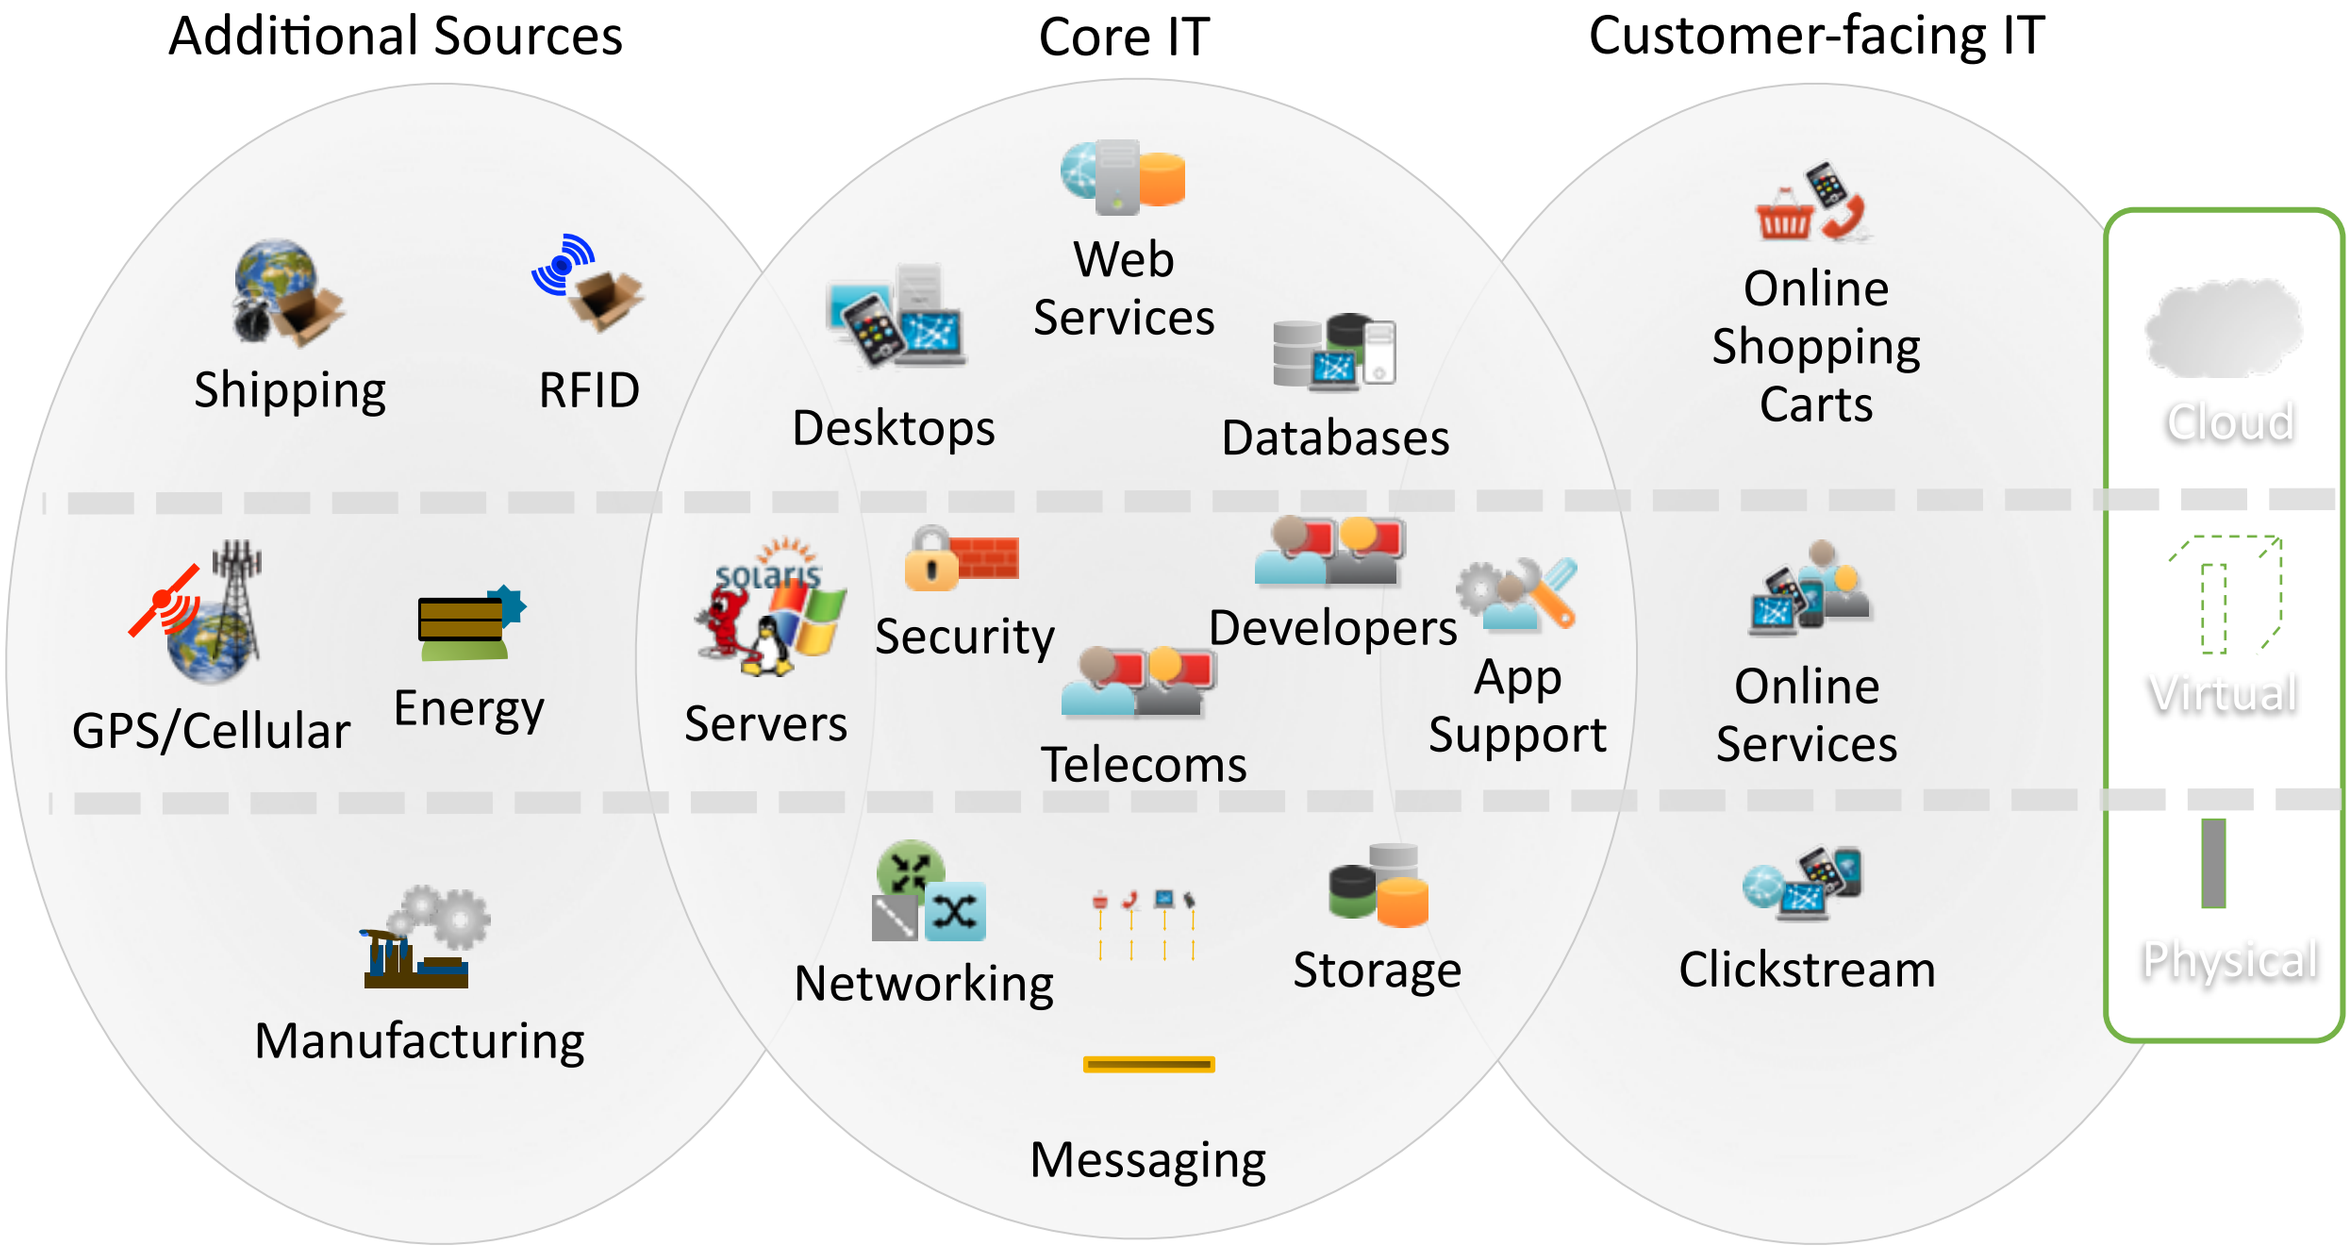

Most Enterprise Data is Machine-generated

Additional Sources

Core IT

Customer-facing IT

Shipping

GPS/Cellular

Manufacturing

Web Services

Messaging

Online Shopping Carts

Clickstream

Cloud

Virtual

Physical

- Fastest growing, most complex segment of digital data
- Generated continuously by websites, communications, networking and complex IT infrastructures
- Contains categorical record of all activity and behavior
- **Value from data largely untapped – extremely difficult to process and analyze by traditional methods and in a timely manner**

Splunk Collects and Indexes Any Machine Data

Customer Facing Data

- Click-stream data
- Shopping cart data
- Online transaction data

Outside the Datacenter

- Manufacturing, logistics…
- CDRs & IPDRs
- Power consumption
- RFID data
- GPS data

Windows

- Registry
- Event logs
- File system
- sysinternals

Linux/Unix

- Configurations
- syslog
- File system
- ps, iostat, top

Virtualization & Cloud

- Hypervisor
- Guest OS, Apps
- Cloud

Applications

- Web logs
- Log4J, JMS, JMX
- .NET events
- Code and scripts

Databases

- Configurations
- Audit/query logs
- Tables
- Schemas

Networking

- Configurations
- syslog
- SNMP
- netflow

Splunk Collects and Indexes Any Machine Data

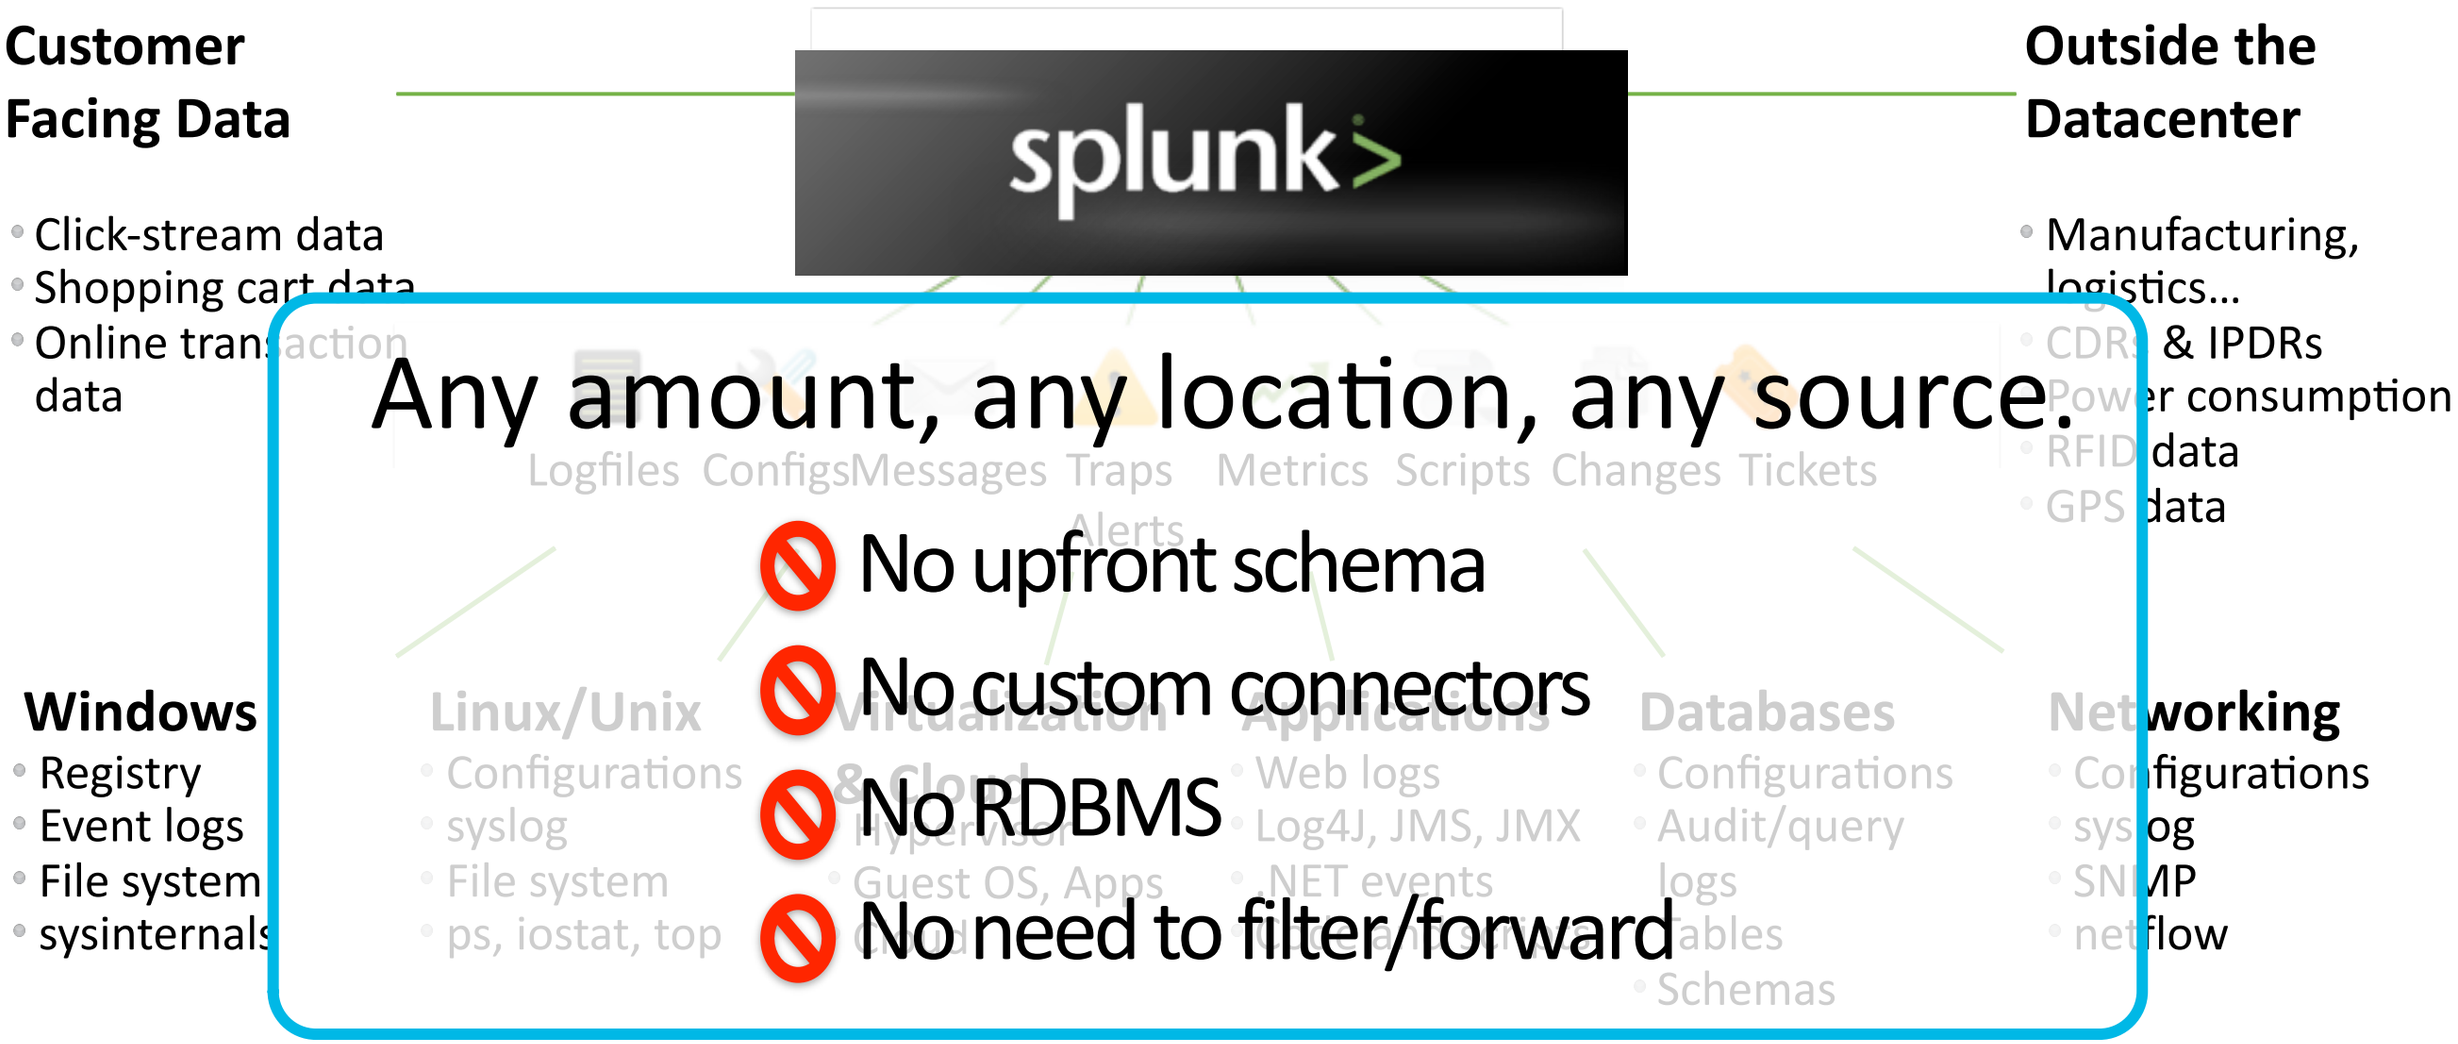

A Single Platform for Operational Intelligence

Single Data Store Single UI Across Use Cases

Three Primary Capabilities

Search / Navigation

- Data drilldown
- “Needle in a haystack”
- Root cause analysis / troubleshooting
- Incident investigations

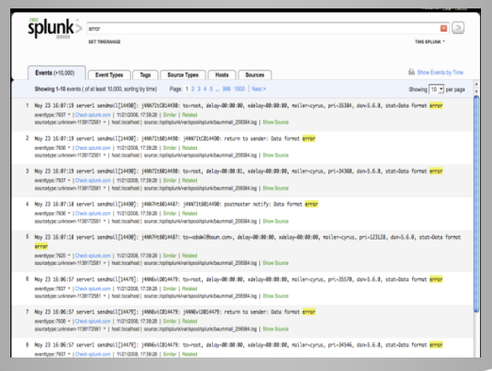

Real-time Visibility

- Live dashboards
- Event correlation
- Monitoring and alerting
- Performance issues
- Transaction levels
- SLA tracking

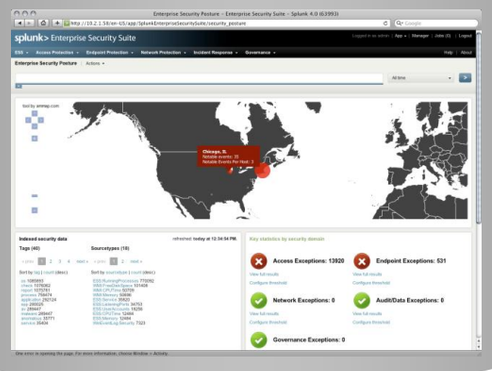

Historical Analytics

- Baseline and thresholds
- Trending
- Operational insights
- Historical patterns
- Compliance reports

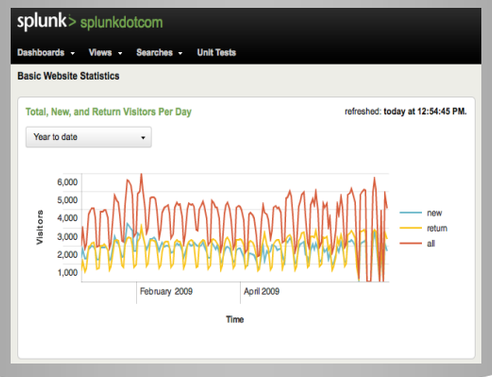

Universal Real-time Indexing

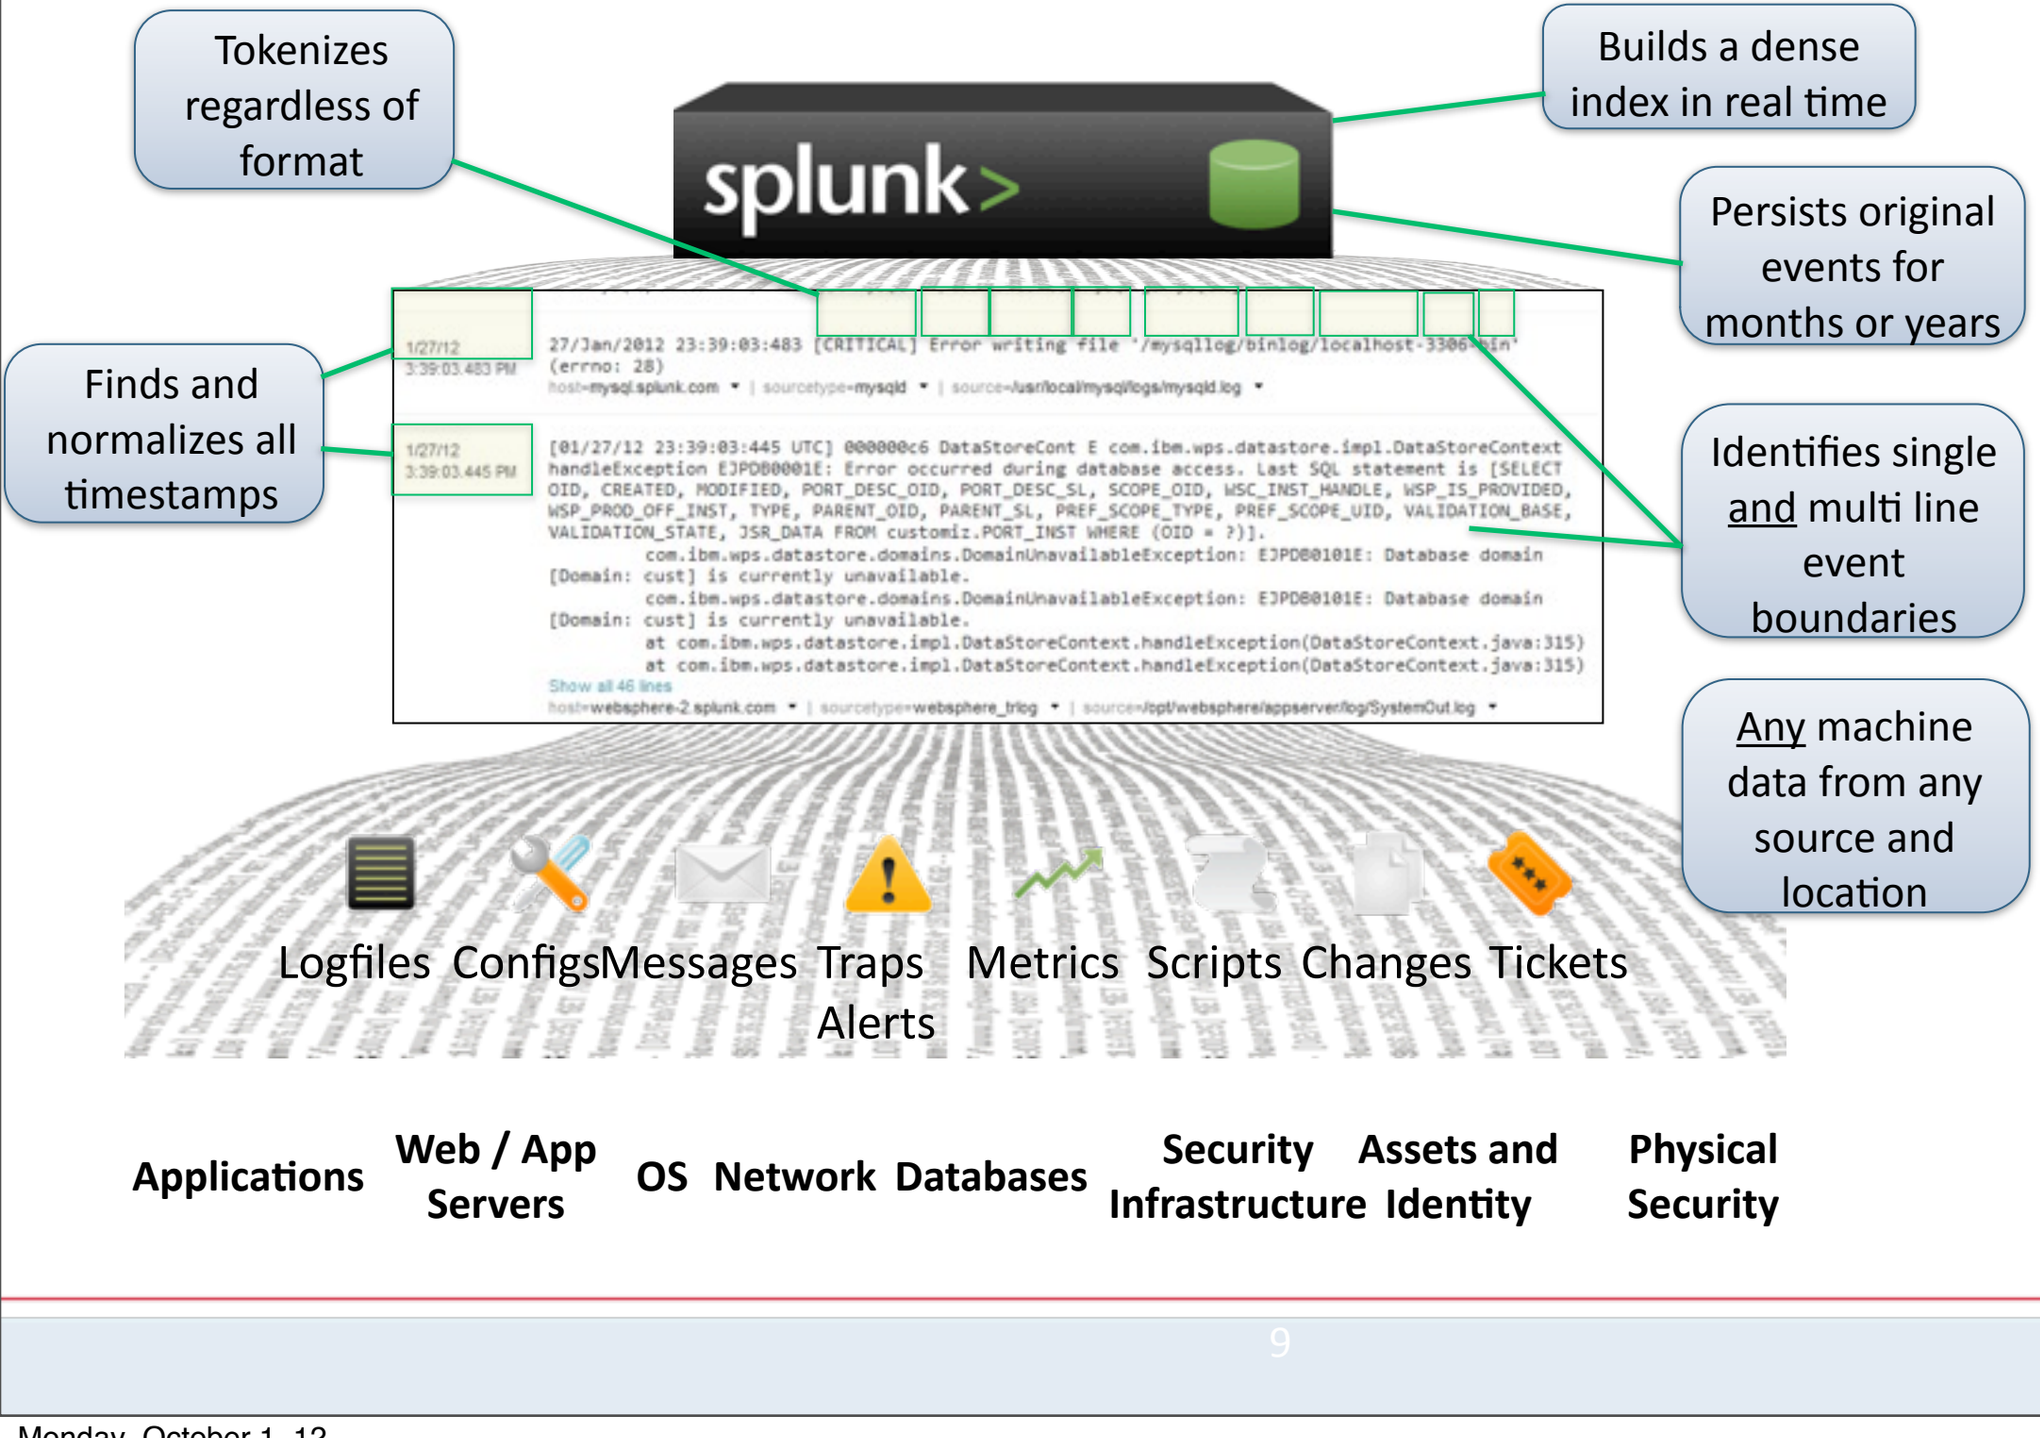

- ✓ Search and ***report*** on <u>any</u> data right away
- ✓ Accurate statistics and time correlation across <u>all</u> sources, even multiline

Resilient Search-Time Knowledge

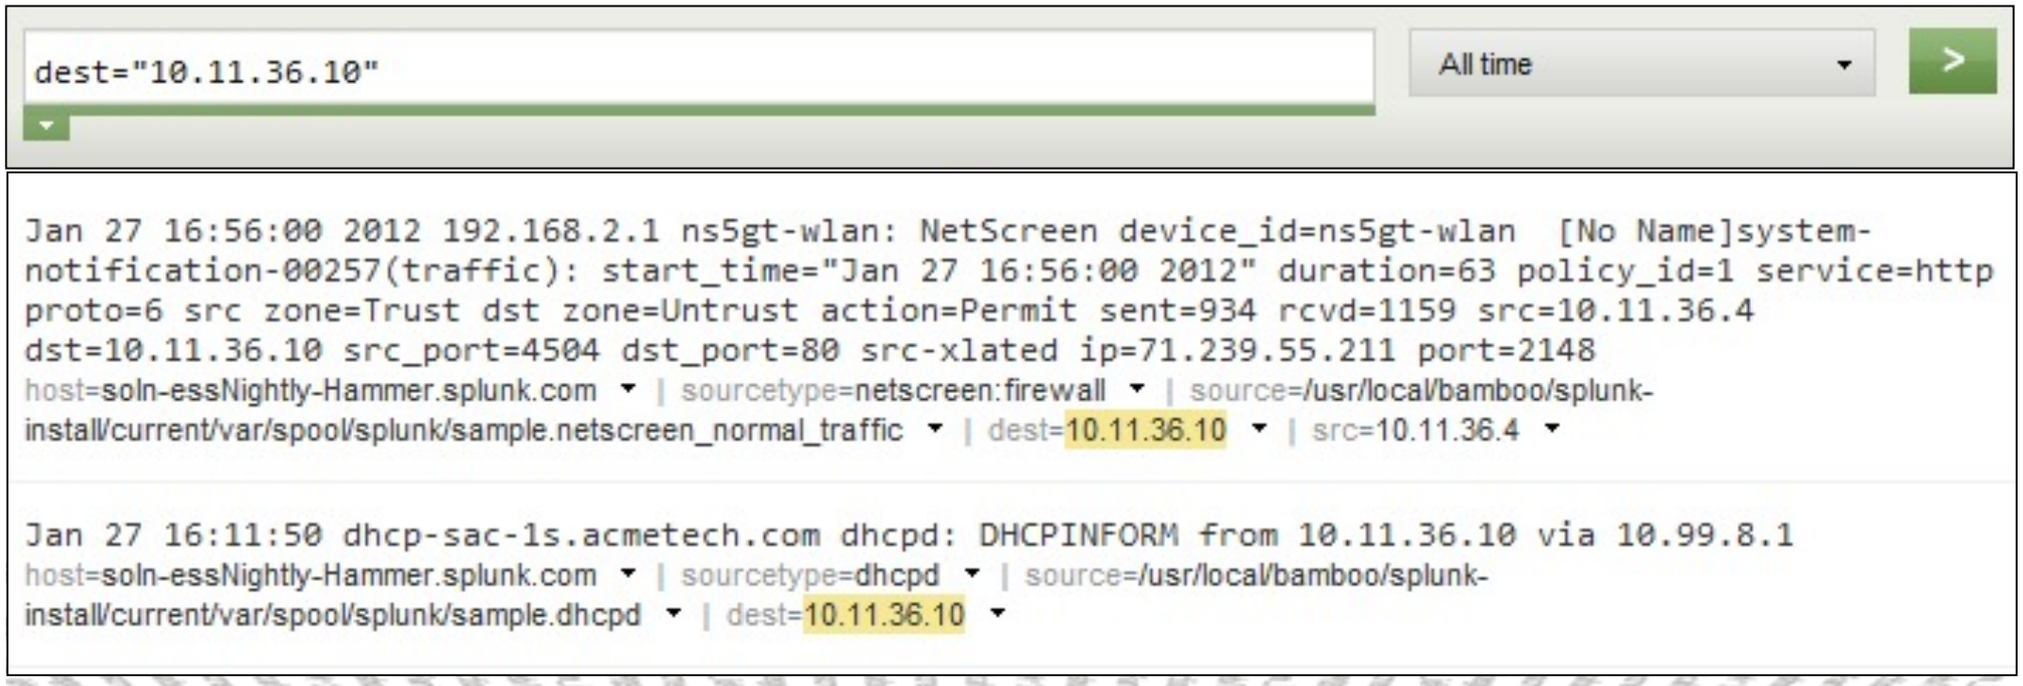

- ✓ Report on any dimension in your data
- ✓ Not limited to pre-recognized formats or a fixed schema
- ✓ Normalize only what you need to when you need to

Search-time Knowledge = On-the-fly Schema

Discovered | **User defined** | **App defined**

Fields | Tags | Lookups | Types

Massive Linear Scalability to Tens of TBs/Day

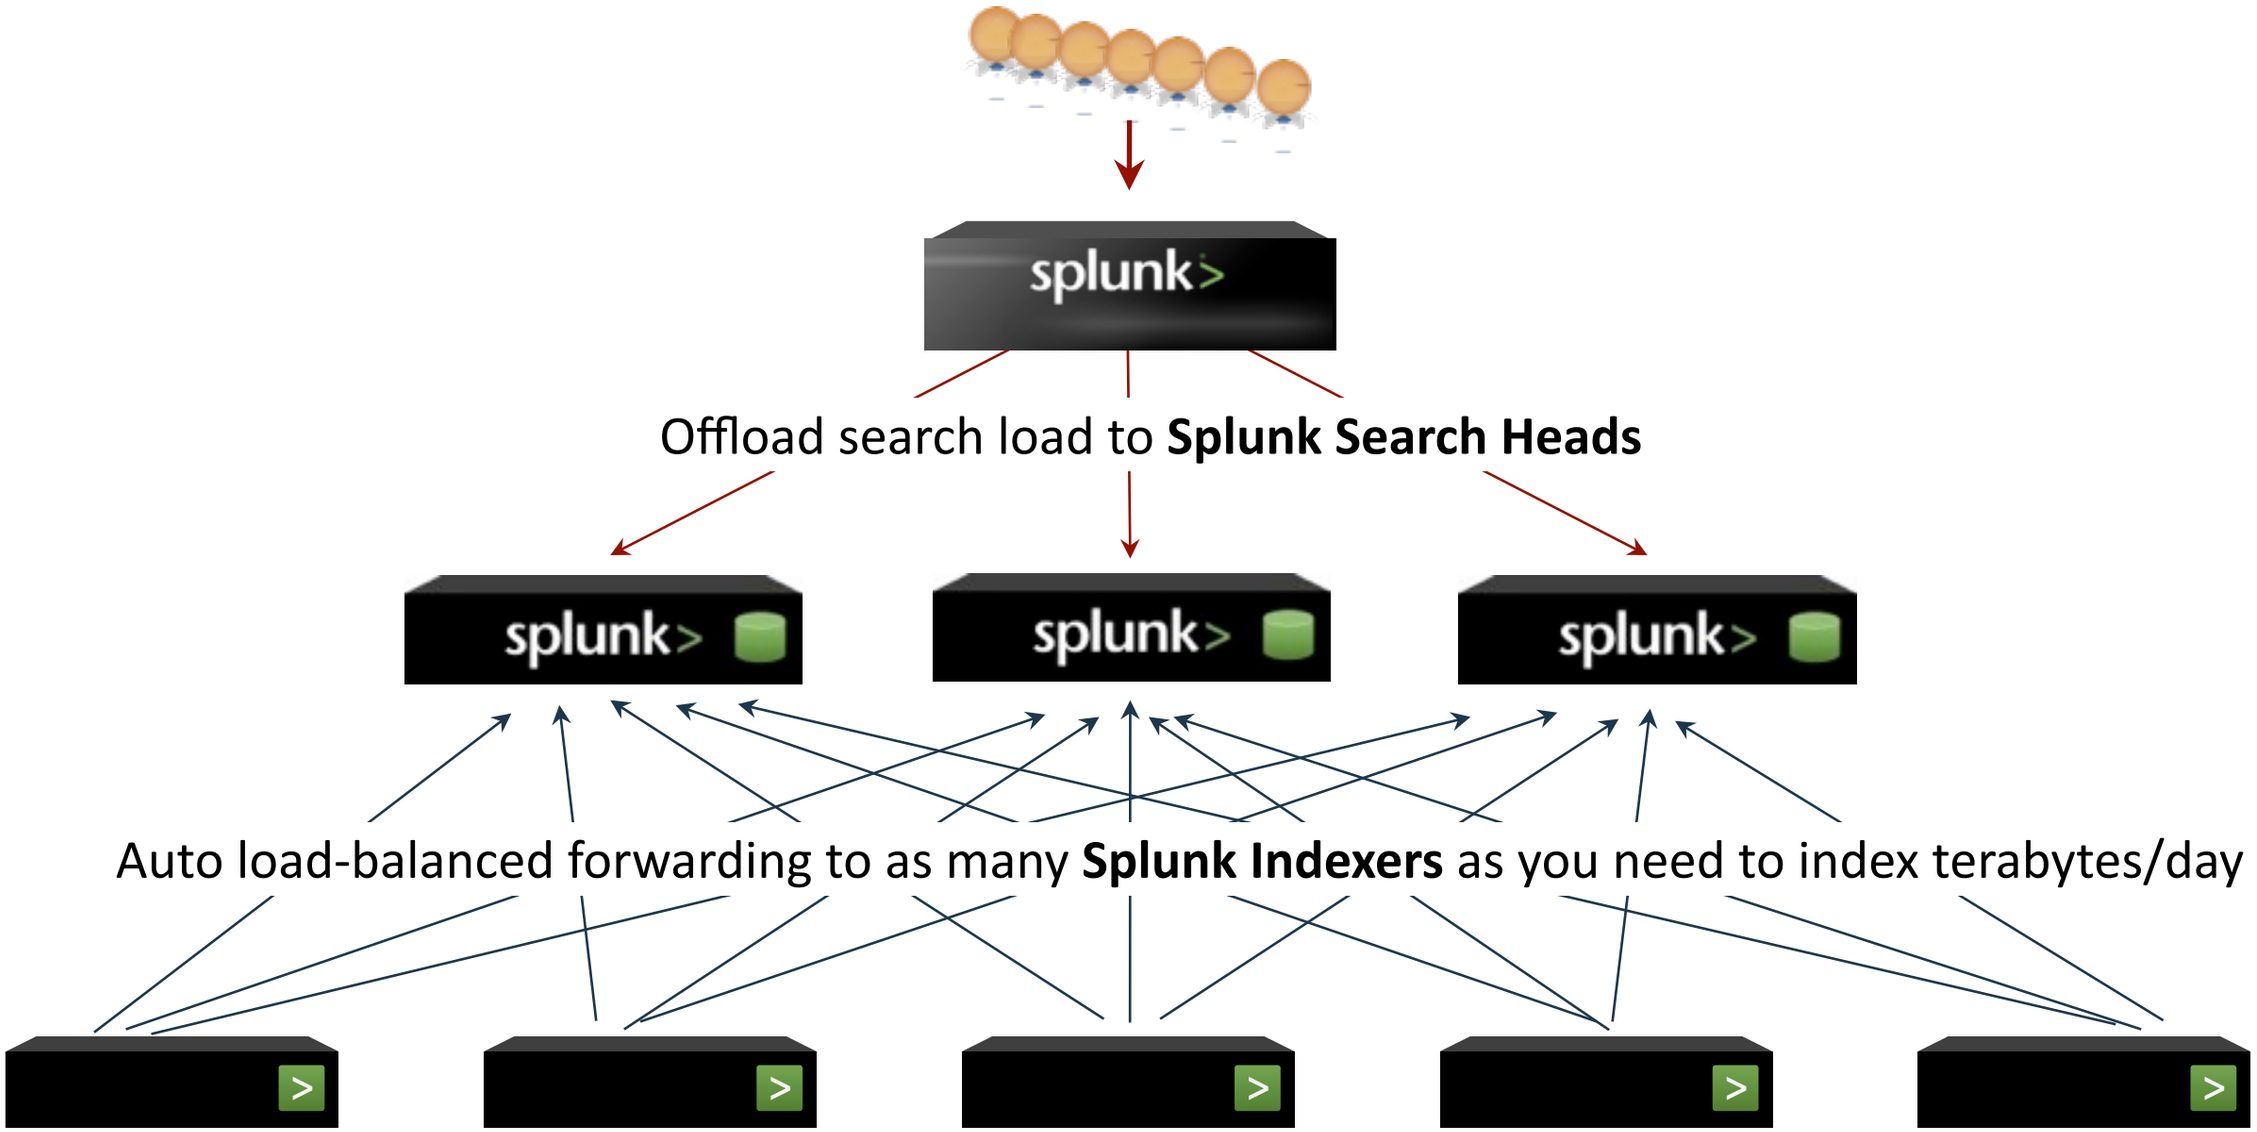

Splunk Apps Let You Do More

Community

Technology Partners

Developers

Splunk Built

Weather	BigFix	Sendmail	PDF Report Server	F5	Radio Stations	WebSphere	XenDesktop NetScaler	Multicast	MS Exchange
Ruby on Rails	Google Maps	Whois lookup	PCI Compliance	Puppet Conf. Mgt	Python Mail	NetFlow	Audible Alerts	Stock Quote	FISMA Monitoring
Twitter	Windows	Nagios	Unix and Linux	Sourcefire	Splunk Monitoring	SNORT	FireEye Malware	POST/GET Rqsts	Citrix NetScaler
Security	Javamail	BlueCoat ProxySG	Solera DeepSee	IMAP	YouTube	Encrypt/Decrypt	Enterprise Security	AS/400 - iSeries	Transaction Profiling
Security	SCOM	TCP/UDP Sending	IronPort WSA	RSS Input	JMS receiver	Geo Location	VMware	Fin. Inf. eXchange	Splunk Mobile

DEVELOPER FRAMEWORK

splunk>

Splunk Apps Let You Do More

Community

Technology Partners

Developers

Splunk Built

300 Apps and growing

DEVELOPER FRAMEWORK

splunk>

Splunk for Exchange

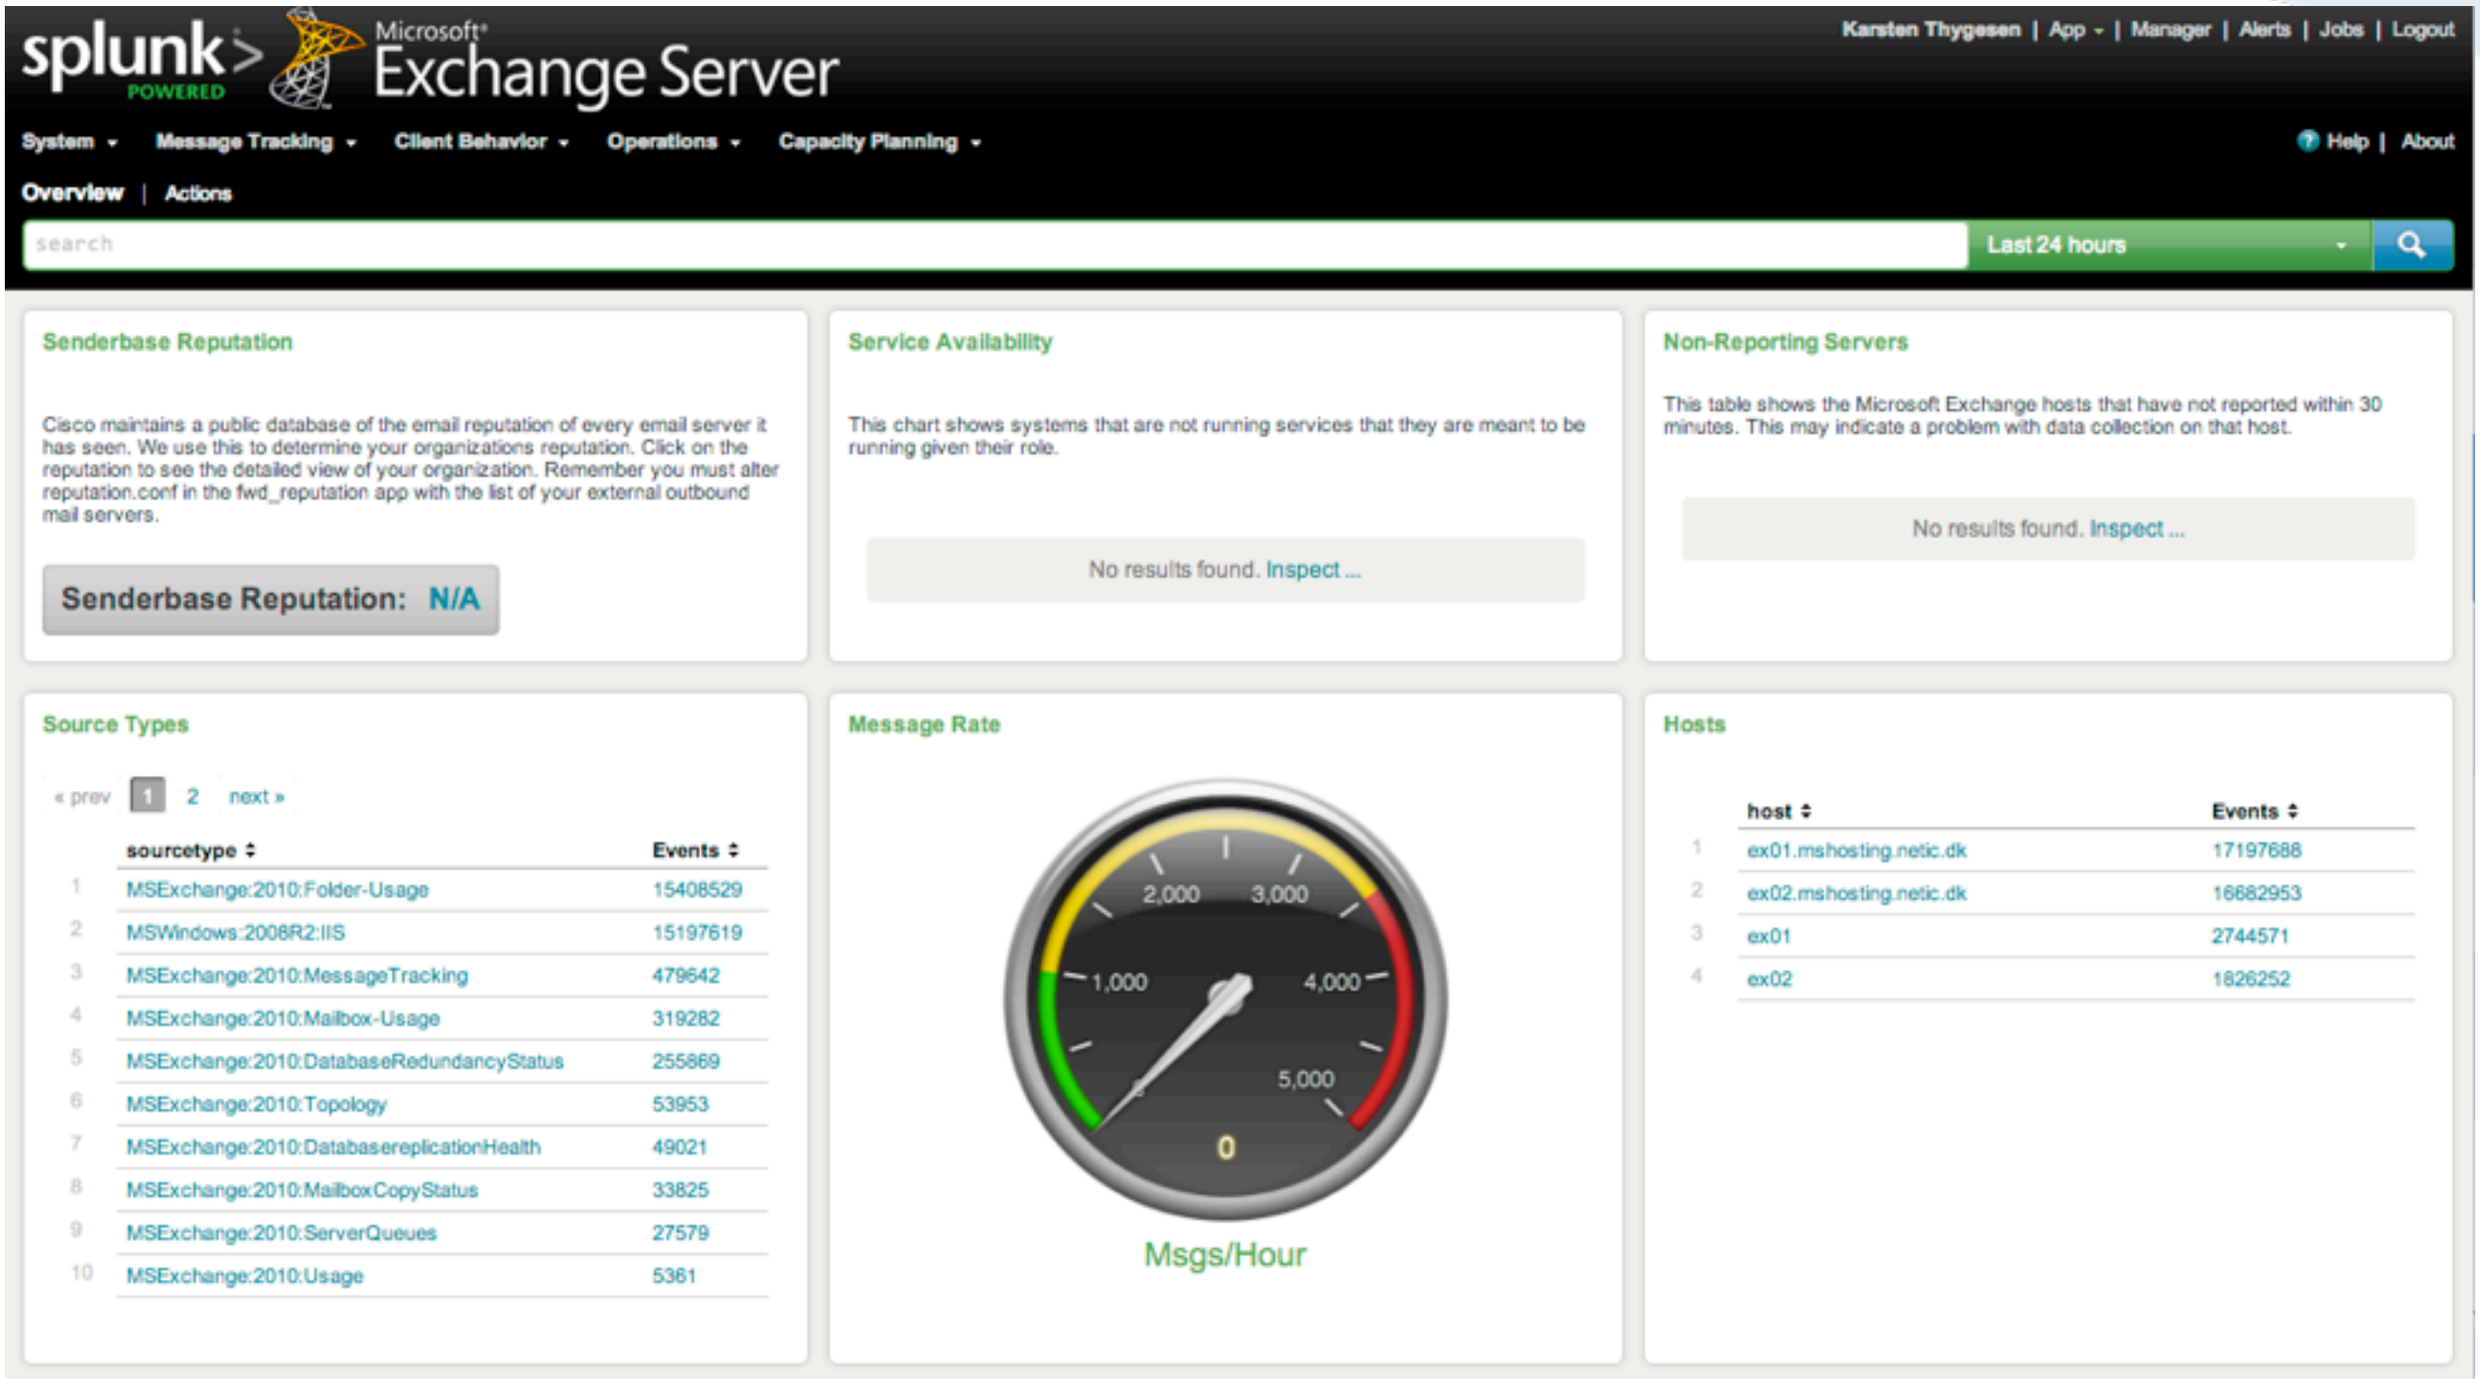

Splunk for Cisco

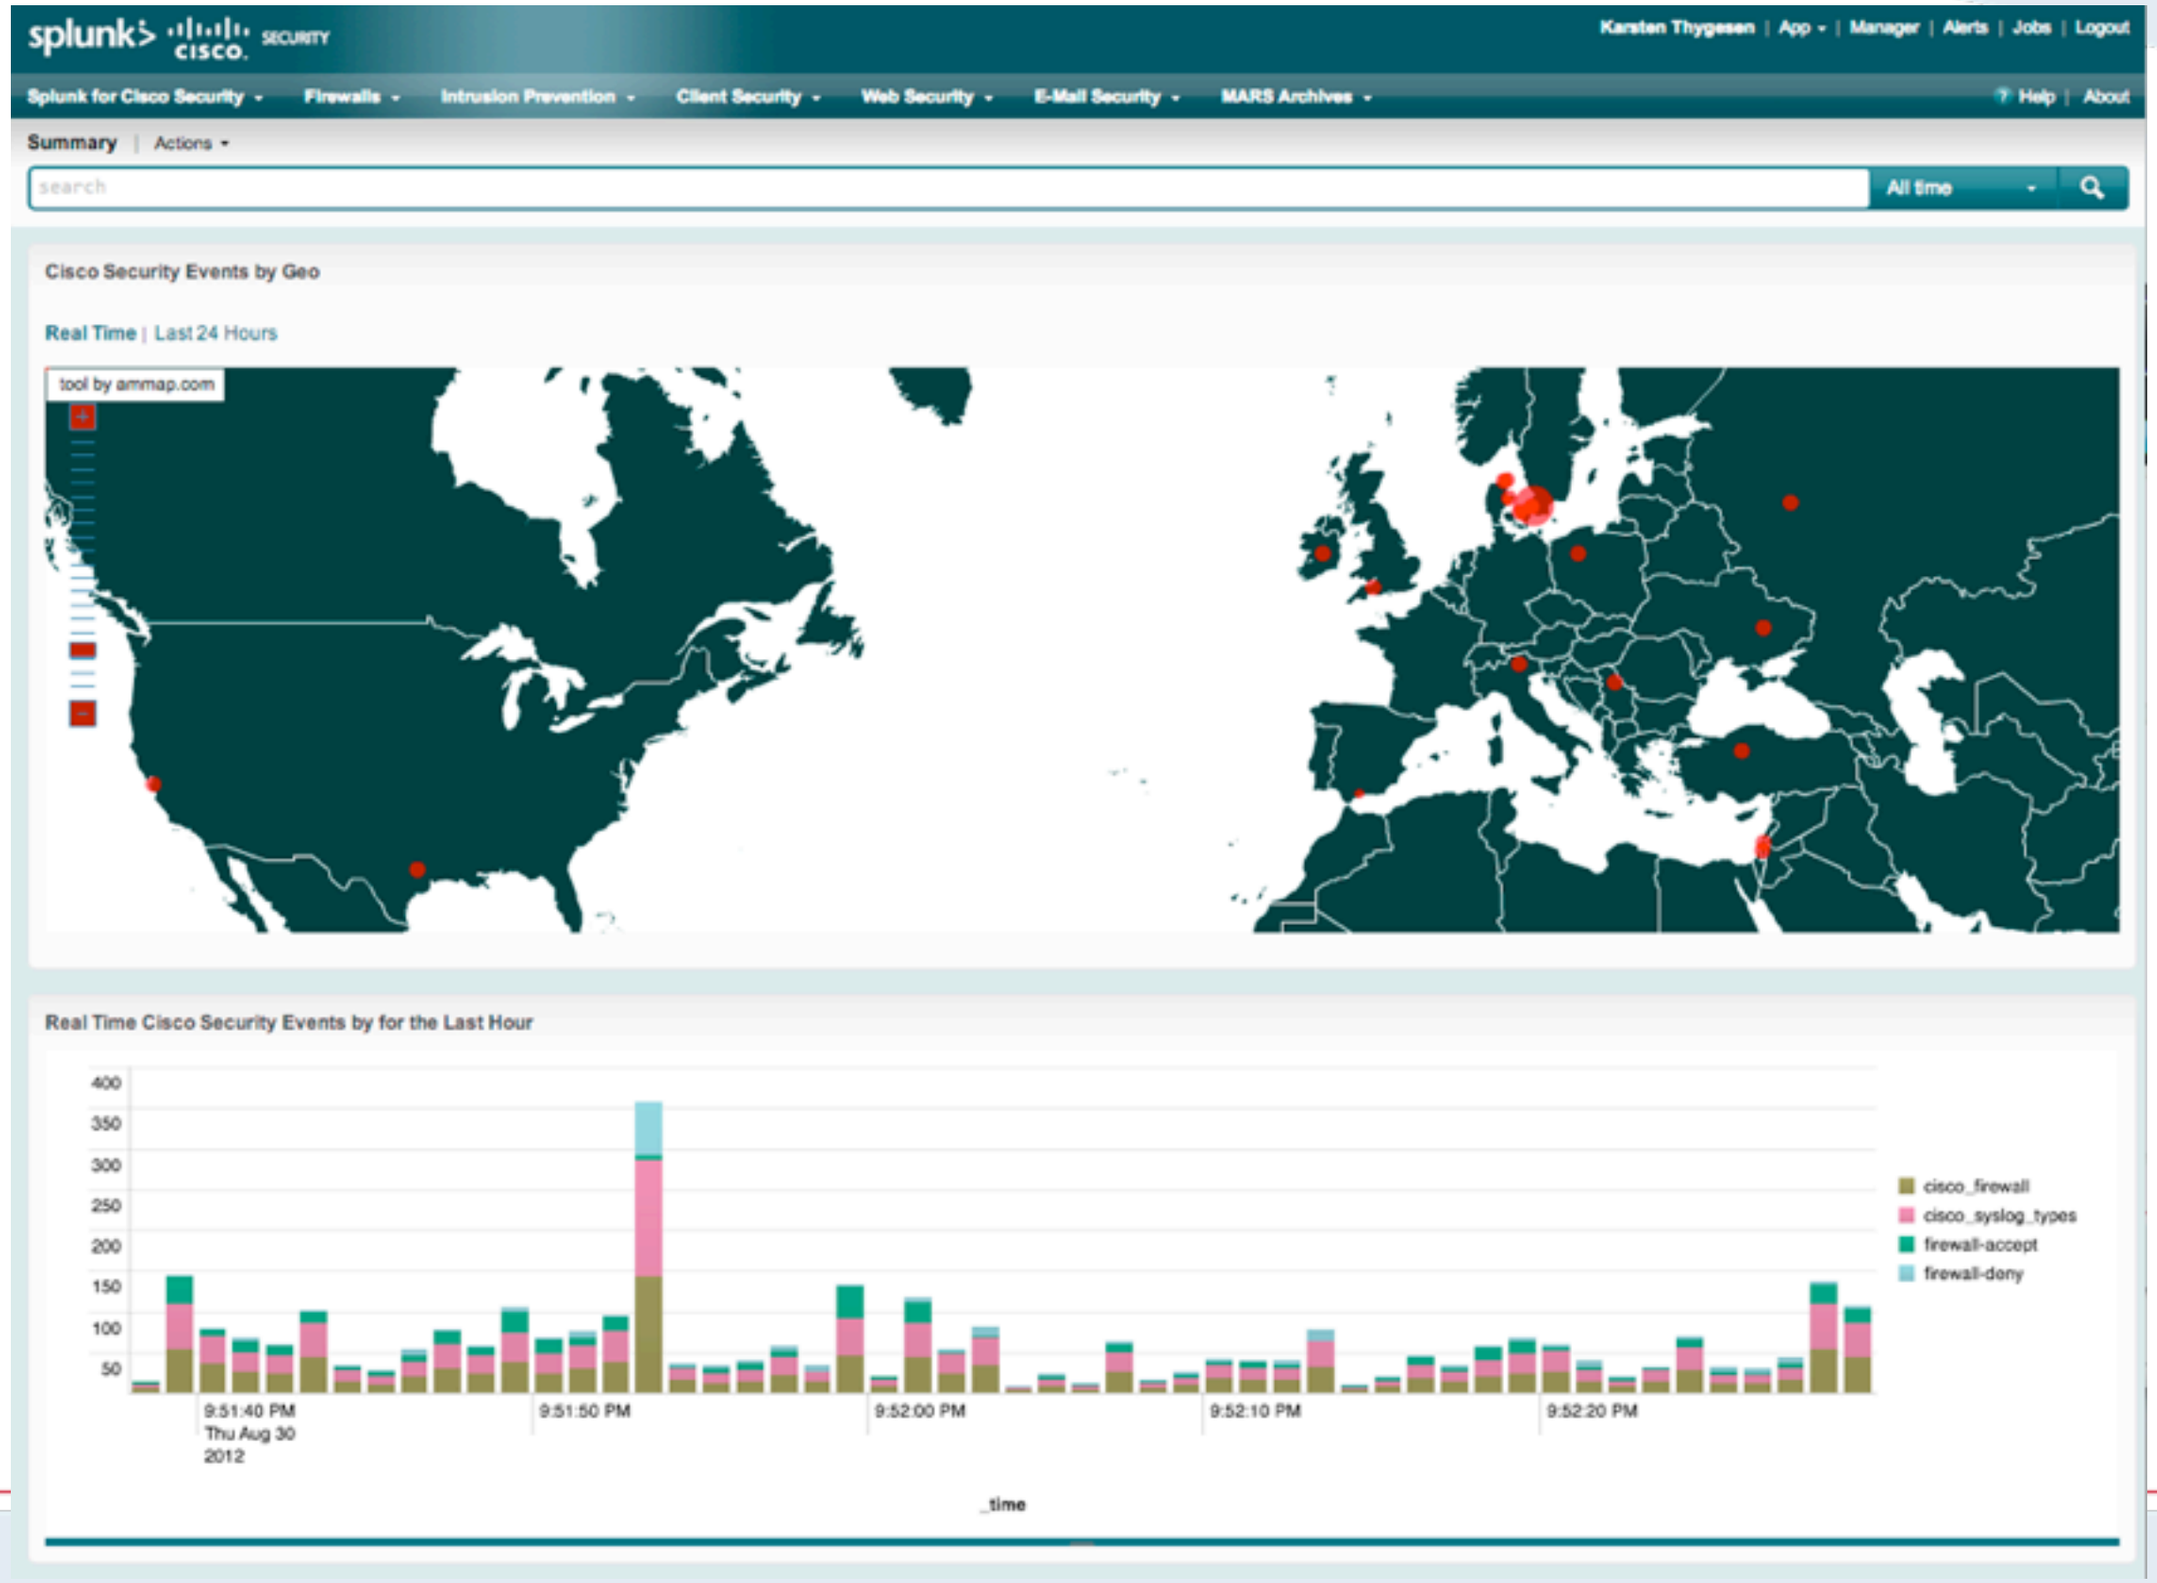

Splunk for VMWare

splunk>vmware

Karsten Thygesen | App | Manager | Alerts | Jobs | Logout

Search | About Splunk App for VMware | Views | Inventory | Performance | Security | Task | Event | Solution Administration | Troubleshooting | Help | About

Capacity Planning for Clusters - CPU Headroom | Actions

This dashboard shows the overall headroom to add more Virtual Machines to a particular cluster. Failover reserve capacity of a cluster is determined such that it accomodates one host going down with all hosts being equal in specification.
Capacity statistics: Shows capacity statistics for a particular statistics. If the "Estimated number of VMs that can be added" is negative, then that is the number of VMs that need to be removed as your cluster is oversubscribed.
Note: Clusters without Hosts, DRS, and/or HA are excluded from this view.

Clusters excluded due to lack of services or hosts: 2

Last 24 hours | Cluster: AMD cluster 1

Capacity statistics for AMD cluster 1 in the last 24 hours

Hosts in Cluster	7
Powered on VM's	213
Average usage (MHz) per VM	179.149504
Total MHz available in cluster	139041
Estimated number of VM's that can be added	452

213 VM's powered on in AMD cluster 1 in the last 24 hours

« prev 1 2 3 4 5 6 7 8 9 10 next »

vm	AvgUsg_mhz	MaxUsg_mhz
gos-app01.dodi-test.netic.dk	20.824716	N/A
test01.hprofilen.regsj.netic.dk	62.416100	N/A
dev01.regsj.netic.dk	40.809297	N/A
svn-internal01.netic.dk	11.660317	N/A
admin.regsj.netic.dk	39.637415	N/A
ext15-niab01.nsp-test.netic.dk	140.951927	N/A
ext15-cniab01.nsp-test.netic.dk	226.073696	N/A
prodcd01.medicin-it.dk	61.909977	N/A
stage01.fmk.netic.dk	91.452381	N/A
stage02.fmk.netic.dk	84.122676	N/A

View results

Currently used MHz and available MHz over time in the last 24 hours

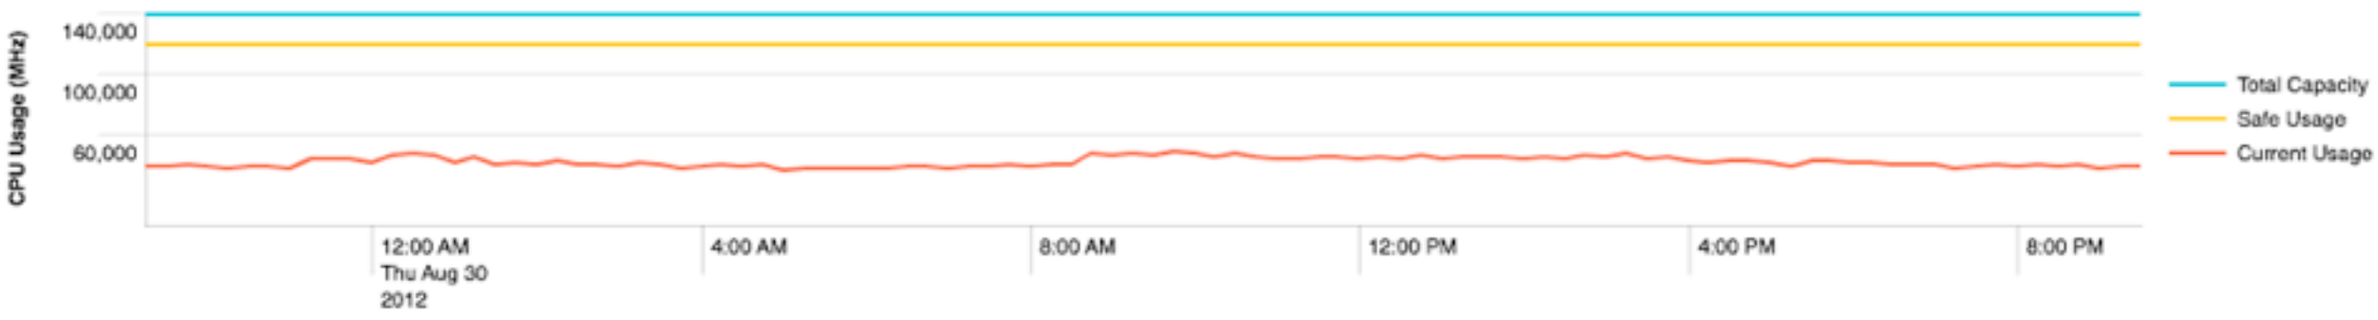

Custom Apps

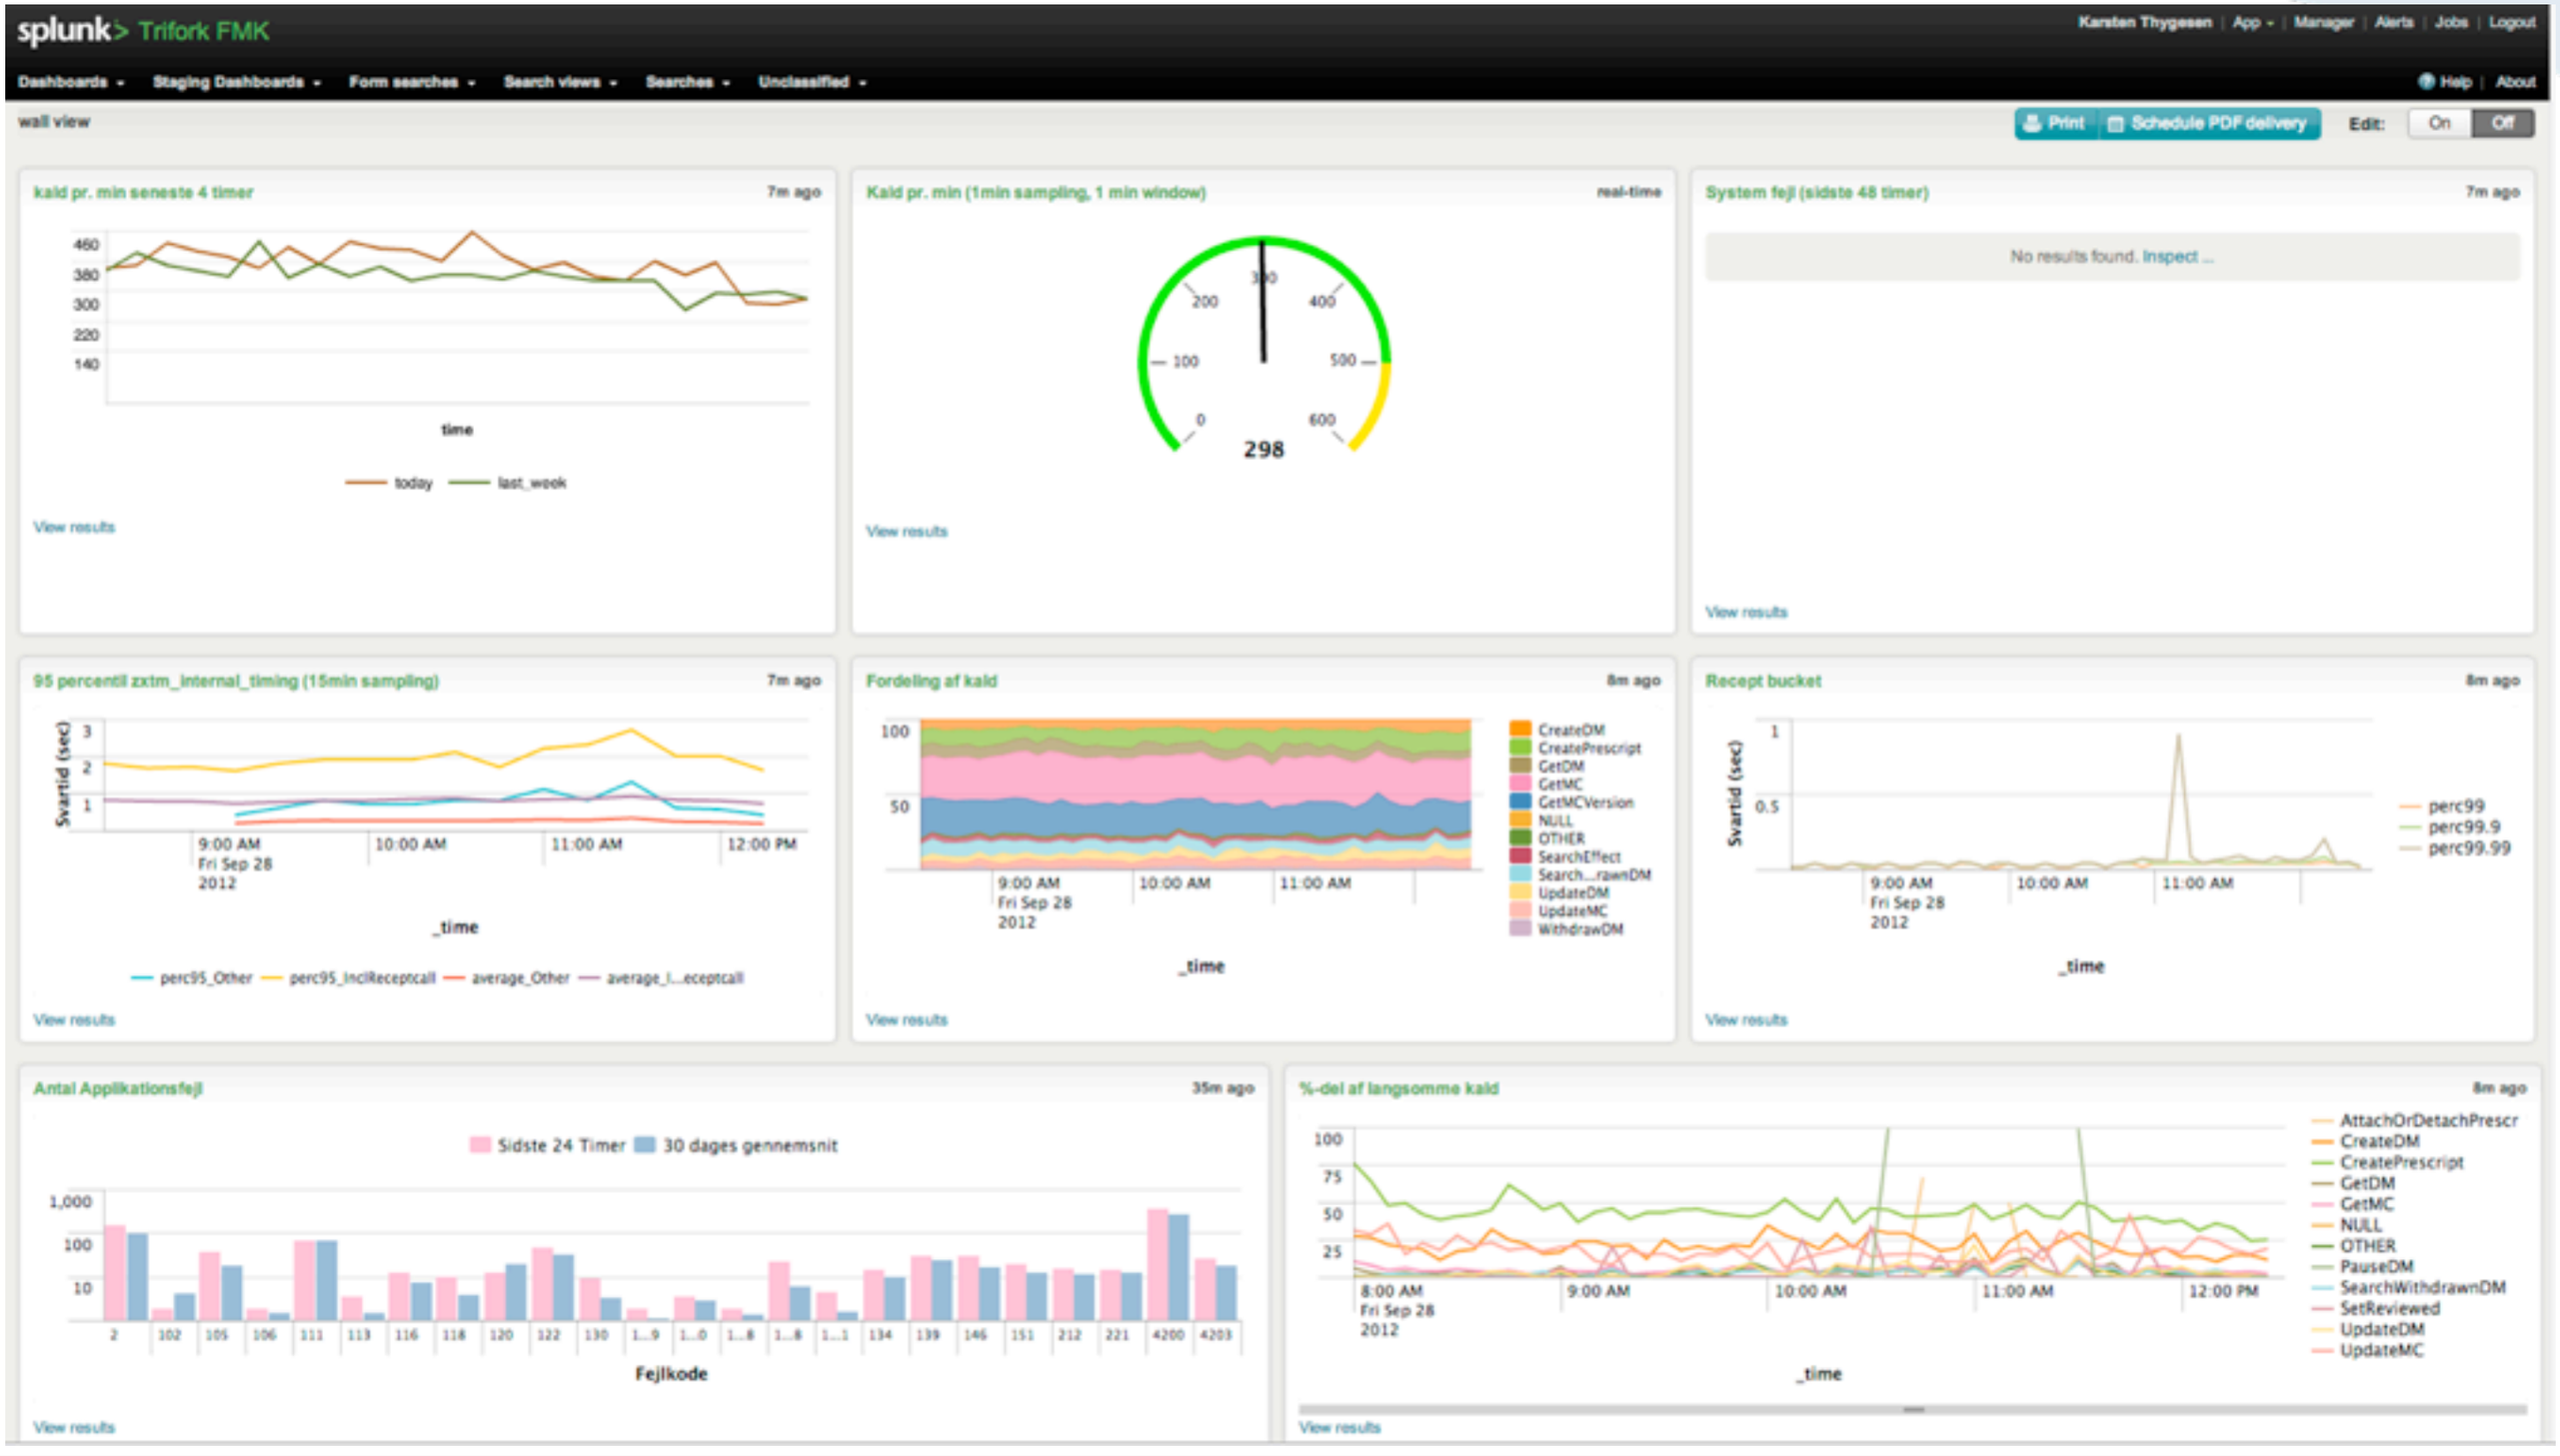
splunk> Trifork FMK
Karsten Thygesen | App | Manager | Alerts | Jobs | Logout
Dashboards | Staging Dashboards | Form searches | Search views | Searches | Unclassified
Help | About
wall view
Print | Schedule PDF delivery | Edit: On Off
kald pr. min seneste 4 timer
7m ago
460
380
300
220
140
time
today
last_week
View results
Kald pr. min (1min sampling, 1 min window)
real-time
0
100
200
300
400
500
600
298
View results
System fejl (sidste 48 timer)
7m ago
No results found. Inspect ...
View results
95 percentil zxtm_internal_timing (15min sampling)
7m ago
Svartid (sec)
3
2
1
9:00 AM Fri Sep 28 2012
10:00 AM
11:00 AM
12:00 PM
_time
perc95_Other
perc95_InclReceptcall
average_Other
average_I...eceptcall
View results
Fordeling af kald
8m ago
100
50
9:00 AM Fri Sep 28 2012
10:00 AM
11:00 AM
_time
CreateDM
CreatePrescript
GetDM
GetMC
GetMCVersion
NULL
OTHER
SearchEffect
Search...rawnDM
UpdateDM
UpdateMC
WithdrawDM
View results
Recept bucket
8m ago
Svartid (sec)
1
0.5
9:00 AM Fri Sep 28 2012
10:00 AM
11:00 AM
_time
perc99
perc99.9
perc99.99
View results
Antal Applikationsfejl
35m ago
Sidste 24 Timer
30 dages gennemsnit
1,000
100
10
2
102
105
106
111
113
116
118
120
122
130
1...9
1...0
1...8
1...8
1...1
134
139
146
151
212
221
4200
4203
Fejlkode
View results
%-del af langsomme kald
8m ago
100
75
50
25
8:00 AM Fri Sep 28 2012
9:00 AM
10:00 AM
11:00 AM
12:00 PM
_time
AttachOrDetachPrescr
CreateDM
CreatePrescript
GetDM
GetMC
NULL
OTHER
PauseDM
SearchWithdrawnDM
SetReviewed
UpdateDM
UpdateMC
View results

What Makes Splunk Different?

What Makes Splunk Different?

Any Data

- Any format and amount of data
- Any source
- Full access to 100% of data for months/years
- Cradle-to-grave data management

What Makes Splunk Different?

Any Data

- Any format and amount of data
- Any source
- Full access to 100% of data for months/years
- Cradle-to-grave data management

Completely Flexible

- Supports analysis, reporting, monitoring across IT silos
- Flexible dashboards present any view for any user
- Adapts to change—schema-on-the-fly design supports new or unexpected data

What Makes Splunk Different?

Any Data

- Any format and amount of data
- Any source
- Full access to 100% of data for months/years
- Cradle-to-grave data management

Completely Flexible

- Supports analysis, reporting, monitoring across IT silos
- Flexible dashboards present any view for any user
- Adapts to change—schema-on-the-fly design supports new or unexpected data

Immediate Results

- Free download, installs in minutes
- Can get started small and grow over time—from laptop to datacenters
- Initial benefits realized in hours or days

Splunk in (FMK) operations

Splunk in (FMK) operations

Challenges

- ISO-27000
- Segregation of duties
- 24x7x365 availability
- Complex setup with
  - 70+ Servers
  - Multiple Datacenters
  - Many technologies

Splunk in (FMK) operations

Challenges

- ISO-27000
- Segregation of duties
- 24x7x365 availability
- Complex setup with
  - 70+ Servers
  - Multiple Datacenters
  - Many technologies

Deploy Process

- Drain single node
- Deploy new release
- Drip traffic (1%)
- Monitor
- Increase traffic
- Monitor
- Full traffic
- Repeat with other nodes

Splunk in (FMK) operations

Challenges

- ISO-27000
- Segregation of duties
- 24x7x365 availability
- Complex setup with
  - 70+ Servers
  - Multiple Datacenters
  - Many technologies

Deploy Process

- Drain single node
- Deploy new release
- Drip traffic (1%)
- Monitor
- Increase traffic
- Monitor
- Full traffic
- Repeat with other nodes

Splunk Advantage

- Dev access to realtime logs
- Operational insight/baseline
- Catch unknown errors
- Alert on complex errors
- Improved and easier reporting
- Adapt to changing logformat
- Infrastructure included

Easy to Get Started

Download from splunk.com

1. Free* download **2. Eat your Machine Data** **3. Start Splunking**

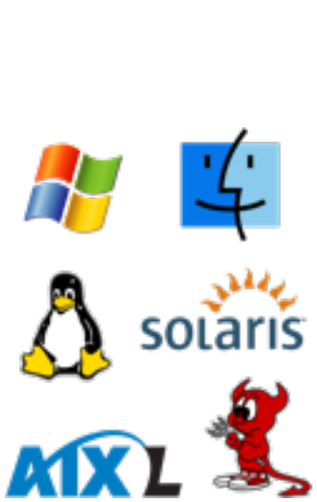

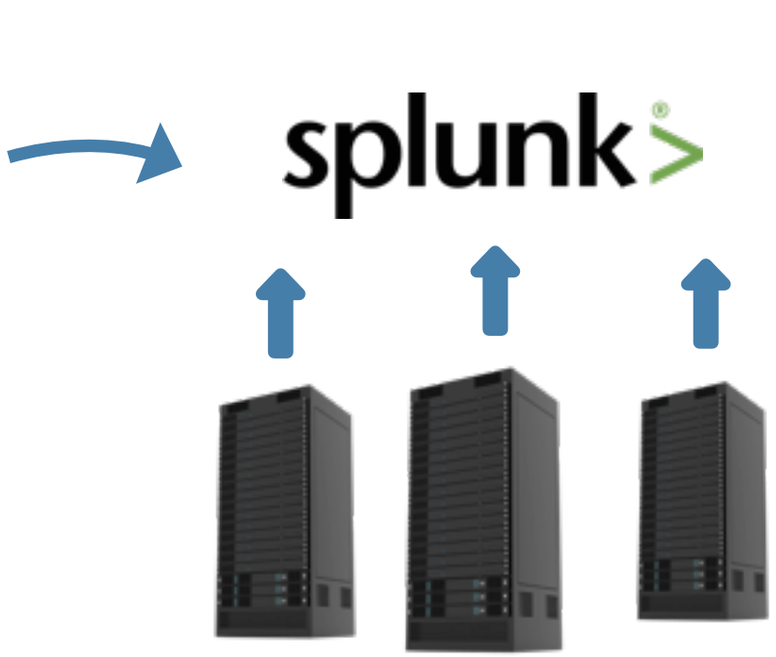

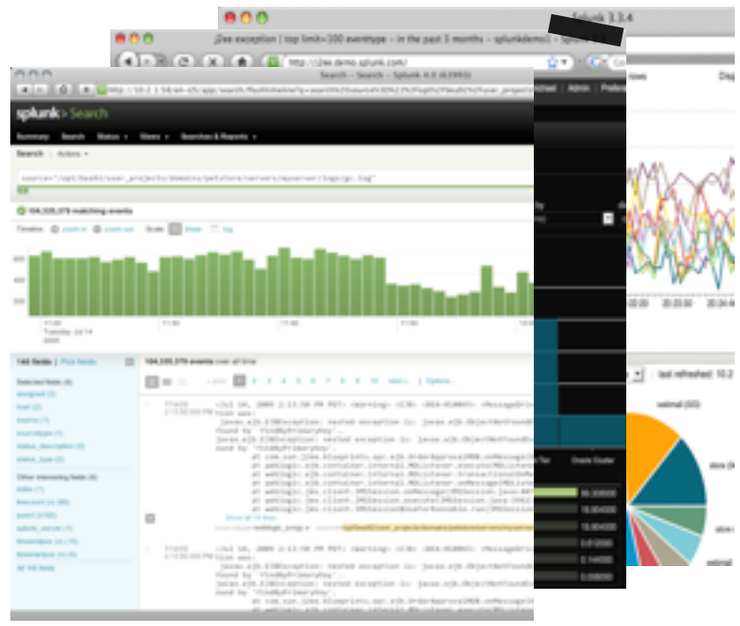

Datacenter

***free up to 500Mb/day**

Questions

Questions

?

!

Come by booth for in-depth demo

Reference Hardware

- Intel x86/64bit
- 8 cores
- 8GB memory
- 800 IOPS disks

Reference Hardware

- Intel x86/64bit
- 8 cores
- 8GB memory
- 800 IOPS disks

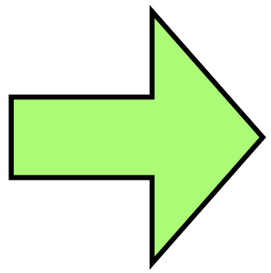

- 100GB/day
- 4 concurrent users