.conf2013

YOUR DATA
NO LIMITS

# End-to-End Monitoring From Scratch

Robert Gustafson, Staff Engineer, DirectTV

#splunkconf

splunk>

# Background: Me

- Staff Engineer at DIRECTV's Los Angeles Broadcast Center
- Three Things I Love
  - USC
  - Detroit Sports
  - Cycling

# Background: Project

- High Revenue
- New Venture
- Limited Monitoring
- Variety of Systems
- Lots of Blind Spots

# Why Splunk Over Other Tools?

- Give it data in any format
- Data, alarming mechanisms, UI's don't have to fit a mold
- All the heavy lifting is already done
- Finding cause faster than humans
- Developer Freedom!

# Challenges Splunk Solved

- Project scope
- Little insight into workings of platform
- Users not nerdy enough to decode log files

# Challenge: Scope

Before expansion: “The world is flat”

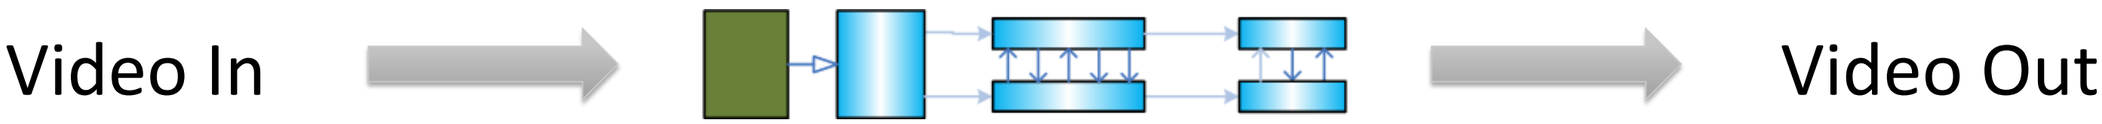

# Challenge: Scope

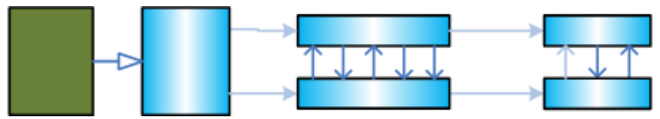

# Challenge: Scope

After expansion: “The world is round”

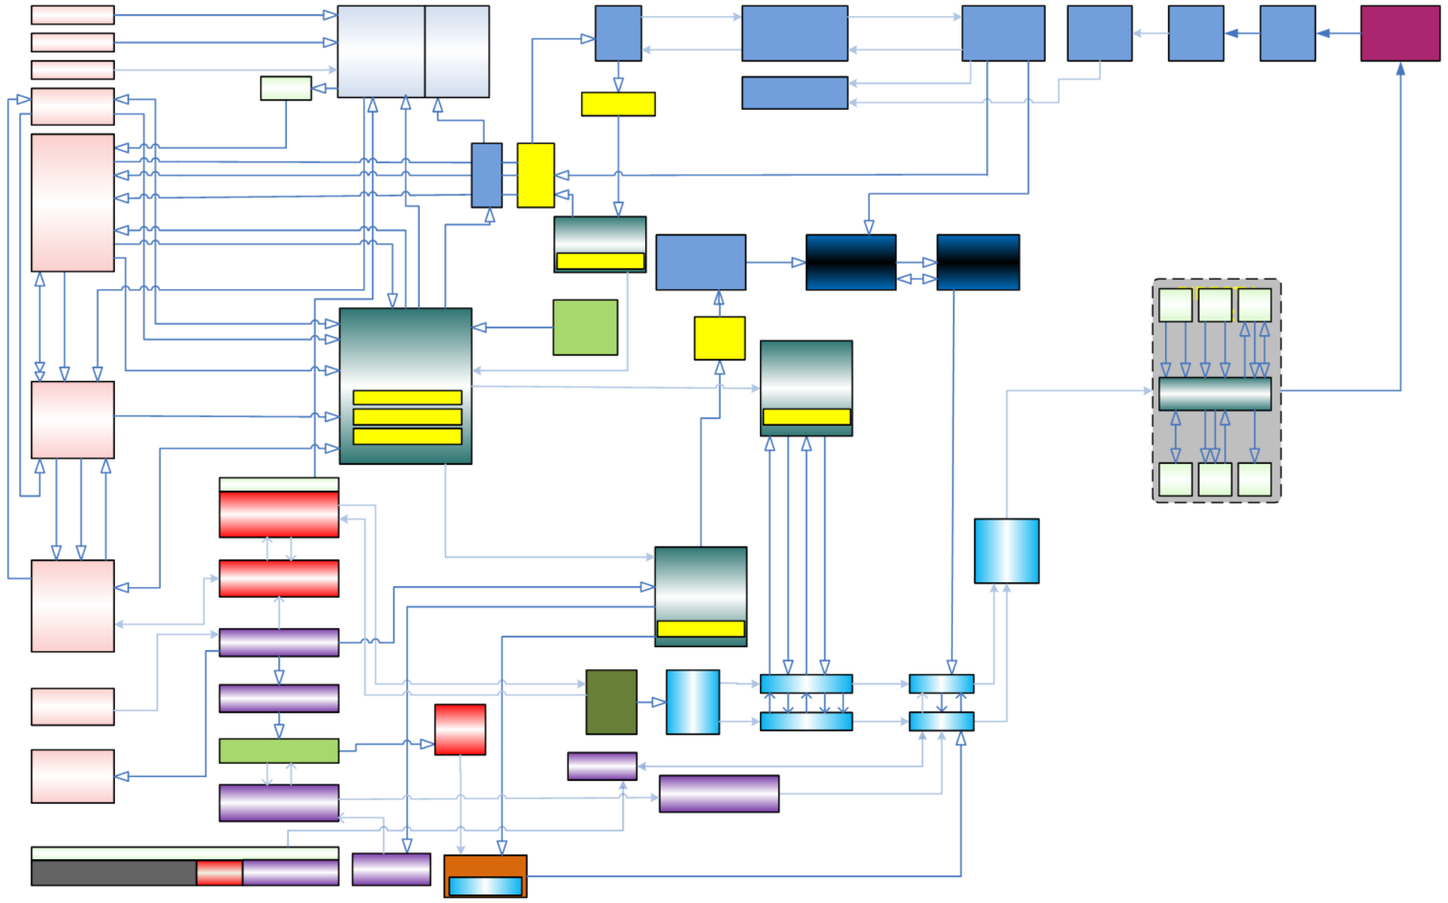

# Challenge: Scope

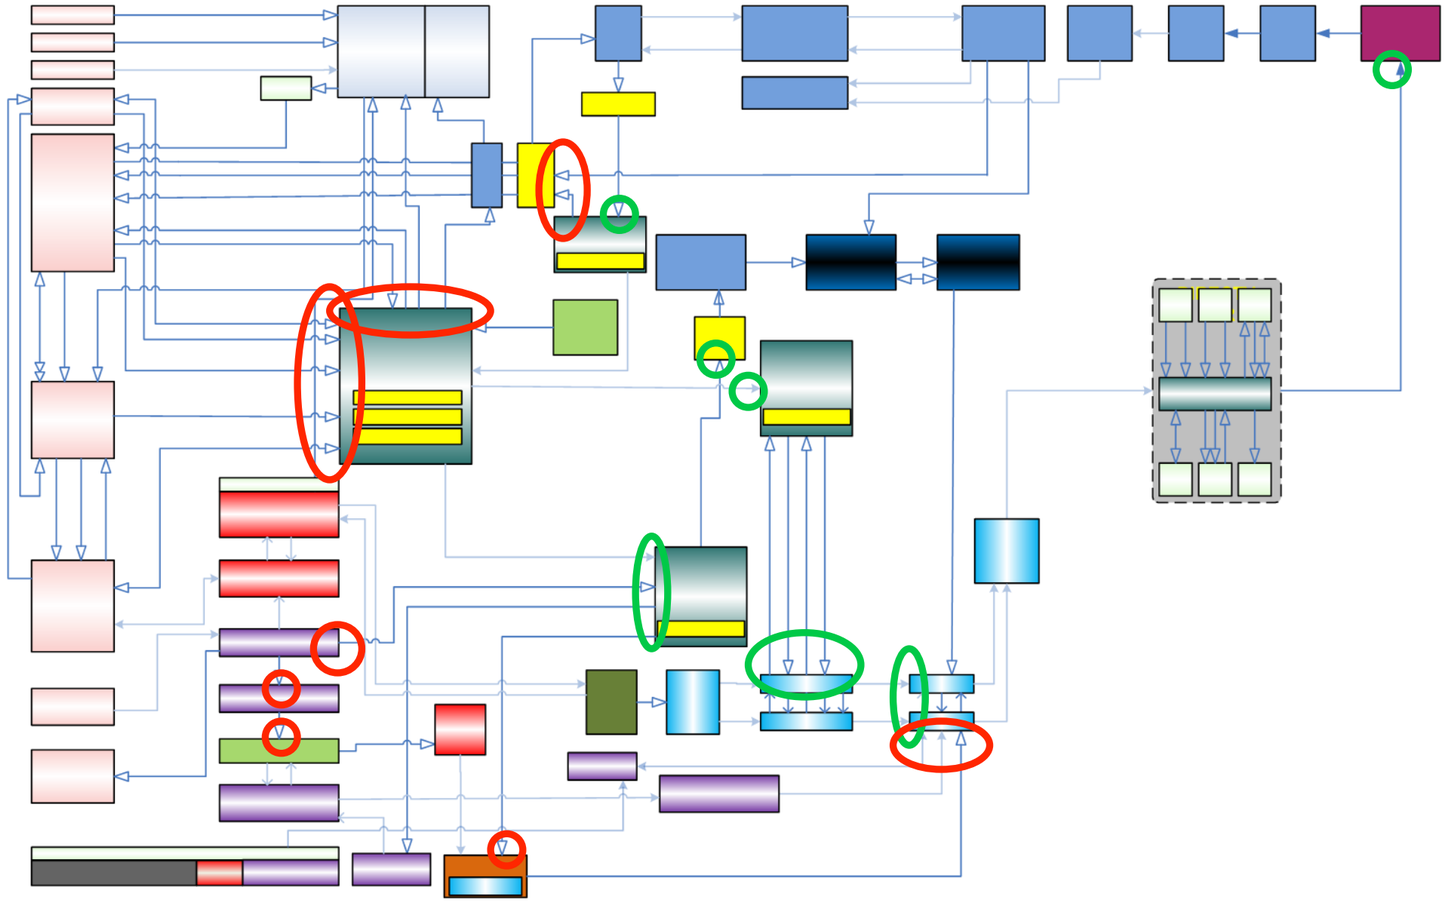

# Challenge: Scope

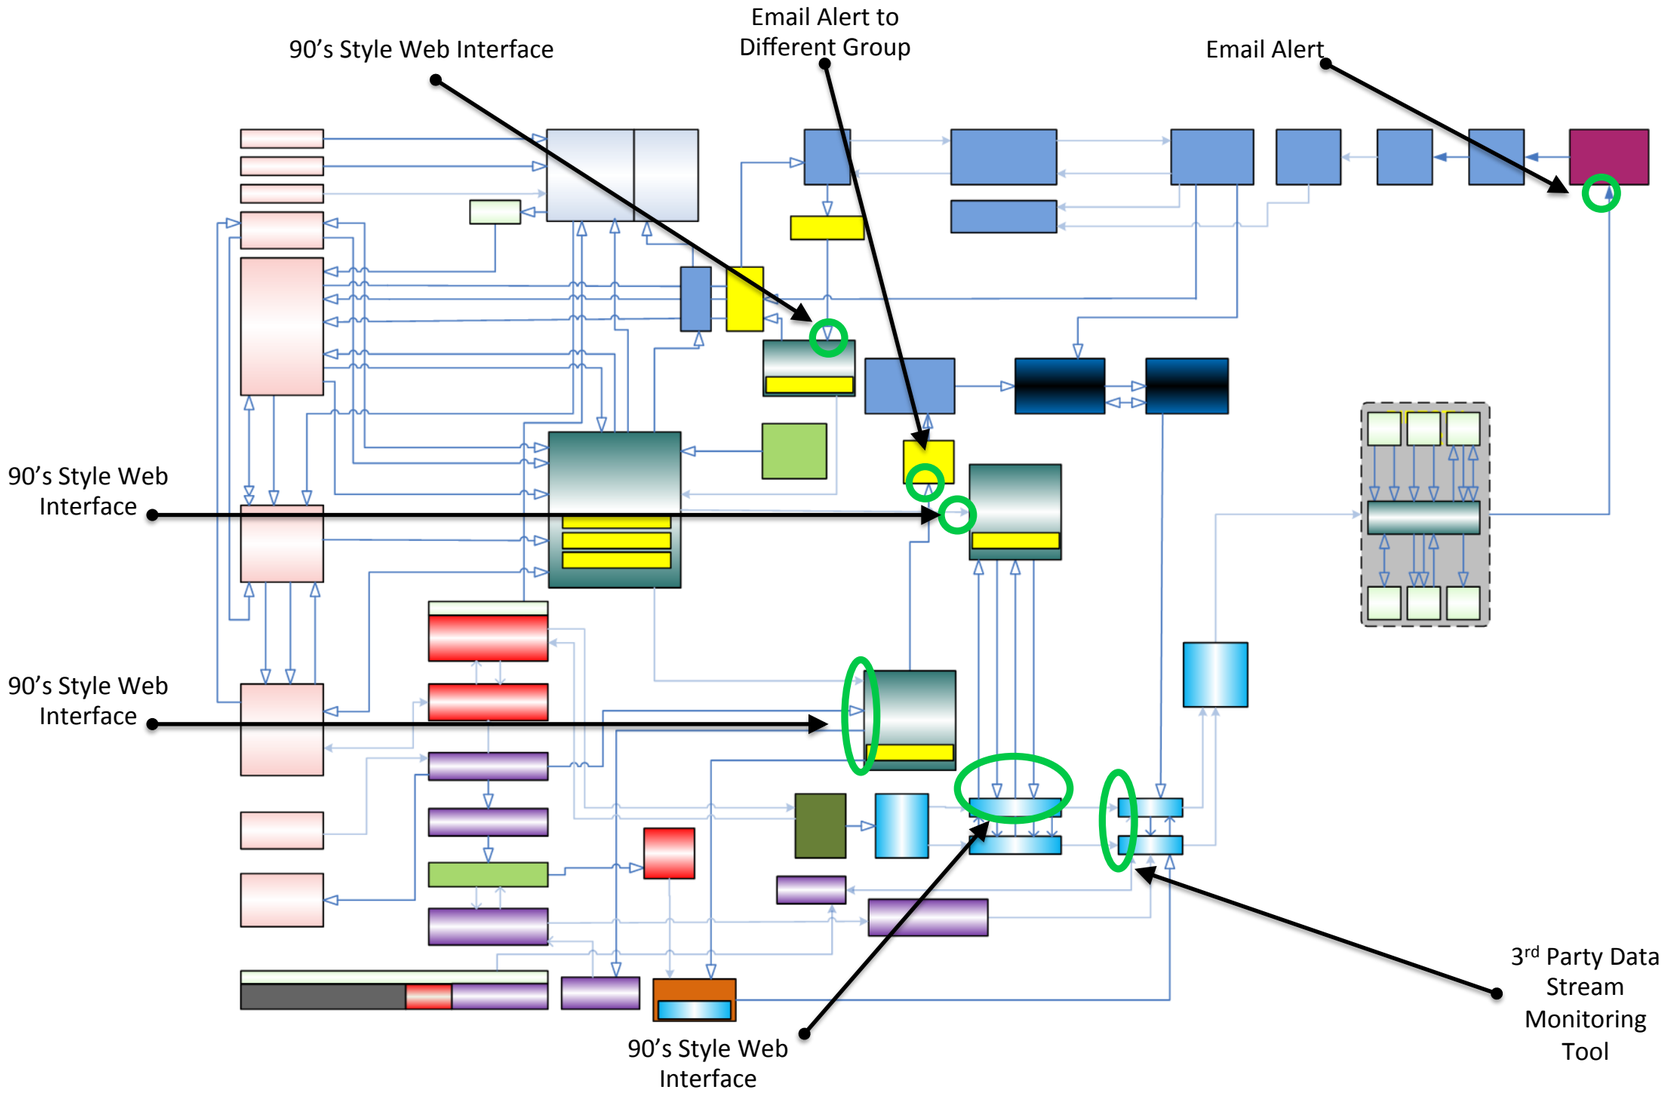

# Challenge: Scope

## Is it really 2013?

| Host | | | | | | Status/Control | Host | | | | | | Status/Control |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| prod-la-01 | ⚠ | ▪ | ▪ | ⛔ | ▪ | Active | prod-la-02 | ✅ | ▪ | ▪ | ✅ | ▪ | Inactive |
| prod-la-01 | ⛔ | ▪ | ▪ | ▪ | ✅ | Inactive | prod-la-02 | ⚠ | ▪ | ▪ | ▪ | ⛔ | Active |
| prod-la-01 | ⚠ | ▪ | ▪ | ▪ | ▪ | | prod-la-02 | ⚠ | ▪ | ▪ | ▪ | ▪ | |
| prod-la-03 | ⚠ | ⛔ | ⚠ | ▪ | ▪ | Online Go Offline | prod-la-04 | ⚠ | ⛔ | ⚠ | ▪ | ▪ | Offline Go Online |
| prod-la-05 | ⚠ | ⛔ | ⚠ | ▪ | ▪ | Online Go Offline | prod-la-06 | ⚠ | ⛔ | ⚠ | ▪ | ▪ | Offline Go Online |
| prod-la-07 | ⚠ | ⛔ | ⚠ | ▪ | ▪ | Online Go Offline | prod-la-08 | ⚠ | ⛔ | ⚠ | ▪ | ▪ | Offline Go Online |
| prod-la-09 | ✅ | ⛔ | ⚠ | ▪ | ▪ | Online Go Offline | prod-la-10 | ✅ | ⛔ | ⚠ | ▪ | ▪ | Offline Go Online |
| prod-la-11 | ⚠ | ⛔ | ✅ | ▪ | ▪ | Online Go Offline | prod-la-12 | ⚠ | ⛔ | ✅ | ▪ | ▪ | Offline Go Online |
| prod-la-13 | ⚠ | ✅ | ✅ | ▪ | ▪ | Online Go Offline | prod-la-14 | ⚠ | ✅ | ✅ | ▪ | ▪ | Offline Go Online |
| prod-la-15 | ✅ | ⛔ | ✅ | ▪ | ▪ | Online Go Offline | prod-la-16 | ✅ | ⛔ | ✅ | ▪ | ▪ | Offline Go Online |
| prod-la-17 | ✅ | ✅ | ⚠ | ▪ | ▪ | Online Go Offline | prod-la-18 | ✅ | ✅ | ⚠ | ▪ | ▪ | Offline Go Online |
| prod-la-19 | ⚠ | ✅ | ✅ | ▪ | ▪ | Online Go Offline | prod-la-20 | ⚠ | ✅ | ✅ | ▪ | ▪ | Offline Go Online |

# Challenge: Scope

What happens when you have no idea what you're looking for??

- Look for events that occurred while systems experienced problems
  - Narrow results using the punct field
  - "`|dedup punct`"
  - "`|stats count by punct`"

# Challenge: Scope

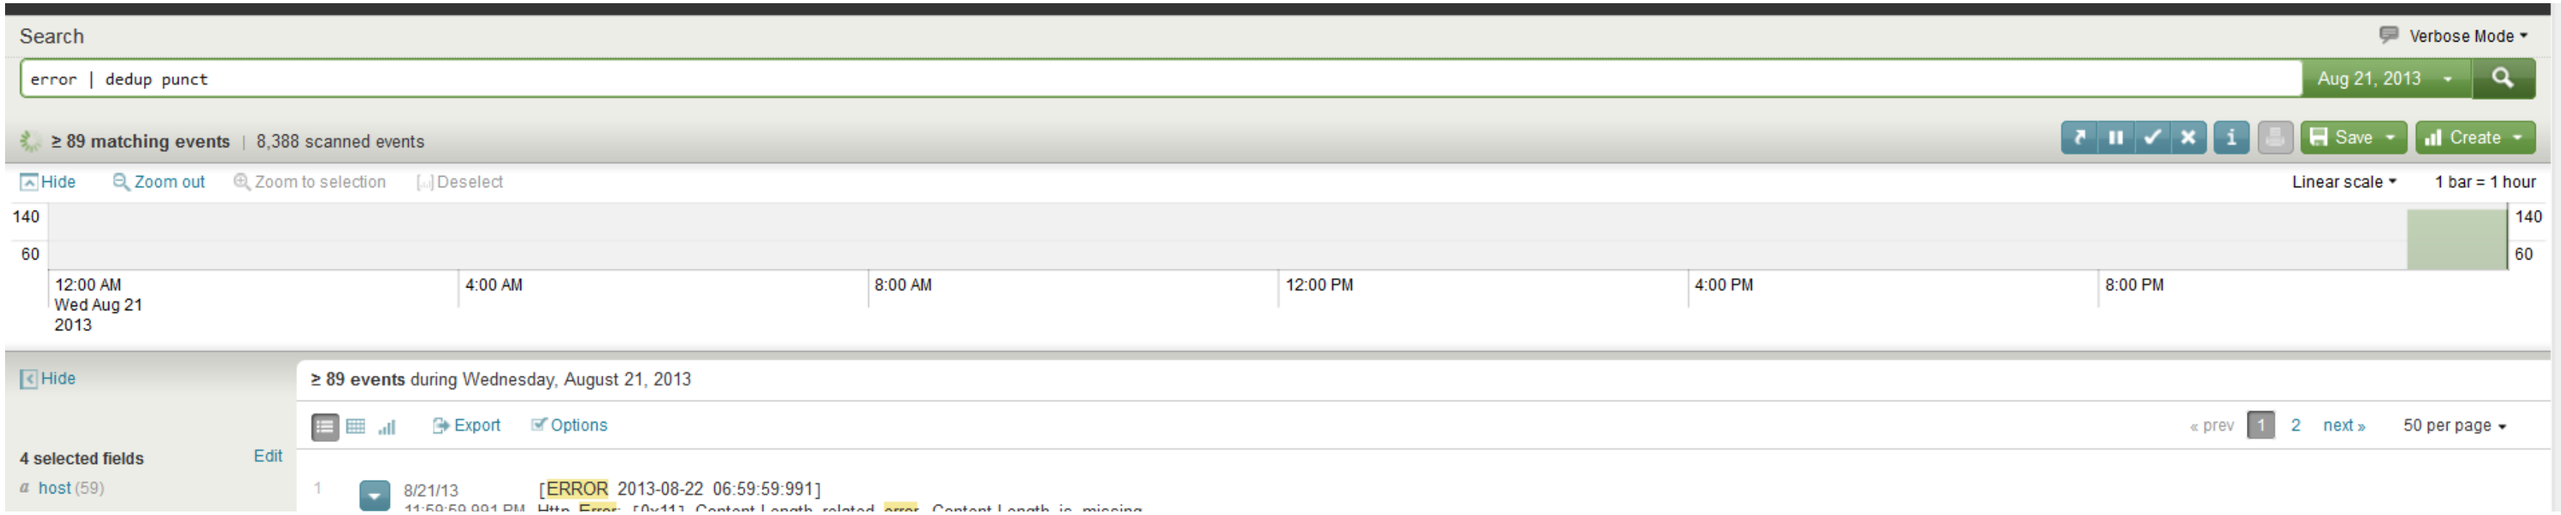

# Ignoring the Garbage

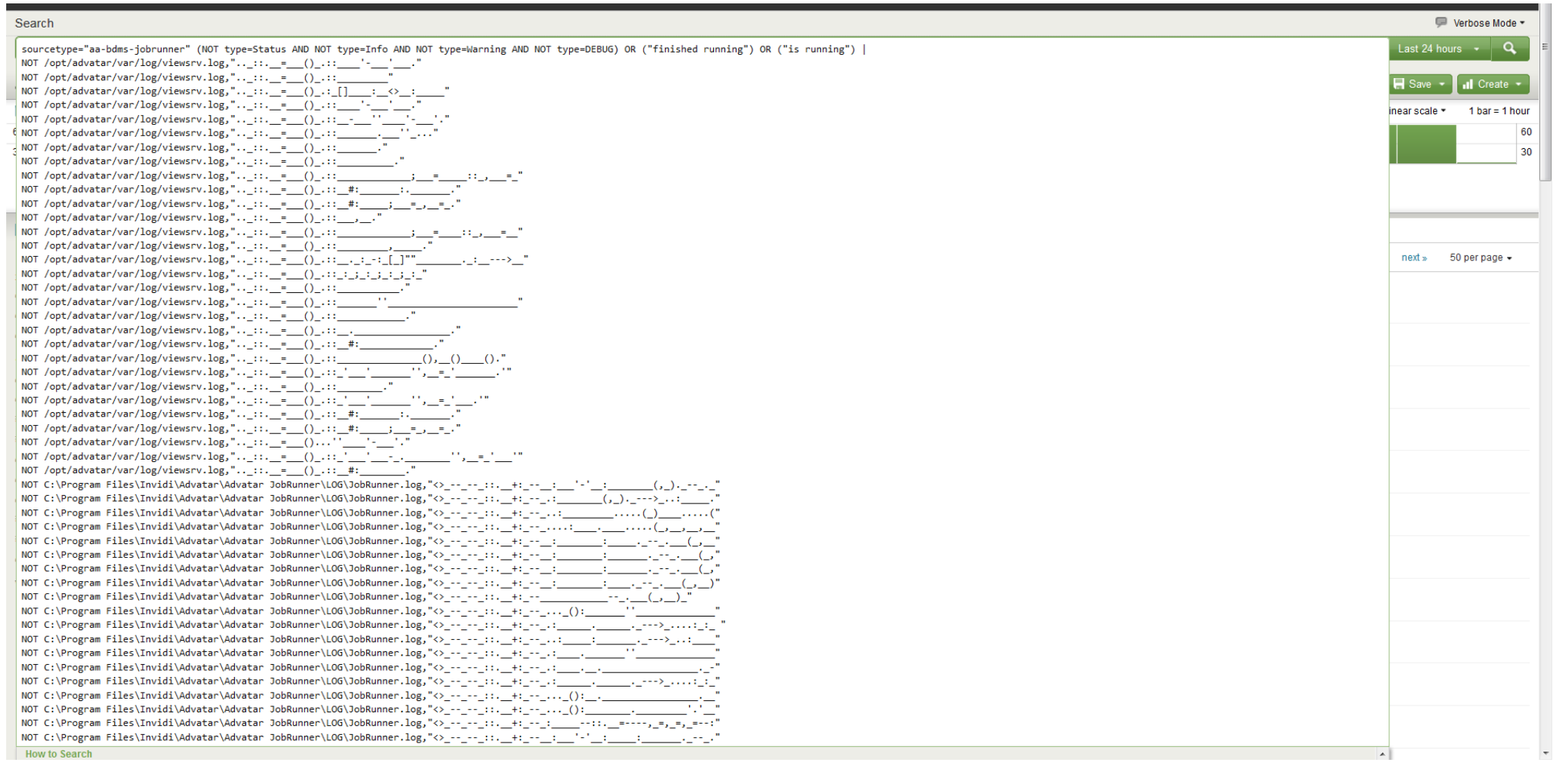

# Ignoring the Garbage

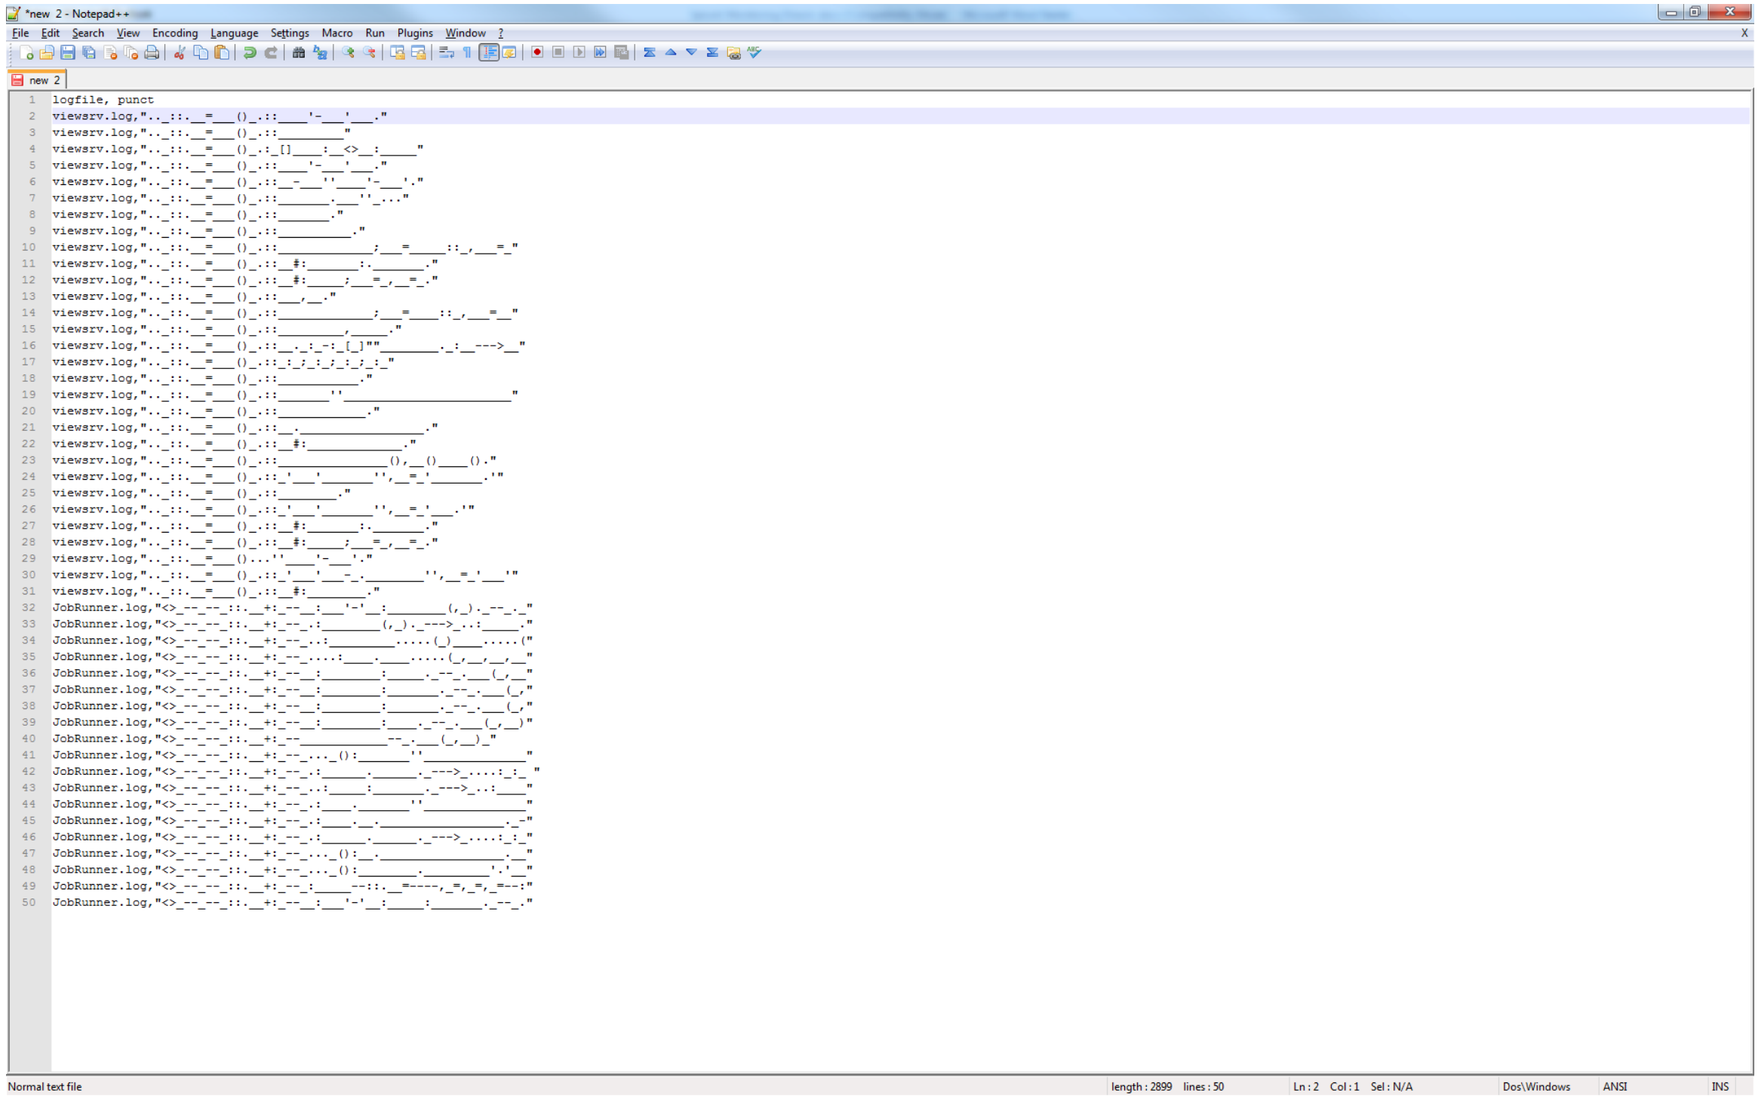

# Ignoring the Garbage

```
Error NOT [inputcsv ignore_errors.csv
| search logfile="JobRunner.log” |
rename ignore_punct AS punct | fields
punct]
```

# Exception Based Monitoring

- What am I looking at? (High Level)
  - Does this log entry directly impact the customer experience?
  - Is it actionable?
- If yes to both: make an alert with a saved search

# Exception Based Monitoring

Saved Search:

- Query to find error condition

```
host=splunkforwarder* ERROR no file found
```

- Secondary query to alert on

```
| stats count | where count > 1
```

- Time Ranges

# Proactive Monitoring

- What events can begin to indicate failure?
  - ex: Steady disk space increase, Repeated timeouts
- Use Splunk Enterprise to set thresholds or look for patterns

# Proactive Monitoring: Patterns

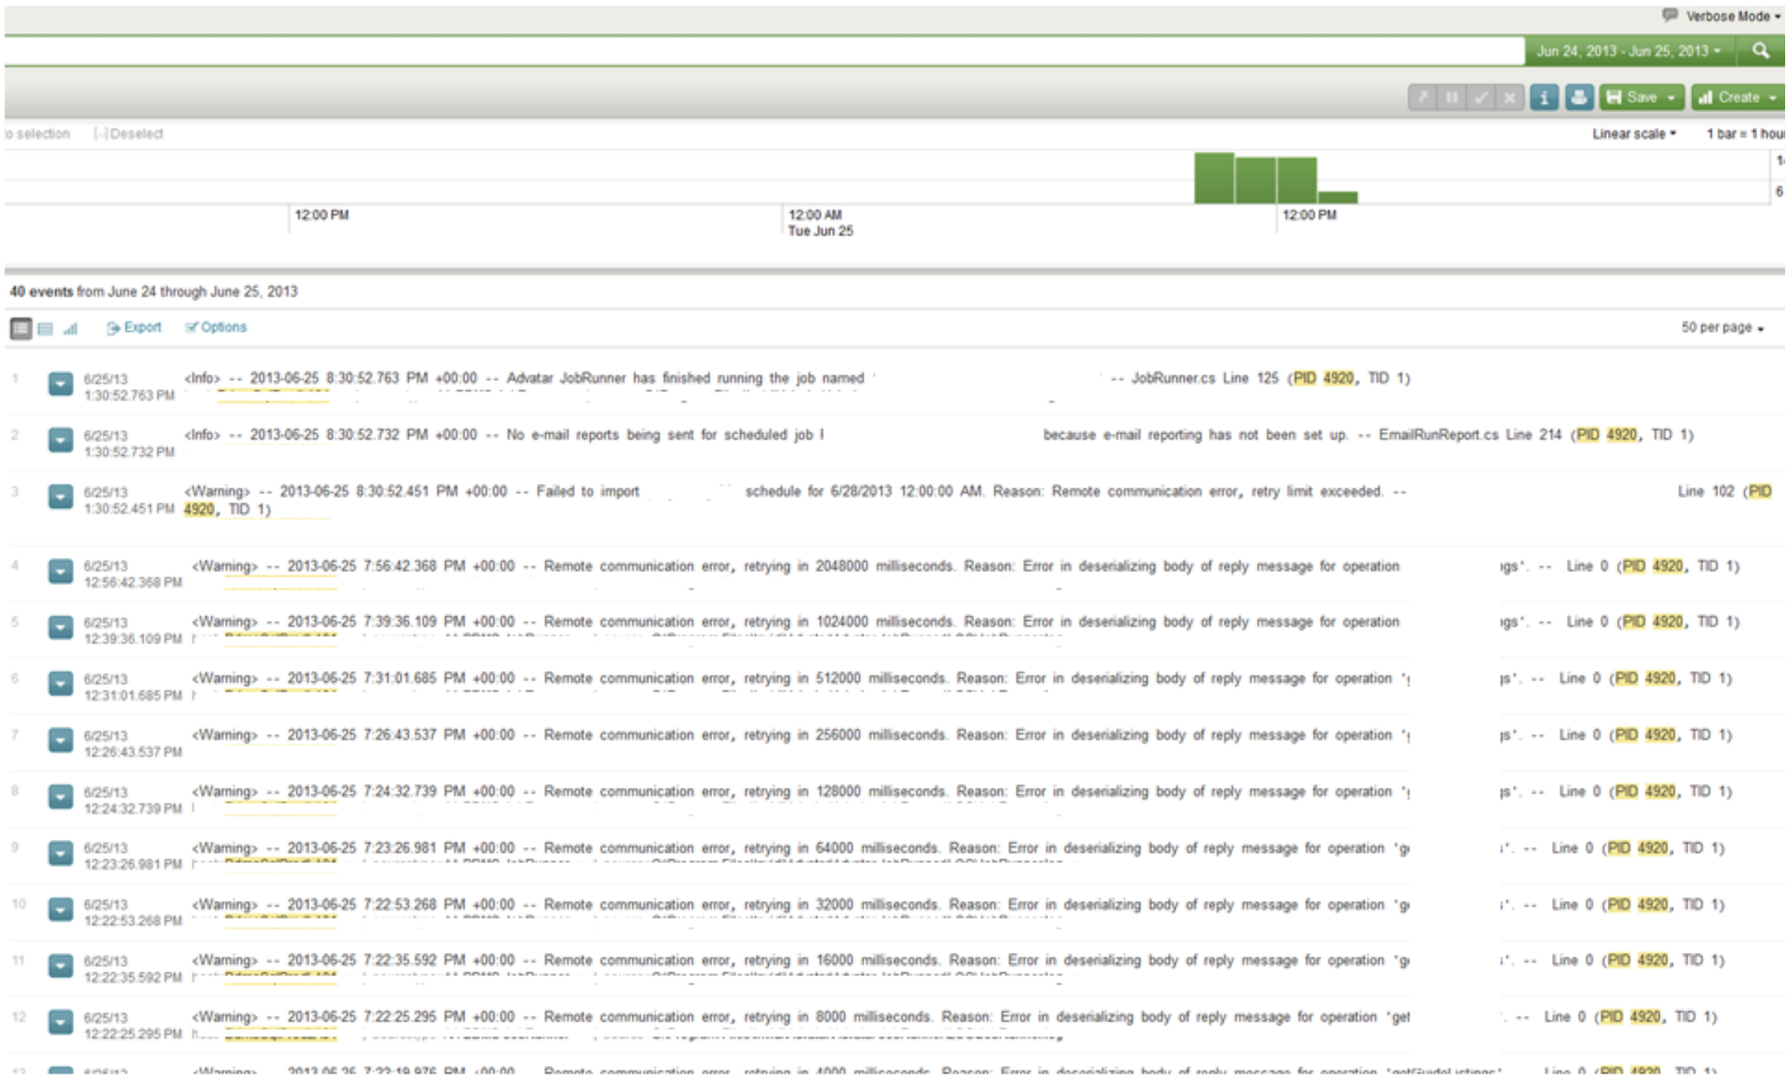

# Proactive Monitoring

- Query to find error condition

```
host=splunkforwarder* WARNING timeout
 | stats count as num_errors_by_host by host
| stats sum(num_hosts_in_error) as total_errors
            count(eval(num_hosts_in_error>0))  as num_hosts_in_error
```

- Secondary query to alert on

```
| where num_hosts_in_error > 10

| where total_errors > 100

| where num_hosts_in_error > 10 AND total_errors > 100
```

- Time Ranges

# Merging Different Data Types

- Multiple servers saying the same in different ways
- Who do you listen to?

# Merging Different Data Types

- Unify the data
- Correlate the data
- Alert on the data

# Merging Different Data Types

```
<tr><td align="center"><font color="#000000">Aug 23 05:25:39.000000000</font></td>
<td align="center"><font color="#000000">Aug 23 05:25:39.000000000</font></td>
<td align="center"><font color="#000000">405889479</font></td>
<td align="center"><font color="#000000">-324 ms</font></td>
<td align="center"><font color="#000000">-prod-la-05</font></td>
<td align="center"><font color="#000000"></font></td>
<td align="center"><font color="#000000">110</font></td>
<td align="center"><font color="#000000">0x100</font></td>
<td align="center"><font color="#000000">10153</font></td>
<td align="center"><font color="#000000">1813253</font></td>
<td align="center"><font color="#000000">Aug 23 05:00:00</font></td>
<td align="center"><font color="#000000">Aug 23 06:00:00</font></td>
<td align="center"><font color="#000000">1</font></td>
<td align="center"><font color="#000000">Successful</font></td>
</tr>
```

2013.08.23 06:13:43 GMT (event) Successful ID= 1813253 pid=14283

**57 results** since 7:48:43 PM August 22, 2013

Export | Options | « prev 1 2 3 4 5 6 next » | 10 per page

Overlay: None

| | _time | ID |
|---|---|---|
| 1 | 8/22/13 10:38:09.000 PM | 1813913 |
| 2 | 8/22/13 10:27:53.000 PM | 1813323 |
| 3 | 8/22/13 10:25:58.000 PM | 1813674 |
| 4 | 8/22/13 10:25:39.000 PM | 1813253 |
| 5 | 8/22/13 10:25:35.000 PM | 1813912 |
| 6 | 8/22/13 10:23:41.000 PM | 1814096 |
| 7 | 8/22/13 10:22:57.000 PM | 1815395 |
| 8 | 8/22/13 10:22:43.000 PM | 1812378 |
| 9 | 8/22/13 10:22:29.000 PM | 1814303 |
| 10 | 8/22/13 10:21:27.000 PM | 1812538 |

# Correlating events

```
sourcetype="box1" | table ID |
join type=left ID [sourcetype="box2"  | eval box2ID=ID | table box2ID,ID] |
table ID, box2ID |
| where ID!=box2ID
```

| ID | box2ID |
|---|---|
| 1 | 1 |
| 2 | |
| 3 | 3 |
| 4 | 4 |
| 5 | 5 |
| 6 | |
| 7 | 7 |

# Splunk!

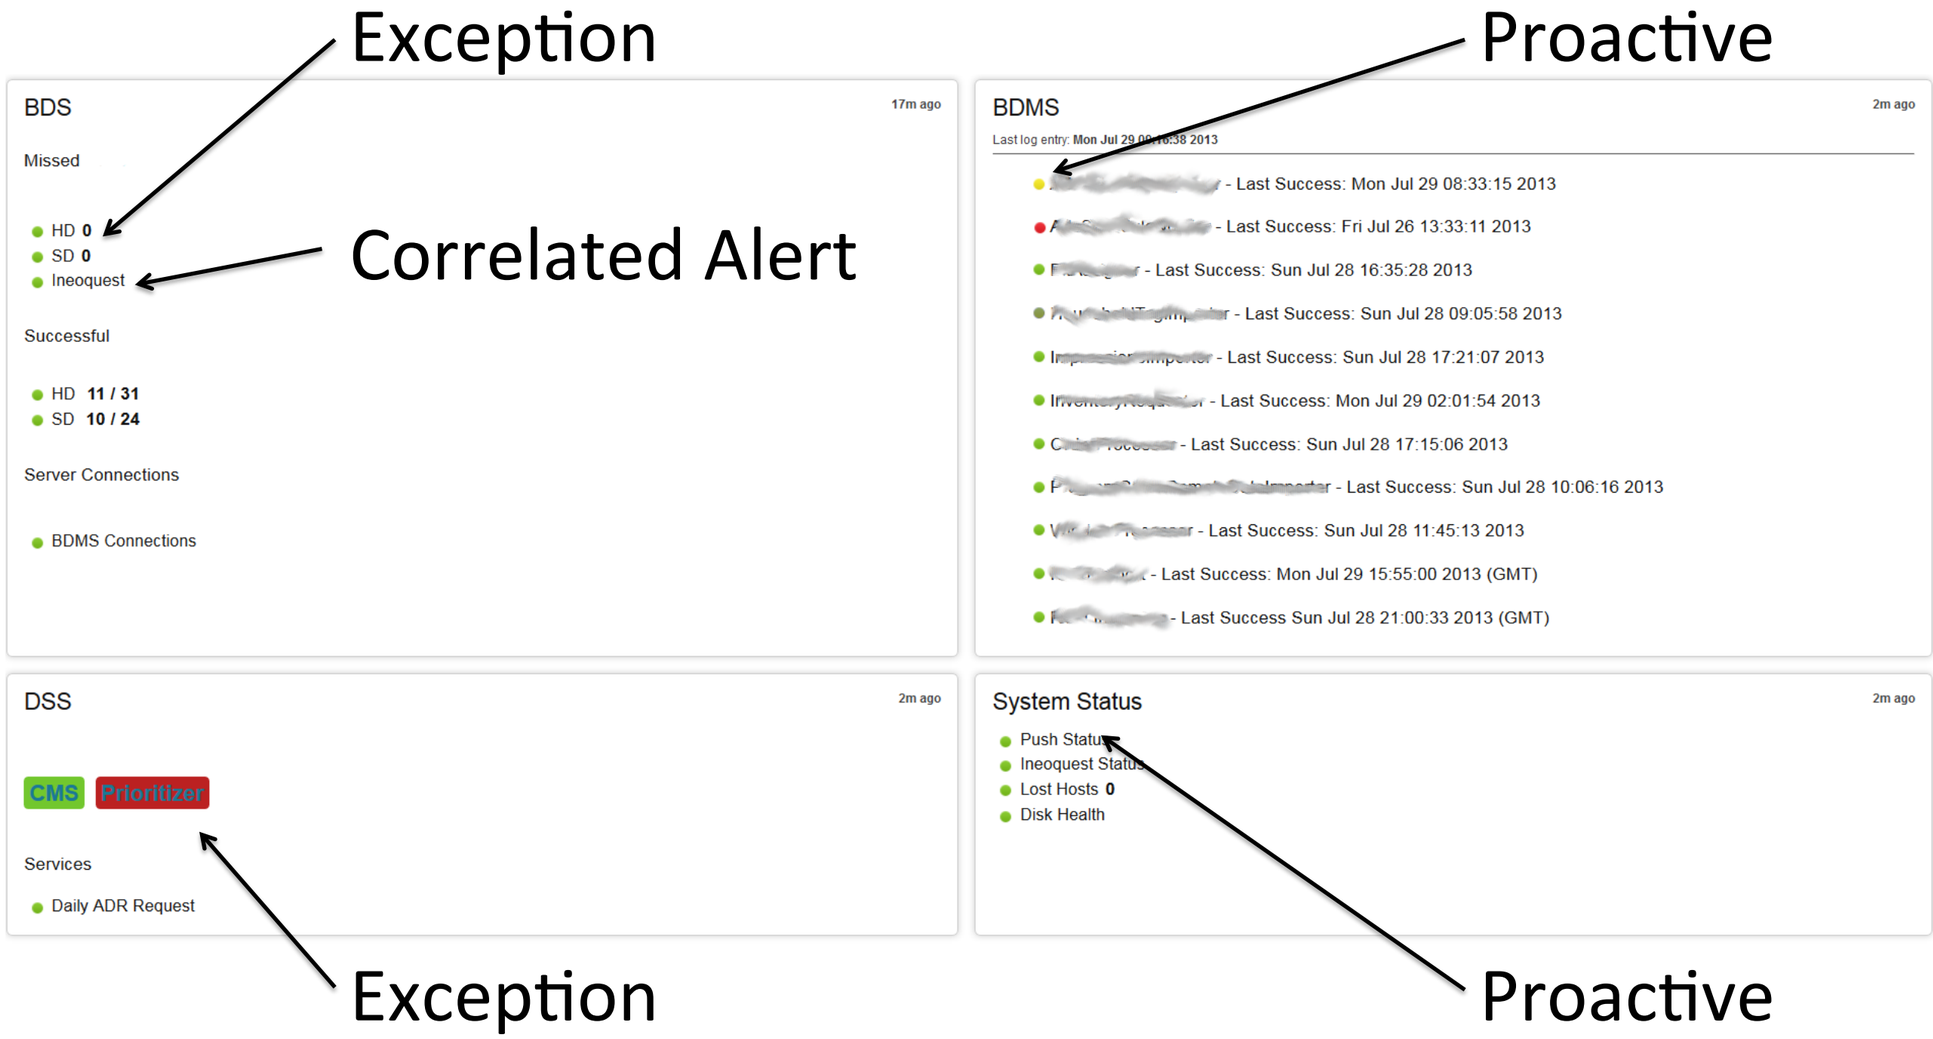
Exception
Proactive
BDS
17m ago
Missed
HD 0
SD 0
Ineoquest
Successful
HD 11 / 31
SD 10 / 24
Server Connections
BDMS Connections
Correlated Alert
BDMS
2m ago
Last log entry: Mon Jul 29 08:33:38 2013
- Last Success: Mon Jul 29 08:33:15 2013
- Last Success: Fri Jul 26 13:33:11 2013
- Last Success: Sun Jul 28 16:35:28 2013
- Last Success: Sun Jul 28 09:05:58 2013
- Last Success: Sun Jul 28 17:21:07 2013
- Last Success: Mon Jul 29 02:01:54 2013
- Last Success: Sun Jul 28 17:15:06 2013
- Last Success: Sun Jul 28 10:06:16 2013
- Last Success: Sun Jul 28 11:45:13 2013
- Last Success: Mon Jul 29 15:55:00 2013 (GMT)
- Last Success Sun Jul 28 21:00:33 2013 (GMT)
DSS
2m ago
CMS
Prioritizer
Services
Daily ADR Request
System Status
2m ago
Push Status
Ineoquest Status
Lost Hosts 0
Disk Health
Exception
Proactive

# Challenge: Not So Nerdy Users

- Present data to users who don't understand the workings of the system
- Don't call it what it is

| **What the log says** | **What user understands** |
|---|---|
| Timeout | Network connectivity |
| Forwarder agent is down | Lost Server |

# Easy to Understand UI

- Start with high level… drilldown to confusion
- Incorporate workflows
- Do everything you can from same interface
- Always present in plain English, use choice words
- Colors over numbers (or colored numbers)
- Separate hardware/network alerts from application alerts

# Drilldown Example

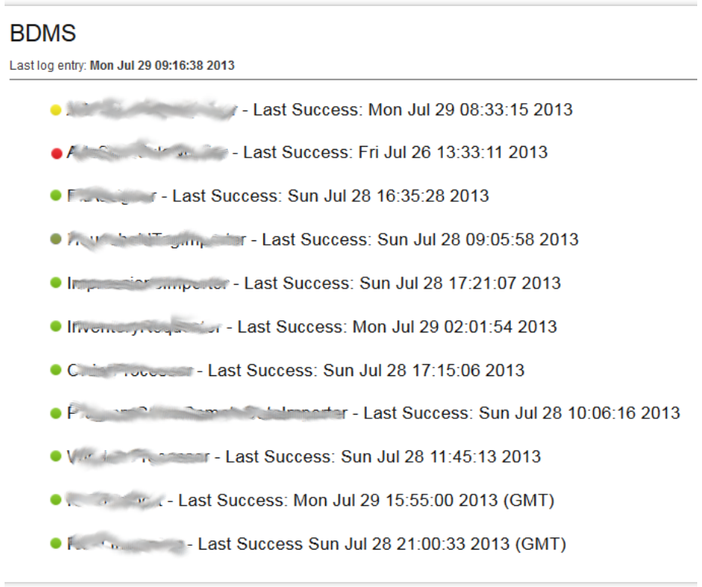


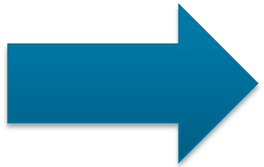


Search:

« prev 1 2 3 4 5 6 next »

| time | type | pid | message |
|---|---|---|---|
| 09/13/13 09:36:18.558 AM | Info | 6456 | |
| 09/13/13 09:36:18.543 AM | Info | 6456 | |
| 09/13/13 09:36:18.527 AM | Info | 6456 | Import successfu |
| 09/13/13 09:36:00.914 AM | Error | 6456 | IMPORT FAILED: |
| 09/13/13 09:36:00.586 AM | Error | 6456 | IMPORT FAILED: |
| 09/13/13 09:36:00.336 AM | Error | 6456 | IMPORT FAILED: |
| 09/13/13 09:35:59.978 AM | Error | 6456 | IMPORT FAILED: |
| 09/13/13 09:35:59.759 AM | Error | 6456 | IMPORT FAILED: |
| 09/13/13 09:35:59.447 AM | Error | 6456 | IMPORT FAILED: |
| 09/13/13 09:35:59.197 AM | Error | 6456 | IMPORT FAILED: |

# Workflow Example

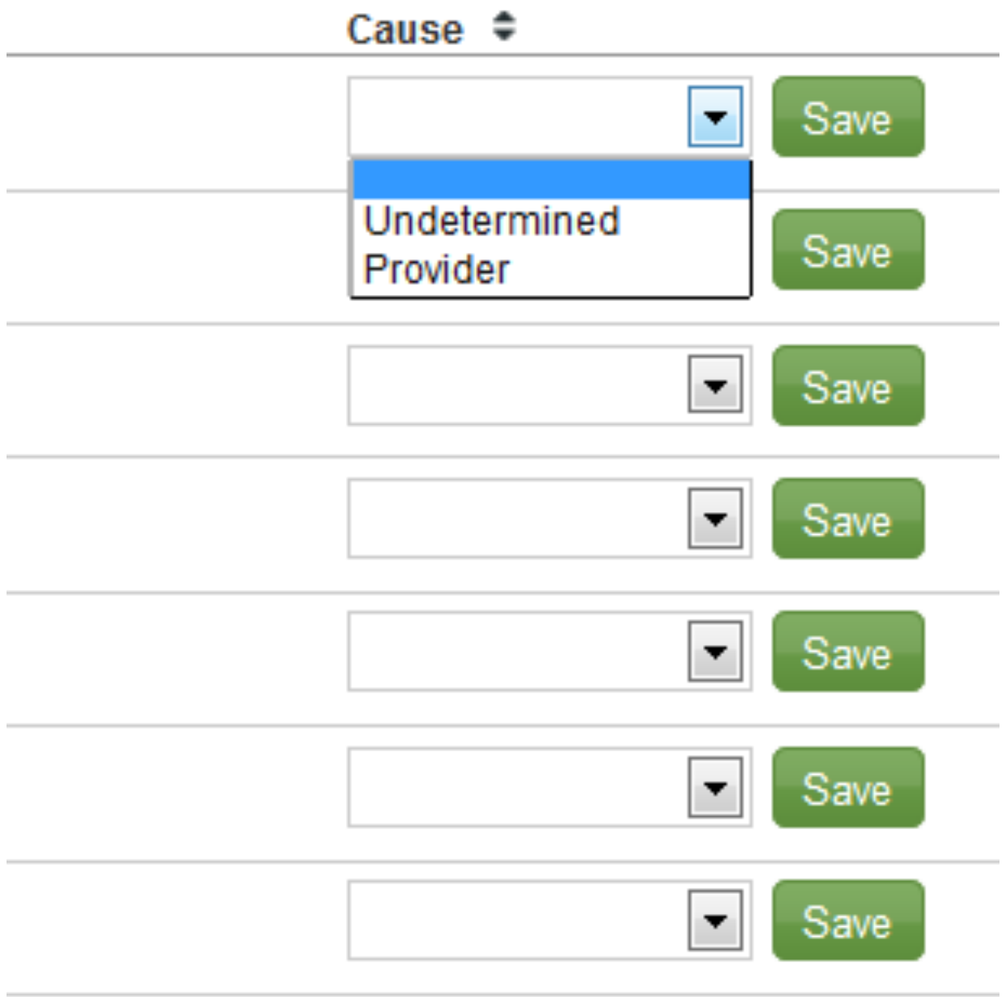

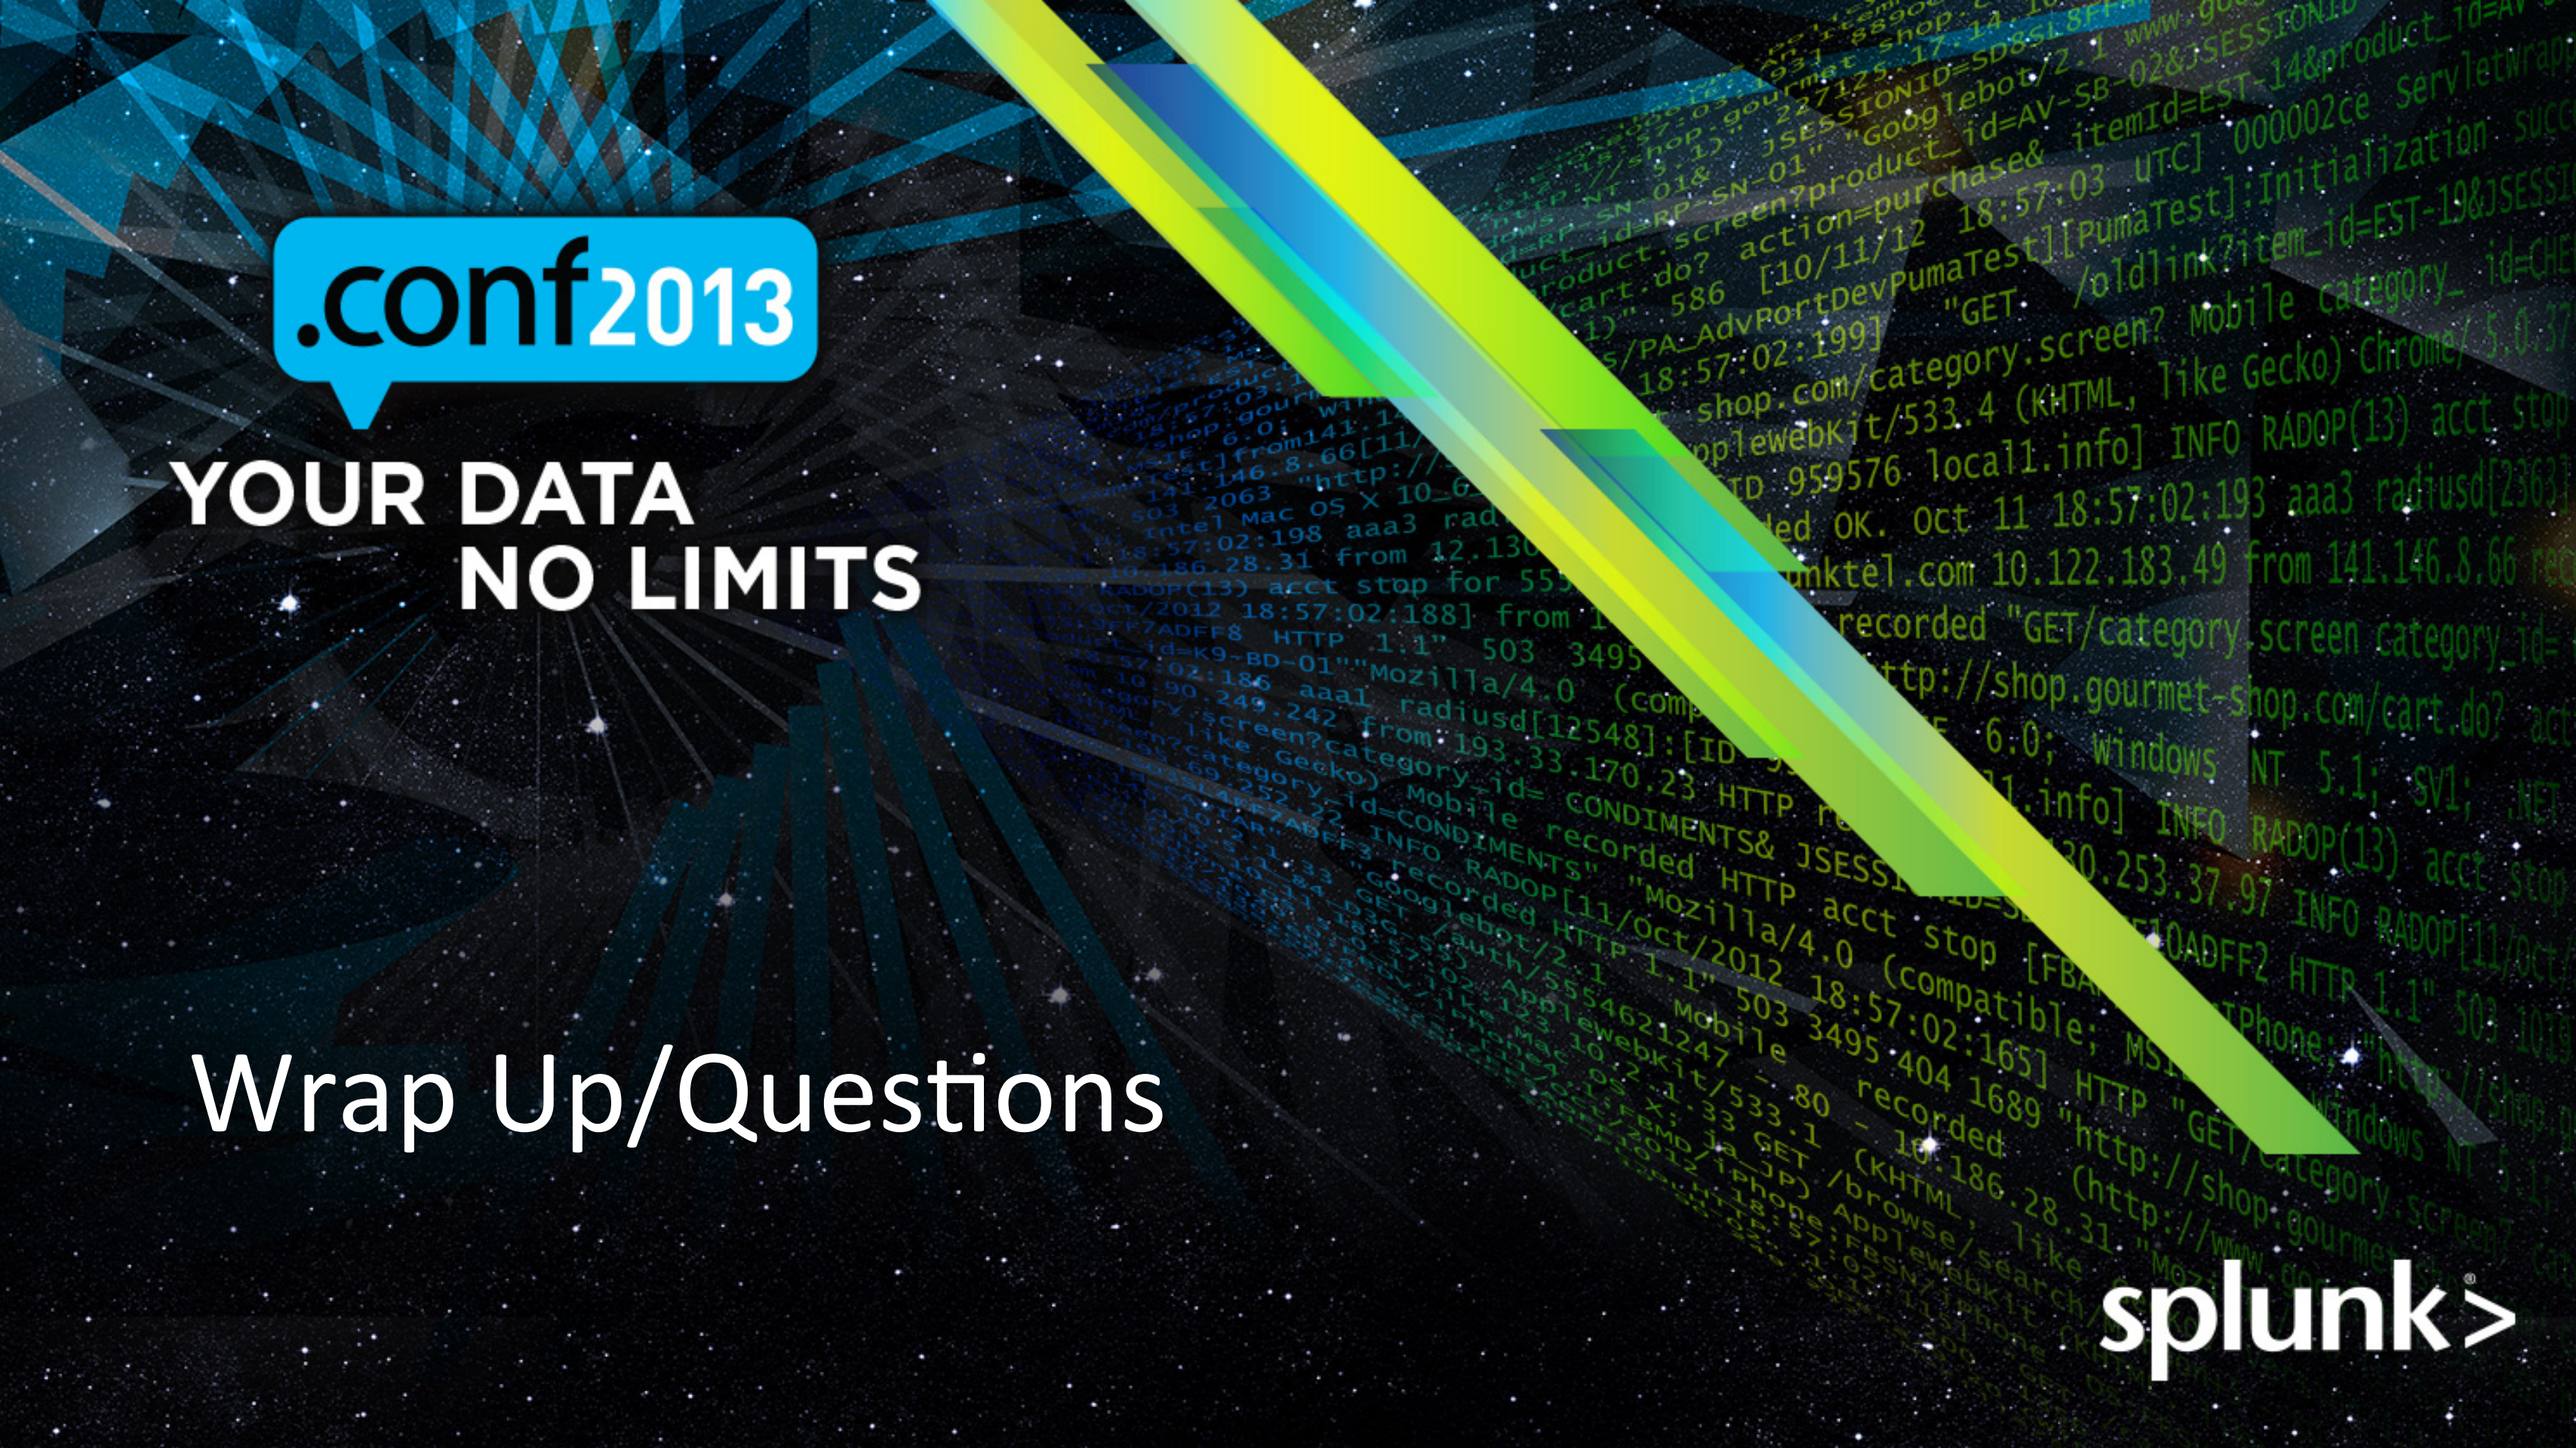
.conf2013
YOUR DATA
NO LIMITS
Wrap Up/Questions
splunk>

# 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!


.conf2013
YOUR DATA
NO LIMITS
THANK YOU
splunk>