

Writing Actionable Alerts

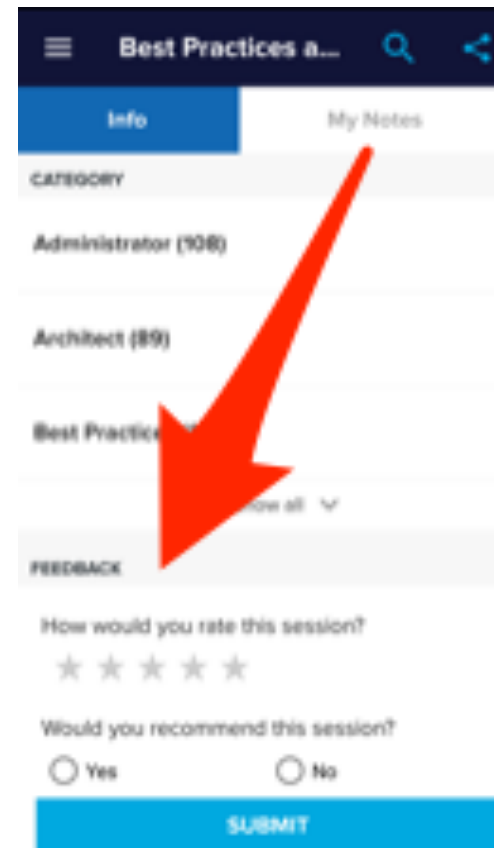
While you get settled...

Download Latest Slides:

<https://splunk.box.com/v/burch-alerts>

or ask a neighbor with flash drive

Load Feedback:



The screenshot shows a mobile application interface for 'Best Practices a...'. It has a dark blue header with a menu icon, a search icon, and a share icon. Below the header are two tabs: 'Info' (selected) and 'My Notes'. The main content area lists categories: 'CATEGORY', 'Administrator (108)', 'Architect (89)', and 'Best Practice'. A red arrow points from the 'FEEDBACK' section to the 'My Notes' tab. The feedback form includes a question 'How would you rate this session?' with five stars, and 'Would you recommend this session?' with 'Yes' and 'No' radio buttons. A blue 'SUBMIT' button is at the bottom.

Related Blog:

Search: blogs.splunk.com writing
actionable alerts

<http://blogs.splunk.com/2016/01/29/writing-actionable-alerts>

Writing Actionable Alerts

Burch

Sales Engineer @ 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.

What's a Burch?

- Senior Sales Engineer in Boston
- Education
 - CS @ Boston University
 - MBA @ Northeastern University
- Splunk Customer
 - Middleware for 8 years (+splunk)
 - **Splunk Admin for 1.5 years (splunk 4.3+)**
- Certs: Knowledge, Admin, Architect
- @Splunk since Dec 14
- Splunkbase apps



About you

- Name
- User?
- Power User?
- Admin?
- Groupie?



Burch's Goal

From spam to glam with Splunk Alerts

eval Agenda = “Maturity Model”

- Stage 1: Message of Concern
- Stage 2: Thresholds
- Stage 3: Relative Percentages
- Stage 4: Average Errors
- Stage 5: Percentiles
- Bonus Stage 6: IT Service Intelligence
- Stage 7: Actionable Alerts

Stage 1

Message of Concern

Attempted Solution

- Created spammy alert:

```
[Spam]
action.email = 1
action.email.priority = 1
action.email.to = welovespam@spam.com
counttype = number of events
cron_schedule = */15 * * * *
dispatch.earliest_time = -15min
dispatch.latest_time = now
enableSched = 1
quantity = 0
relation = greater than
search = index=_internal error
```

- Weak search definition

Result

index=_internal error

6,500 events (7/5/16 12:53:20.000 PM to 7/5/16 1:08:20.000 PM) No Event Sampling

Events (6,500) Patterns Statistics Visualization

Format Timeline Zoom Out Zoom to Selection Deselect

1 minute per column

6,500 errors over last 15 min

unrealistic to act upon

```
2016-07-05 13:08:00,067 ERROR pid=2084 tid=MainThread file=util.py:__call__:153 | Failed to execute function=run, error=Traceback (most recent call last):
File "/opt/splunk/etc/apps/Splunk_TA_aws/bin/splunktalib/common/util.py", line 150, in __call__
return func(*args, **kwargs)
File "/opt/splunk/etc/apps/Splunk_TA_aws/bin/aws_cloudwatch.py", line 92, in run
acconf.AWSCloudWatchConf, "aws_cloudwatch", logger)
host = wicket ; source = /opt/splunk/var/log/splunk/splunk_ta_aws_cloudwatch_main.log ; sourcetype = aws:cloudwatch:log
```


Obvious Improvements

- Scope of problem is large
 - Solution: indexed fields (index, source, sourcetype, and/or pattern)
- Problem: “error” matches more than desired
 - Solution: bind with fields like log_level=“error”
- Result: Stronger search ignores benign results

– `index=_internal sourcetype=splunkd source!="*splunkforwarder*" log_level=ERROR`

Stage 2

Thresholds

Attempted Solution

- Only alert if more than “arbitrary” # occurrences / time
 - Arbitrary = perception of healthy

```
index=_internal sourcetype=splunkd source!="*splunkforwarder*" log_level=ERROR  
| stats count  
| where count>20
```

or...

Alert

Condition


if number of events

is greater than

20

Result & Obvious Improvements

- Ignores volume variances of different types of errors
 - Web errors rarely happen but server errors happen often
- Fluctuations relative to usage
 - Threshold too small or large during peak or minimal usage, respectively
 - Static thresholds not adjusting with business growth or decline



Stage 3

Relative Percentages

New Concept

```
eval goal_attacking = coalesce( spam, system )
```

SPAM

- Normalize against # of errors
- Ignore non error events
- log_level=ERROR

- Good for clean up
- Bad for permanent


SYSTEM


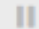






- Normalize to all events
- Include all error + non error events
- log_level=*

- Good for permanent
- Bad for clean up



Attempted Solution

```
index=_internal sourcetype=splunkd source!="*/splunkforwarder/*" log_level=*
| stats count, count(eval(log_level=="ERROR")) AS error_count by component
| where ( error_count / count ) > .50
```

Last 15 minutes 

✓ 303,891 events (7/5/16 3:54:42.000 PM to 7/5/16 4:09:42.000 PM) No Event Sampling        Smart Mode 

Events Patterns **Statistics (2)** Visualization

100 Per Page  Format  Preview

component	count	error_count
DeployedServerclass	936	936
ExecProcessor	488	486

Result & Obvious Improvements

- Huge improvement
 - Less spam
 - Adjusts because normalized to volume
- What if that's normal?
 - Then persistent alerts that should be ignored = spam + noise!
- Arbitrary static percentages

Stage 4

Average Errors

Attempted Solution

Current period vs historical average

```
index=_internal sourcetype=splunkd source!="*/splunkforwarder/*" log_level=ERROR
| bin span=5min _time
| stats count by _time, component
| stats latest(count) as current_count, avg(count) as historical_count by component
| where current_count > historical_count
```

Last 7 days ▾



✓ 619,072 events (7/19/16 4:00:00.000 PM to 7/26/16 4:58:23.000 PM)

No Event Sampling ▾

! Job ▾



💡 Smart Mode ▾

Events

Patterns

Statistics (13)

Visualization

10 Per Page ▾

✎ Format ▾

Preview ▾

< Prev

1

2

Next >

component	current_count	historical_count
AdminHandler:PersistMessages	16	8.500000
Application	30	10.000000
ApplicationUpdater	77	27.333333
ArchiveContext	6	4.666667

Result


- Adjusts with changes in environment!
- Slow
 - Summary Indexing?
 - Acceleration?
- How often alert?
 - Definition of average!

Hold Up!

Statistics Detour

Statistics Detour

Historical # of errors / 5 min period



Statistics Detour

11

87

56

31

5

19

67

21

18

77

54000

0 to 10 | 11 to 20 | 21 to 29 | 30 to 39 | 40 to 49 | 50 to 59 | 60 to 69 | 70 to 79 | 80 to 89 | 90 to 99

Statistics Detour

At what value does this become actionable?

Min
Average
Max

18

19

5

11

21

31

56

67

77

87

0 to 10 | 11 to 20 | 21 to 29 | 30 to 39 | 40 to 49 | 50 to 59 | 60 to 69 | 70 to 79 | 80 to 89 | 90 to 99

Statistics Detour

What if we could skim off outliers?

Alert at *near* max?

18

19

5

11

21


31

56

67

77

87



0 to 10 | 11 to 20 | 21 to 29 | 30 to 39 | 40 to 49 | 50 to 59 | 60 to 69 | 70 to 79 | 80 to 89 | 90 to 99

Statistics Detour

`perc<X>(Y)` = Returns the X-th percentile value of the numeric field Y, where X is an integer between 1 and 99. The percentile X-th function **sorts the values** of Y in an increasing order. Then, if you consider that 0% is the **lowest** and 100% the **highest**, the functions picks the **value that corresponds to the position** of the X% value.

18

19

5

11

21


31

56

67

77

87



0 to 10 | 11 to 20 | 21 to 29 | 30 to 39 | 40 to 49 | 50 to 59 | 60 to 69 | 70 to 79 | 80 to 89 | 90 to 99

Statistics Detour

`perc90(this_result_set) = ?`

18

19

5

11

21


31

56

67

77


87








0 to 10 | 11 to 20 | 21 to 29 | 30 to 39 | 40 to 49 | 50 to 59 | 60 to 69 | 70 to 79 | 80 to 89 | 90 to 99

Statistics Detour


```
| makeresults count=11
| streamstats count
| eval count = case( count == "11" , "18" , count == "1" , "5" , count == "2" , "11" , count == "3" ,
"19" , count == "4" , "21" , count == "5" , "31" , count == "6" , "56" , count == "7" , "77" , count
== "8" , "87" , count == "9" , "54000" , count == "10" , "67")
| stats perc90(count)
```

Last 15 minutes 

✓ 1 result (7/31/16 2:46:51.000 PM to 7/31/16 3:01:51.000 PM) [No Event Sampling](#) [Job](#)      [Smart Mode](#)


Events | Patterns | **Statistics (1)** | Visualization

[10 Per Page](#) [Format](#) [Preview](#)

perc90(count) 


87

Warning: Assumption



Shout out to Xander!

Warning: Heavy Tails




Warning: Reality



What percentile is appropriate given this distribution?

Know Thy Data

```
index=_internal sourcetype=splunkd
source!="*/splunkforwarder/*"
| bin span=5min _time
| stats count AS group by _time
| bin span=1000 group
| stats count by group
| sort group
```



Stage 5

Percentiles

Attempted Solution

- Current period's error rate vs. historical error rate
 - by error category (component)

```
index=_internal sourcetype=splunkd source!="*/splunkforwarder/*"  
log_level=ERROR  
  | bin span=5min _time  
  | stats count by _time, component  
  | stats perc95(count) AS perc95_count, latest(count) AS  
current_count by component  
  | where current_count > perc95_count
```

- Performance?

Summary Indexing Solution

- **Generate malleable historical data**

```
index=_internal sourcetype=splunkd source!="*/splunkforwarder/*"  
log_level=ERROR  
| bin span=5min _time  
| sistats count by _time, component
```

- **Alert upon historical data**

```
index=summary_internal sourcetype=stash source="my search name"  
| stats count by _time, component  
| stats perc95(count) AS perc95_count, latest(count) AS  
current_count by component
```

The Lasso Approach

- Triage Strategy
- Perimeter around errors
- Tighten lasso by reducing percentile
- Rinse & repeat



Alternatives

- Address most common errors first
 - Start at 5th percentile and work up
- Normalization Frames:
 - Same errors
 - All errors
 - All events
 - Time windows (e.g. work hours)

Result

- Adjusts with changes in environment!
- Requires Maintenance
 - Power User skillz
 - Summary Indexing
- Not period time adjusted
 - Fluctuations in business day or period

Bonus Stage 6


IT Service Intelligence

Why ITSI?

Make **alerting** accessible, usable and valuable to everyone!

Quantile, Range, and STDDDEV. Oh my!

Preview Aggregate Thresholds




Configure Thresholds for Time Policies

Weekdays, 12AM-8AM	Policy type? Quantile
Weekdays, 6PM-12AM	Thresholds are computed from data. Parameter associated with the labels is the quantile value between 0 and 1. 0.25 would equal the 25th percentile of the data, 0.5 would be the median or 50th percentile, and 0.75 would be the 75th percentile
Weekdays, 8AM-6PM	<input type="checkbox"/> Critical 0.9
Weekends	<input type="checkbox"/> Medium 0.75
Default	<input type="checkbox"/> Normal 0.5
	<input type="checkbox"/> Normal


View data from **Last Thursday** between **7:00 - 8:00 hours**

The zoomed-in chart shows a data series for a specific time period. The y-axis labels are 1203.2, 820, 555, and 0. The data points fluctuate around the thresholds.

Adaptive Thresholds




Anomaly Detection



Stage 7

Actionable Alerts

Actionable Alerts Made Easy



What Now?

Related breakout sessions and activities...

- Rate this! (be honest)
- More talks:
 - conf.splunk.com/speakers.html
 - Search for
 - ▶ Burch
 - ▶ Champagne
 - ▶ Optimization
 - ▶ Practices
 - ▶ tips
 - ▶ Worst



Free Discussion

Questions, ideas, experiences
...have you?



THANK YOU

.conf2016

