Security Ninjutsu Part Five – .conf18, Orlando Florida

Search Techniques Used

stats+eval

```
… | stats values(eval(if(result=1, user, null))) as successful_users
```

Summary Indexing

```
… | stats count dc(…) values(…) […] by user
  | collect index=summary
```

- Prime Directive of Summary Indexing: anything before "by" is nearly free

First Time Seen Detections

```
… | stats min(_time) max(_time) by user
  | where 'min(_time)' > relative_time(now(), "-1d@d")
```

Rare Events Detections

```
… | stats count latest(_time) as latest by user
  | eventstats sum(count) as total
  | where count / total > 1 / 20000 AND
              latest > relative_time(now(), "-1d@d")
```

Variations of First Time Seen

tstats, peer groups, lookup cache

Notable of Notables

```
index=risk earliest=-30d
| stats sum(risk) as current_risk
        sum(eval(if(_time<relative_time(now(), "@d"))) as old_risk
| where current_risk > 500 OR old_risk > 3000
```

Recommended Resources

.conf18 Talks

- SEC1479 - Say Goodbye to Your Big Alert Pipeline, and Say Hello to Your New Risk-Based Approach
- SEC1583 - Turning Security Use Cases Into SPL
- SEC1039 - Detection Technique Deep Dive
- SEC1355 - Hunting the Known Unknown: Microsoft Cloud
- SEC1674 - From Threat Modelingto Automated Response
- SEC1378 - Splunking the Endpoint IV: A New Hope

Earlier Ninjutsu Talks

(Particularly part four)
https://www.davidveuve.com/splunk.html

Splunk Hunting Blog Series

https://www.splunk.com/blog/2017/07/06/hunting-with-splunk-the-basics.html

How Search Works

https://www.davidveuve.com/searchdeepdive

Techniques Detailed in Security Ninjutsu Part Four

Recommended reading for anyone deeply interested in advanced searching.

Each term below refers to a technique header in the PDF

Apps

Splunk Security Essentials

URL Toolbox

Splunk ES Content Update

Intermediate

- Common Information Model
- eval
- Multi-Value Fields
- stats
- stats on stats
- Formatting a Table
- Multi-Scenario Alerts
- Inline Comments
- Tuning
- Stats+eval
- Override Urgency / Severity / Risk
- Common Apps
- Risk
- Subsearches

Advanced

- Summary Indexing
- Lookup Caching
- Confidence Checking
- Managing Alert Fatigue
- Transaction (in a good way)
- First Time Seen Detection
- Time Series Detection
- Time Series + First Time Seen Detection

Ninja

- tstats
- Timestamps & Timestamps
- Advanced Search Commands
- Metacharacteristics
- Machine Learning Toolkit Numeric Clustering
- Approach to Analytics

End-to-End

- When Log Sources Go Quiet

Types of Correlation Events

Alerts

High Fidelity, Actionable, Automatable

These are aligned to traditional correlation searches whose results you would review.

In the future, expect most of these alerts to be automated through tools like Phantom, as you drive your SOC operational maturity.

Risky Events

Can be aggregated, and provide context

Most behavioral searches fall into this category. The basic premise for risky events are those that you would look at and say "that's interesting" but wouldn't send on their own to the SOC. New API calls, logins from different countries, etc.

Contextual Events

Will never generate a SOC ticket, but provide a useful head start

These alerts can come from any inherently noisy sources, such as behavioral searches or just low fidelity indicators. Because of their volume and the correspondingly low confidence, they're not a part of any ticket, but they may provide a good starting point for an investigation.

Getting Started Plan

1. Get Familiar with First Time Seen, and Time Series Detections
   - Look through Splunk Security Essentials, and pick a couple of your low volume data sources (e.g., printer logs, cloudtrail for specific APIs, etc. – effectively that you can search over 30-90 days without waiting an hour for the search to complete)
   - Grab a couple of the first time seen, and time series detections to get comfortable with them. Just don't expect particularly useful results yet, you're only getting used to the general format.
2. Build out a first draft dataset on your dataset with 4-8 meaningful fields, but don't worry about acceleration or complexity.
3. Run one or two behavioral searches on this dataset – don't worry if it takes a long time, the goal here is understand the premise and start seeing some results. If you're looking for threshold based detections (such as rarity but not first time seen, or standard deviation based time series analysis, set that bar low so that you see something.
4. Work on accelerating this dataset. Leverage the techniques discussed both in this slide deck, and Security Ninjutsu Part Four.
5. Leverage the analytic patterns to build out 5-10 different analytics that you find casually interesting. Don't send these to the SOC, just report them somewhere, such as in a dedicated summary index that you use to test. (End with something like | eval search="mySearchName" | addinfo | collect index=myTestIndex)
6. If any of those analytics are high fidelity enough to send directly to the SOC, do so. Use whatever structure you already have in place for correlation searches.
7. For the others, build out some of the risk-based analytics discussed in step three to detect aggregations. Send these alerts to yourself to begin with, before passing along to the SOC. Evaluate whether they seem to be legitimate.
8. Review the aggregated alerts that are being generated, and be prepared to tweak the thresholds.
9. Keep an eye for any behavioral detections that don't alert at all, or that create a huge volume of alerts, and be prepared to tweak those individual alerts.

Notes