

Forward-Looking Statements



This presentation may contain forward-looking statements regarding future events, plans or the expected financial performance of our company, including our expectations regarding our products, technology, strategy, customers, markets, acquisitions and investments. These statements reflect management's current expectations, estimates and assumptions based on the information currently available to us. These forward-looking statements are not guarantees of future performance and involve significant risks, uncertainties and other factors that may cause our actual results, performance or achievements to be materially different from results, performance or achievements expressed or implied by the forward-looking statements contained in this presentation.

For additional information about factors that could cause actual results to differ materially from those described in the forward-looking statements made in this presentation, please refer to our periodic reports and other filings with the SEC, including the risk factors identified in our most recent quarterly reports on Form 10-Q and annual reports on Form 10-K, copies of which may be obtained by visiting the Splunk Investor Relations website at www.investors.splunk.com or the SEC's website at www.sec.gov. The forward-looking statements made in this presentation are made as of the time and date of this presentation. If reviewed after the initial presentation, even if made available by us, on our website or otherwise, it may not contain current or accurate information. We disclaim any obligation to update or revise any forward-looking statement based on new information, future events or otherwise, except as required by applicable law.

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. We undertake no obligation either to develop the features or functionalities described or to include any such feature or functionality in a future release.

Splunk, Splunk>, Data-to-Everything, D2E and Turn Data Into Doing are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. ©2021 Splunk Inc. All rights reserved.

Supercharge Your Risk Based Alerting (RBA) Implementation

SEC1162A

Teresa Chila

Cybersecurity Data Scientist | Chevron

This document is intended only for use by Chevron for presentation at .conf21 and inclusion by Splunk on a conference website that is available to the public. No portion of this document may be copied, displayed, distributed, reproduced, published, sold, licensed, downloaded, or used to create a derivative work, unless the use has been specifically authorized by Chevron in writing.





Teresa Chila

Cybersecurity Data Scientist | Chevron

This document is intended only for use by Chevron for presentation at .conf21 and inclusion by Splunk on a conference website that is available to the public. No portion of this document may be copied, displayed, distributed, reproduced, published, sold, licensed, downloaded, or used to create a derivative work, unless the use has been specifically authorized by Chevron in writing.

Making Your SIEM Smarter

This session offers tips to make your RBA deliver more value

Reduce Alert Fatigue



Spend time on meaningful signals

Expand Detection Coverage



No more ignoring high volume alerts. Make the noise work for you!

Increase Analytic Time Range & Boundary



Ability to look across much larger time range for low-and-slow and across multiple data sources

Agenda


- 1) Super Quick RBA Overview
- 2) How to Decide on Alert Threshold
- 3) How to Combat Repeating False Positives
- 4) Integration with SOAR
- 5) Key Takeaways



Super Quick RBA Overview


Traditional Notables:

- Represents one type of detection/anomaly
- Very atomic, lack context
- High volume detections are usually ignored
- Varying degrees of fidelity



RBA Notables:

- Composed of **one or more Risk Events** (each risk event is a detection/anomaly)
- Correlate many risk events based on host or user
- Generate a notable when the composite Risk Score or MITRE ATT&CK chain exceeds a threshold



RBA pipeline



Impact of RBA

Do more for less

↓ **800x**


800 times reduction
of raw alerts to
notables

↑ **200**

Monitor 200 more
detection types

0

Zero staff increase




Better employee
morale

The SOC Analyst worked one notable with 3 related risk events instead of 3 different alerts.

Without RBA, these 3 alerts may be worked by separate analysts, which means less context and duplication of effort.



Type of risk events in RBA



Risk score calculation

Based on domain knowledge + historical effectiveness

How likely is this a malicious event?



risk_score = risk_rule_**impact**_num * risk_rule_**confidence**_num * ((**risk_modifier_count** * .25)+1)



If this event is malicious, how bad the damage is?



Based on criticality of the entity (system or user)

Is this a critical asset?

Is this an admin user?



How to Decide on Alert Threshold

Are we alerting too much?
Are we alerting too little and missing crucial alerts?

splunk>

.conf21

Non-ideal RBA notable

Is this an “investigation worthy” notable?

Score is inflated due to event 2 and 4 duplication. They are really the same event.

Event 1 shouldn't be there

Event 3 provides context only

An RBA notable consists of the following risk events (risk_score_sum “108” > threshold “100”):

- Risk Event 1 (risk_score=60) → medium confidence, but false positive, don't care, not a concern
- Risk Event 2 (risk_score=20) → low confidence
- Risk Event 3 (risk_score=8) → informational
- Risk Event 4 (risk_score=20) → similar to Event 2 but different alert type, essentially referencing the same event

Solutions

Manual hard work is still needed

Metadata tagging -- Initial work to assign **confidence** and **impact** rating to individual risk event is **crucial**

- Use historical alert volume to help with your decision
- Confidence and impact can be dynamic and context sensitive
 - `eval impact=if(match(action, "allowed"), "high", "low")`

Still have to **tune individual risk rule** to minimize false positives

Non-ideal RBA notable

Example 2

Multiple occurrences of the same risk event type, all low confidence

An RBA notable consists of the following risk events (threshold “100”):

- 10 occurrences of same Risk Event type → low confidence
 - risk_score “20” * 10 = 200

Non-Monolithic Approach to Threshold

More knobs you can turn

Model & MITRE Tactic Count

modelCount: number of distinct
detection type e.g. `dc(search_name)`

tacticCount: number of distinct Mitre
Tactic type e.g. `dc(tactic)`

| where modelCount>1 AND tacticCount>1

Stepwise Risk Score Threshold

```
eval threshold = case(  
  match(confidence,"high"), 80,  
  match(confidence,"medium"), 100,  
  1==1, 200)
```

Use `match()` because
confidence is a multi-value
field aggregating multiple
events' value



How to Combat Repeating False Positives

Dynamic Suppression/Throttling based on Prior Verdict

splunk>

.conf21

How to Combat Repeating False Positives

Cause:

- When FP appears, you didn't immediately create a rule to suppress it, or tune it in the search, or put in a whitelist
- Why?
 - It's someone else's job, don't know how, maybe it's just a one-time thing, etc.
- Out-of-the-box ES Throttling does not take prior verdict into account



"SIEM Babysitting"

Solution:

- Augment the RIR SPL to look at prior notable verdict to allow for automatic suppression and more flexible throttling

Pre-requisites

Having standards that all Risk Rules adhere to is important

Add ES custom Incident Review status:

- Closed - True Positive (101)
- Closed - False Positive (102)

Establish standard fields:

- risk_object
- detection_type (or search_name)
- threat_object

Dynamically suppress new notable if:

- Old notable is False Positive, and/or
- Old notable occurred within last x time period, and
- Old notable references the same risk object, detection type, and threat object (your choice)

SPL

Can be saved as a macro `rir\_suppress` to be included in the RIR

...<code to generate RBA notable>...

|multireport

[| where 1==1

multireport part 1: get a copy of all records

Join multi-value into one line for ease of grouping

| fillnull value="n/a" threat_object | eval type_1line=mvjoin(detection_type,"|"), threat_1line=mvjoin(threat_object,"|")]

[| stats values(risk_object) as risk_object

| eval search_str="(risk_object=".mvjoin(risk_object," OR risk_object=").")"

multireport part 2: format search string e.g. risk_object=A OR risk_object=B OR risk_object=C

| fields search_str

| map

macro in ES

[search earliest=-30d [| stats count | eval query=\$search_str\$ | return \$query] source=*RIR* `notable`

| fillnull value="n/a" threat_object | eval type_1line=mvjoin(detection_type,"|"), threat_1line=mvjoin(threat_object,"|")

| eval time_ago_seconds=now()-_time

| stats min(time_ago_seconds) values(detection_type) values(threat_object) values(status) by risk_object type_1line threat_1line

| rename values(*) as * | rename min(*) as *

Suppress based on time

| eval disposition1=if(time_ago_seconds<(60*60*24*1),"suppress",null())

| eval disposition2=case(match(status,"True Positive"), "keep", match(status,"False Positive"), "suppress", status="New", "keep")

| eval disposition=coalesce(disposition1,disposition2)

Suppress if prior notable is False Positive

| fields risk_object type_1line threat_1line disposition]]

| stats values(min_time) values(risk_message) values(risk_score_sum) values(modelCount) values(detection_type) values(tactic)

values(technique) values(threat_object) values(disposition) values(custom_fields) by risk_object type_1line threat_1line

| rename values(*) as *

| search (NOT disposition="suppress") AND risk_message=*

Keep the original record (with risk message) where disposition is not suppress

Last stats joins part 1 & 2 together based on risk object, detection type, and threat object

.conf21

Use map to execute subsearch on prior notables



Integration with SOAR

splunk>

.conf21


How to process RBA notables in Splunk SOAR

Challenge:


- Many types of different risk events in a RBA notable, very diverse

Solution:

- Assign each detection type to a response category
- Create SOAR sub-playbooks for each response category
- A master RBA playbook that directs the calls to sub-playbooks



RBA Notables generated in ES




RBA SOAR Playbook



Example 1

SOAR


RBA Notable



Example 2

SOAR

RBA Notable



Run sub-playbooks:

Endpoint File &
Process
Response

Suspicious Logon
Response

Network Indicator
Response

What's in a response sub-playbook?

Endpoint File & Process

- Is the file still there?
- Collect file sample
- Process run history
- Containment options:
 - Kill the process
 - Stop & quarantine the file

Network Indicator

- Reputation checks
- Have we seen this indicator before?
- Who else has accessed this domain/IP/URL
- Containment options:
 - Block the indicator

Key takeaways

- 1) Metadata tagging and tuning are important
Spend time assigning confidence and impact rating, as well as tune individual detections
- 2) Establish standard fields in Risk Events
This will enable you to be creative in your RIR SPL such as Dynamic Suppression
- 3) Don't be intimidated by processing RBA notables in Splunk SOAR



Risk-based alerting (RBA) at .conf21

SEC1249A - Accenture's Journey to RBA with Splunk Enterprise Security and Beyond

SEC1163A - Proactive Risk Based Alerting for Insider Threats

SEC1162A - Supercharge Your Risk Based Alerting (RBA) Implementation

SEC1271 - What's New in Splunk Enterprise Security?

SEC1590C - Augmented Case Management With Risk Based Analytics and Splunk SOAR

SEC1800A - Implementing Zero Trust: From Hype to Reality

SEC1466A - A Deep-Dive Into How Zoom Is Building Its World-Class Detection Pipeline in Response to the Zoom-Boom!



Thank You

Please provide feedback via the
SESSION SURVEY

