# Risk-Based Alerting Helps SOCs Focus on What Really Matters

## Splunk's Jesse Trucks on Latest Tools for Reducing False Positives, Ticket Fatigue

**Jesse Trucks**

*Trucks has been helping to make the magic at Splunk for the past six years. He has over 20 years of experience in IT and security operations.*

Detection tools can potentially overwhelm security operation center analysts with alerts, many of which are false positives, leading to ticket fatigue and missed attacks. **Jesse Trucks**, minister of magic at Splunk, says the latest risk-based alerting technology is helping SOCs focus on the threats that really matter.

In this interview with Information Security Media Group, Trucks discusses:

- The common challenges with alerts that security operations teams and analysts face;
- How risk-based alerting works to reduce false positives and create more high-fidelity tickets;
- Specialized tools, services and training to help organizations quickly implement risk-based alerts and see results.

## Risk-Based Alerting

**TONY MORBIN:** What is the concept of risk-based alerting, and how is it different from what organizations do now when they prioritize alerts?

**JESSE TRUCKS:** Traditionally in a SIEM, you have various detections that create alerts that need to be investigated. Then an analyst in the security operations center reviews the alert, investigates it and looks at the data to see whether or not it's a false positive, or if it needs to be triaged for further investigation. If you turned on all of the detections in the SIEM to alert every little thing that represents a potential security problem, you'd get hundreds of tickets a day – even in a modest-size environment. You have ticket overload. An analyst can only investigate 17 to 25 tickets a day, depending on what they are, even if many are false positives and can be quickly dispatched.

Risk-based alerting changes this because you don't have as many small rules to look at these little, discrete detections and create an event that does not always result in a ticket. So, when you have these risk events, you'll have a collection of things happening that indicates a higher fidelity or higher probability of something to investigate. It's a true positive and an actual security incident that somebody needs to look into, whether that's malware, an insider threat or another threat.

With risk rules, you can expand the number of detections you have to very large volume but have a smaller volume of tickets than you used to have because it groups them together with the intelligence under the hood.

## Solving Alert Fatigue and False Positives

**MORBIN:** What problems does risk-based alerting solve?

**TRUCKS:** It's twofold. One is ticket fatigue by the SOC. If you have five people in a SOC and you have 250 tickets in a day, that's too many. But if you have 250 or even 500 risk events in a day, but you only have 60, 70 or 100 tickets, then you can work those. In enterprise security, we call a ticket a notable event – something that needs to be investigated. With risk-based alerting, you reduce the number of tickets, so you can focus on the things that matter. There are fewer false positives and collectively you have a higher confidence level and a higher probability for true positives. So, it solves the problem of alert fatigue and high false positives.

## Peeking Under the Hood

**MORBIN:** How does risk-based alerting work? How does it differentiate between a true positive and a false positive?

**"With risk-based alerting, you reduce the number of tickets, so you can focus on the things that matter. There are fewer false positives and collectively you have a higher confidence level and a higher probability for true positives."**

**TRUCKS:** Traditionally in enterprise security, you have a correlation search that can create a notable event or change the risk score on an object, which is either a system or an account. That's the traditional paradigm, similar to the legacy SIEM approach of one ticket per rule. With risk-based alerting, you have many small detections that look for very discrete, individual things and create risk events. The risk events go into an index, a data store, and then they are related to risk objects. A risk object is a process, a file name, an account ID, a system IP address or a host name – something that can be associated with a log event.

There are risk rules that create risk events and risk-notable correlation searches that look at a collection of risk events. If you meet a certain threshold, which you can configure, that threshold will create a notable event, called a risk notable. It's a ticket that shows up as a view that a SOC analyst expands, so they can see all the associated risk events through time, as well as how the event is related to the individual risk scores that have been added to all those risk objects and how those relate to this search. It shows you what needs to be investigated.

It's a combination of more searches than you used to have for more discrete things that are easier to identify small scores for, and an aggregate of all of those together, that creates a small number of high-fidelity tickets because they have higher fidelity information. It reduces your probability of false positives, because you have so much information that says, "If this many things are happening with this risk score, then you have a better chance of this being actually activated or acted on."

## Tuning Risk Levels Up or Down

**MORBIN:** Can you turn the dial up or down to suit your own risk appetite or change how it works in your environment?

**TRUCKS:** Yes to both. You can determine the risk score added to any risk object for a risk rule. Some defaults are based on our experience with thousands of other events across dozens and hundreds of customers. You might turn the risk score up for a particular activity or system that is more important to you and more potentially dangerous if compromised. You also can change the thresholds and mechanisms you use for creating the risk notable, or the ticket, for the SOC analyst. The threshold could be when the risk score for an object goes over 100, or when the rate of change of a risk score goes beyond a certain amount.

Every environment can be completely customized. However, out of the box, it has some base standards that work for a lot of customers. There's not a lot of tuning required to make it operational. You can be operational with risk-based alerting very quickly, in a matter of weeks.

## Combating Email-Based Attacks

**MORBIN:** Can you share an example of how this process works in practice?

**TRUCKS:** In a sequence of events on a system, such as email, the email is received and opened, an attachment is opened, and immediately the exploit executes and changes a privilege. Then there's a connection that sends data to another

> **“The benefit from Splunk Enterprise Security is work reduction. It takes less work to tune the SIEM for security monitoring and alerting. You’re not chasing false positives. ... You spend time on things that actually need investigation.”**

system. Each one of those is a discrete event. When an email comes in, we’re going to add 10 to the risk score for the user ID and the system it’s on. When an attachment is opened, we’re going to add 15. When an exploit is run, we’re going to add 30. So, each event adds to the risk score, and at some point in that sequence, you’ll reach a threshold. You also can have it create a risk notable with the aggregate risk events today over the last eight hours or whatever schedule you want.

You can also create that for systems you want to review no matter what. In either case, it goes over a threshold, as in the email sequence example, and it looks like a kill chain sequence. Each one of those risk rules and events is related to the MITRE ATT&CK matrix. Once you create a risk notable, you can see that, for example, in a sequence of events, three things happened in parallel. We know exactly what system to look at and we can use our security orchestration and automation to lock a system down, pull data out of it, close it, stop it or put it into a different network.

## Less Work, Faster Resolution

**MORBIN:** What results can security teams expect to see by adopting this approach?

**TRUCKS:** They can spend more focused time on things that matter. It gives them time outside of just chasing tickets and trying to quell the noise. There’s less work needed to tune the system to reduce false positives, because you get less of them. The benefit from Splunk Enterprise Security is work reduction. It’s less work to tune the SIEM for security monitoring and alerting. You’re not chasing false positives and investigating four to seven things in a day that take you 10 to 20 minutes to open up, investigate, verify, document and close. Instead, you spend time on things that actually do need investigation, and because you have more information about what’s happening, your investigation time is shortened.

Also, the remediation piece of it goes faster because the investigation time is shorter and you see more actual events that happen. Nothing’s going to slip through the cracks because of ticket noise. For instance, you might have four or five tickets that are related in an old paradigm, where you spend your efforts, maybe using more than one analyst, looking at something and then discover that two things are related. In that scenario, which analyst takes over the overall event? Instead, you have all of those events in one spot, so one analyst can see the relationship and move to remediation and recovery faster.

## The Splunk Approach

**MORBIN:** How is Splunk helping its customers move toward risk-based alerting?

**TRUCKS:** We start with professional services through Splunk or our partners. We sit down with customers to understand what things are in their environment – what are the low-hanging fruit, the most important critical assets to convert to using risk-based alerting, and what types of technologies are they using? Then, we help them implement it, either by doing it for them or helping them get it up and running. Once it's enabled and on, they're using the new paradigm, and we help teach them how to use it. Then they can continue growing and tuning it themselves – adding risk rules, different risk events and more. Or they can come back and get our assistance, depending on their needs or their budget capabilities.

In addition, to understand how to use risk-based alerting, SOC analysts can take a free RBA workshop on how RBA works, and they can interact with RBA in our nonproduction environment that we create for the workshop. Each analyst is trained on the process of working with RBA, and they have knowledge that they can take directly to production. Whether they're already using RBA or considering it, our RBA workshops and other security workshops are free.

## About ISMG

Information Security Media Group (ISMG) is the world's largest media organization devoted solely to information security and risk management. Each of our 28 media properties provides education, research and news that is specifically tailored to key vertical sectors including banking, healthcare and the public sector; geographies from North America to Southeast Asia; and topics such as data breach prevention, cyber risk assessment and fraud. Our annual global summit series connects senior security professionals with industry thought leaders to find actionable solutions for pressing cybersecurity challenges.

(800) 944-0401 • sales@ismg.io

BANK INFO SECURITY® · CU INFO SECURITY® Just for Credit Unions · GOV INFO SECURITY® · HEALTHCARE INFO SECURITY®

infoRisk® TODAY · CAREERS INFO SECURITY® · Data Breach. Prevention. Response. Notification. TODAY · CyberEd.io

iSMG
INFORMATION SECURITY
MEDIA GROUP

902 Carnegie Center • Princeton, NJ • 08540 • www.ismg.io