splunk> turn data into doing™

WHITE PAPER

Close the Gaps in Your Defenses With Automated Detection, Investigation and Response

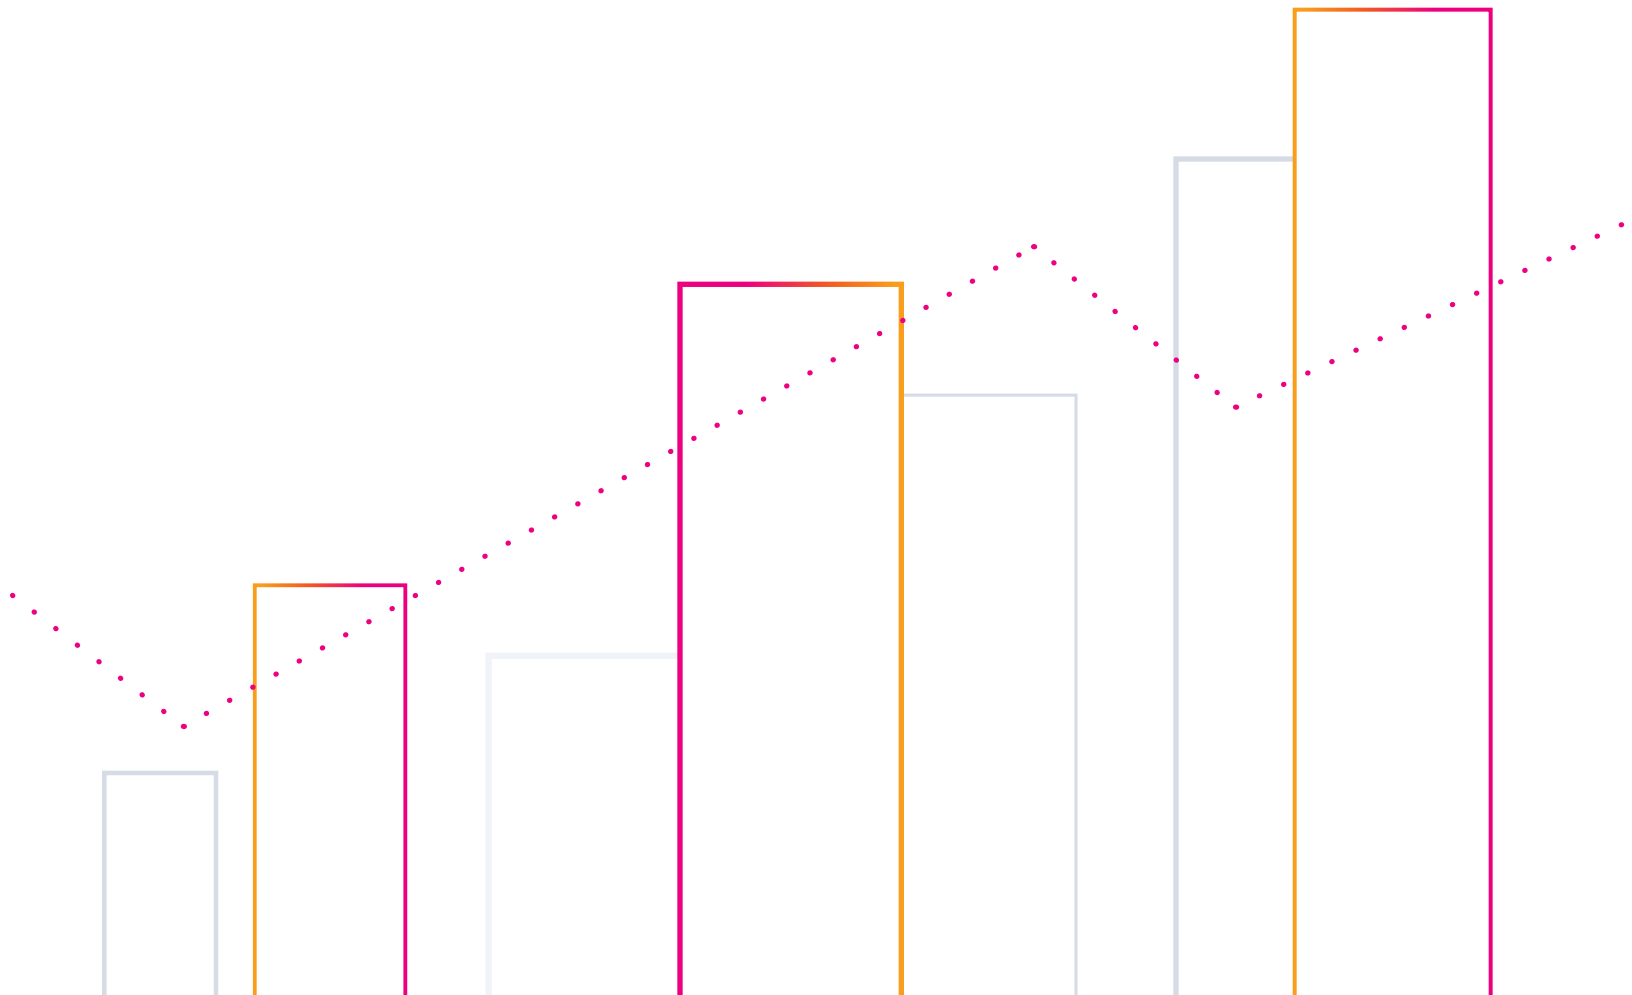

There are thousands of single-purpose security tools on the market. Some search for threats in on-prem environments only. Others only monitor the cloud. Others are focused solely on email security. The list goes on and on.

But even if your company could afford to invest in every single-purpose tool on the market, there would probably still be holes in your defenses. Products from disparate vendors are often unable to communicate with one another, making it difficult to correlate their findings. The alerts from these tools would come without context, turning them into background noise. The disconnect would erode the value of your investment and create vulnerabilities in your security posture.

Regardless of this fact, the dearth of unified solutions forces many companies to cobble together a collection of tools in this fashion — an expensive proposition that, sadly, yields poor results.

This e-book explores the anatomy of a truly integrated defense, including automated detection, investigation and response. The goal is for you to start thinking about where the gaps might be in your own environment, so you can build a strategy to lock them down for good.

What would a fully integrated defense look like?

A fully integrated, automated defense would have the following characteristics:

- **It would be able to ingest all of the various data sources in your organization.** As mentioned above, having tools that can't talk to each other gets in the way of understanding what's really happening in your environment. You need to be able to correlate data sources and get alerts that demonstrate this context.
- **It would produce robust analytics that have the potential to drive actual decisions.** To do effective analysis, you need to be able to monitor your entire environment, including your cloud infrastructure. Being able to deploy detections and investigations that work across all of your on-prem and cloud infrastructure is key.
- **It would eliminate as much manual work as possible.** Today's enterprises produce so much data that it's impossible for mere humans to triage every alert or investigate every finding. What's more, the optimal workflow for detection, investigation and response may not be a straight line. For an analyst, this could take hours. For a fully optimized SIEM leveraging automation, minutes or seconds. Additionally, there is a dearth of available security professionals and it's difficult for most organizations to stay fully staffed.
- **It would send fewer, but more meaningful, alerts.** If you're an analyst, you already know that alert fatigue is a real thing. When you're getting constantly bombarded by alerts, you end up ignoring them (or you'd never have time to do anything else!). Alerts that show history, correlations and context are more conducive to action.

Before You Begin

If automating the entire end-to-end security workflow was cut-and-dried, we'd all have done it ages ago. The fact of the matter is that there are formidable issues to think through before you even start making plans.

Consider the following:

- **How do you know if you can trust your software vendor?** Outside of the obvious (reputation, reviews, notable customers and the like), the vendor should be able to show you how they developed and tested their detections. For example, the Splunk Security Research Team has developed a framework that allows replication, verification and data production of attacks in a shareable, community-friendly fashion. The Splunk open-source **Attack Range** allows the use of adversarial simulation engines — along with tools for measurement, translation and recording in defense technologies that help streamline the process of creating defense artifacts (signatures, detection, investigation, analytics, playbooks and so on). It also provides simulation and verification of the entire cycle of a threat.
- **Do you have access to the data you need?** Your analysis is only as good as your data, so you need to ensure that a.) your organization actually has the data you need for your detections, and b.) you have permissions and access to said data. After that, you'll need to evaluate data quality and ensure that the software you choose can ingest all of your various data sources — structured, unstructured, in motion and at rest.
- **Can you take action on your findings?** You may be able to see your way clear to deploying automated detection and investigation. But what about response? Depending on the type of response you want to trigger (closing ports, reconfiguring the firewall, filing a ticket), you may need additional privileges. You'll also need human resources in place, since most breaches require manual intervention, both to decide that the findings are of concern, as well as to think through the actions you wish to take.

Even if you can't immediately automate your responses in most situations, you can still consider increasing your velocity by automating some of the more low-risk, routine tasks, such as sending alerts or filing tickets.

The Splunk SIEM Approach

While there are many vendors on the market, only a few, like Splunk, automate the entire end-to-end workflow of detection, investigation and response.

The Splunk Security Operations Suite helps you leverage data (both security and non-security related) across your organization to enhance threat detection and response. It provides context and streamlines security operations by helping you collect, aggregate, de-duplicate and prioritize threat intelligence from multiple sources.

The platform enables you to scale and solve a wide range of security use cases. It may be deployed on-premises, in the cloud or in hybrid environments. It integrates machine learning to help detect anomalies and reduce complexity, speeding up threat and attack investigation and response. It also leverages Splunk Phantom playbooks that assist in all phases of the workflow.

But the "secret sauce," in terms of automation, lies within Use Cases and Analytic Stories.

What Is an Analytic Story?

Splunk Analytic Stories are themed security guides that provide you with tactics, techniques and methodologies that help with detection, investigation and response. Analytic Stories also contain easy-to-read background information and guidance, which provide key context for motivations and risks associated with attack techniques, as well as pragmatic advice on how to combat those techniques.

Each story is mapped to various frameworks, including MITRE ATT&CK, Lockheed Martin Kill Chain phases, CIS (Center for Internet Security) controls and NIST (National Institute of Standards and Technology) and include the following content objects.

1. **Detection:** Out-of-the-box detection techniques in the form of detection searches or machine-learning models.
2. **Investigation:** Searches and/or Splunk Phantom playbooks that help the analyst determine the importance of the results of the detections. They may also gather collaborative evidence and additional contextual information.
3. **Response:** These help the analyst conduct specific response actions to remediate the incident.

Analytic Stories are categorized by use case and can be accessed via the Splunk Enterprise Security (ES) Use Case Library, the Splunk **Enterprise Security Content Updates (ESCU)** app, or the open-source Splunk **Security Content Exchange**.

What's Inside an Analytic Story

Each Analytic Story consists of the following elements that provide key insights into each story and its significance as a security methodology.

Description: Brief, top-level overview of the objective of the Analytic Story.

Narrative: Details that provide a better understanding of the attack and/or methodology. Where appropriate, these will include:

- Historical context and evolution
- Associated risk(s) and/or exploits identified in the wild
- Potential impact/rationale for implementation
- Key security value/benefits
- Other associated methods and/or interdependencies

Framework mapping: Categorization within MITRE ATT&CK, corresponding Kill Chain phase and CIS control(s).

Data Model: Relevant data model(s) needed for normalization.

Technologies: Examples of applicable data sources for the Analytic Story.

References: Pointers to additional reading and/or technical resources that provide background, context, links to related methods and/or other relevant information.

Analytic Story searches: Each Analytic Story contains different types of searches, all designed to help with critical tasks performed by security teams at various operational stages. Depending on the Analytic Story, there may be multiple searches provided for each stage.

Each search includes a description and a plain-language, non-technical explanation called "Explain It Like I'm 5." It also displays its associated SPL and provides implementation guidance and known false positives.

- **Detection searches:** Each Analytic Story contains one or more detection searches. These searches are designed to detect activities, events or behavior that are associated with known issues and/or threats. Each search contains supporting details, such as framework mapping, relevant data models/technologies and other related specifics (such as confidence level or a list of at-risk assets).
- **Investigative searches:** Analytic Stories may contain multiple investigative searches that gather context, perform verification steps and/or collect specific types of evidence related to the story.
- **Supporting searches/baseline searches:** These searches provide operational methods that support detection, investigation and response, such as building lookup files. They may also provide other techniques for ensuring proper execution of detection or investigation searches, as well as other operational tasks.
- **Portfolio-wide applicability:** Where appropriate, Analytic Stories incorporates machine-learning models and playbooks from Splunk UBA and Splunk Phantom. This allows security teams to apply the Analytic Story across the entire portfolio of Splunk security products.

You could review these searches and select the individual detections you want to run, but the real power is running the entire Analytic Story end-to-end. Let's look at an example, using the Analytic Story "DNS Hijacking."

This Analytic Story helps users secure their environments against DNS hijacking by detecting and investigating suspicious activities in their environments. It includes the following:

Detection searches:

- **Clients connecting to multiple DNS servers:** This search allows you to identify the endpoints that have connected to more than five DNS servers and has made DNS Queries over the timeframe of the search.
- **Detect hosts connecting to dynamic domain providers:** Malicious actors often abuse legitimate Dynamic DNS services to host malicious payloads or interactive command-and-control nodes. Attackers will automate domain-resolution changes by routing dynamic domains to countless IP addresses to circumvent firewall blocks and blacklists, as well as frustrate a network defender's analytic and investigative processes. This search will look for DNS queries made from within your infrastructure to suspicious dynamic domains.
- **DNS query requests resolved by unauthorized DNS servers:** This search will detect DNS requests resolved by unauthorized DNS servers.
- **DNS record changed:** The search takes the DNS records and the results of the discovered DNS records lookup to see if any records have changed within the last day.

Investigative search:

- **Get DNS server history for a host:** While investigating any detections, it is important to understand which and how many DNS servers a host has connected to in the past. This search uses data that is tagged as DNS and gives you a count and list of DNS servers that a particular host has connected to within the previous 24 hours.

Support search:

- **Discover DNS records:** Discover the DNS records and their answers for domains owned by the company using network traffic events. The discovered events are exported as a lookup.

Playbook:

- **DNS hijack enrichment playbook:** This playbook can be set to run when the detection search "DNS Record Changed" is used to identify DNS record changes for cloud and corporate domains used in your environment. The detection search is dependent on a feature called "Discover DNS Records," which finds the common DNS responses for the last 30 days of monitored corporate domains and cloud providers. These responses are stored in a lookup file. The playbook starts with the changed DNS records and uses MaxMind, whois, Censys, Malware Domain List and PassiveTotal to gather attributes of the DNS records for comparison against expected values.

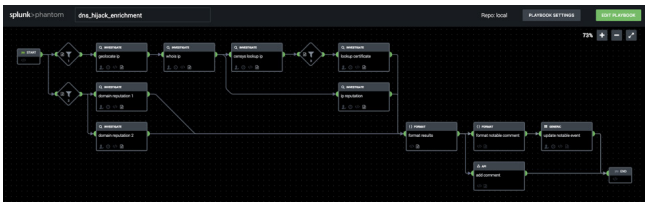

All of these objects can be enabled and configured to be run individually, but their real power is when they are used together, passing data forward when it's relevant to the next step (i.e., detection results are fed to investigation, which looks more deeply if they meet specific criteria; results of investigations trigger response, where appropriate).

The End-to-End Workflow

The **Analytic Story Execution** (ASX) open-source app runs an Analytic Story end-to-end, generating use-case relevant context and correlation when events are generated. This application gives you the tools to make the execution of an Analytic Story in Splunk an automated process.

In the example of the DNS hijacking story, the detection searches are designed to find suspicious DNS activity, such as an unusual number of DNS connections and queries, connections to suspicious dynamic DNS providers, altered DNS records and so on. When these detections return worrisome findings, you'd want to dig deeper. Unless you truly have nothing else to do, there is no reason why you would want to manually run the investigative search (but you want that info!). So ASX carries out that step. In this case, the investigative search provides the DNS server history for the host.

The support search also generates data used by the playbook. "Discover DNS Records," as mentioned above, returns a list of altered DNS records. The playbook takes those records and checks them against industry lists of malicious domains.

If you wanted to move beyond investigation and into response, you could use a Splunk Phantom playbook to file a ticket or initiate another relevant action.

ASX was designed to be easy to use. It's as easy as selecting an Analytic Story and clicking "Submit!" It uses two simple commands:

1. `Detect:` This is a **custom search generating command** that runs all baseline and detection searches in an Analytic Story.
2. `Investigate:` Another custom search generating command that runs all investigative searches on all the entities generated using the `detect` command. Investigate is a streaming command that is executed after the `detect` command has generated the result object.

Getting Started

As mentioned earlier, in addition to Splunk ES and **ESCU**, you can access Analytic Stories via the open-source Splunk **Security Content Exchange**. This Github project was designed to bring the community together to improve our collective defenses. By sharing research and analytics, we can help the entire industry craft more effective strategies. This project provides a mechanism to facilitate this exchange.

If you have any questions about how to use Splunk Analytic Stories, the Attack Range or ASX, feel free to email **research@splunk.com**.

2020-Splunk-SEC-Close-The-Gaps-In-Your-Defenses-102-WP