The Essential Guide to **Security**

How to Get Started Using Splunk for Security to Solve Your Everyday Challenges

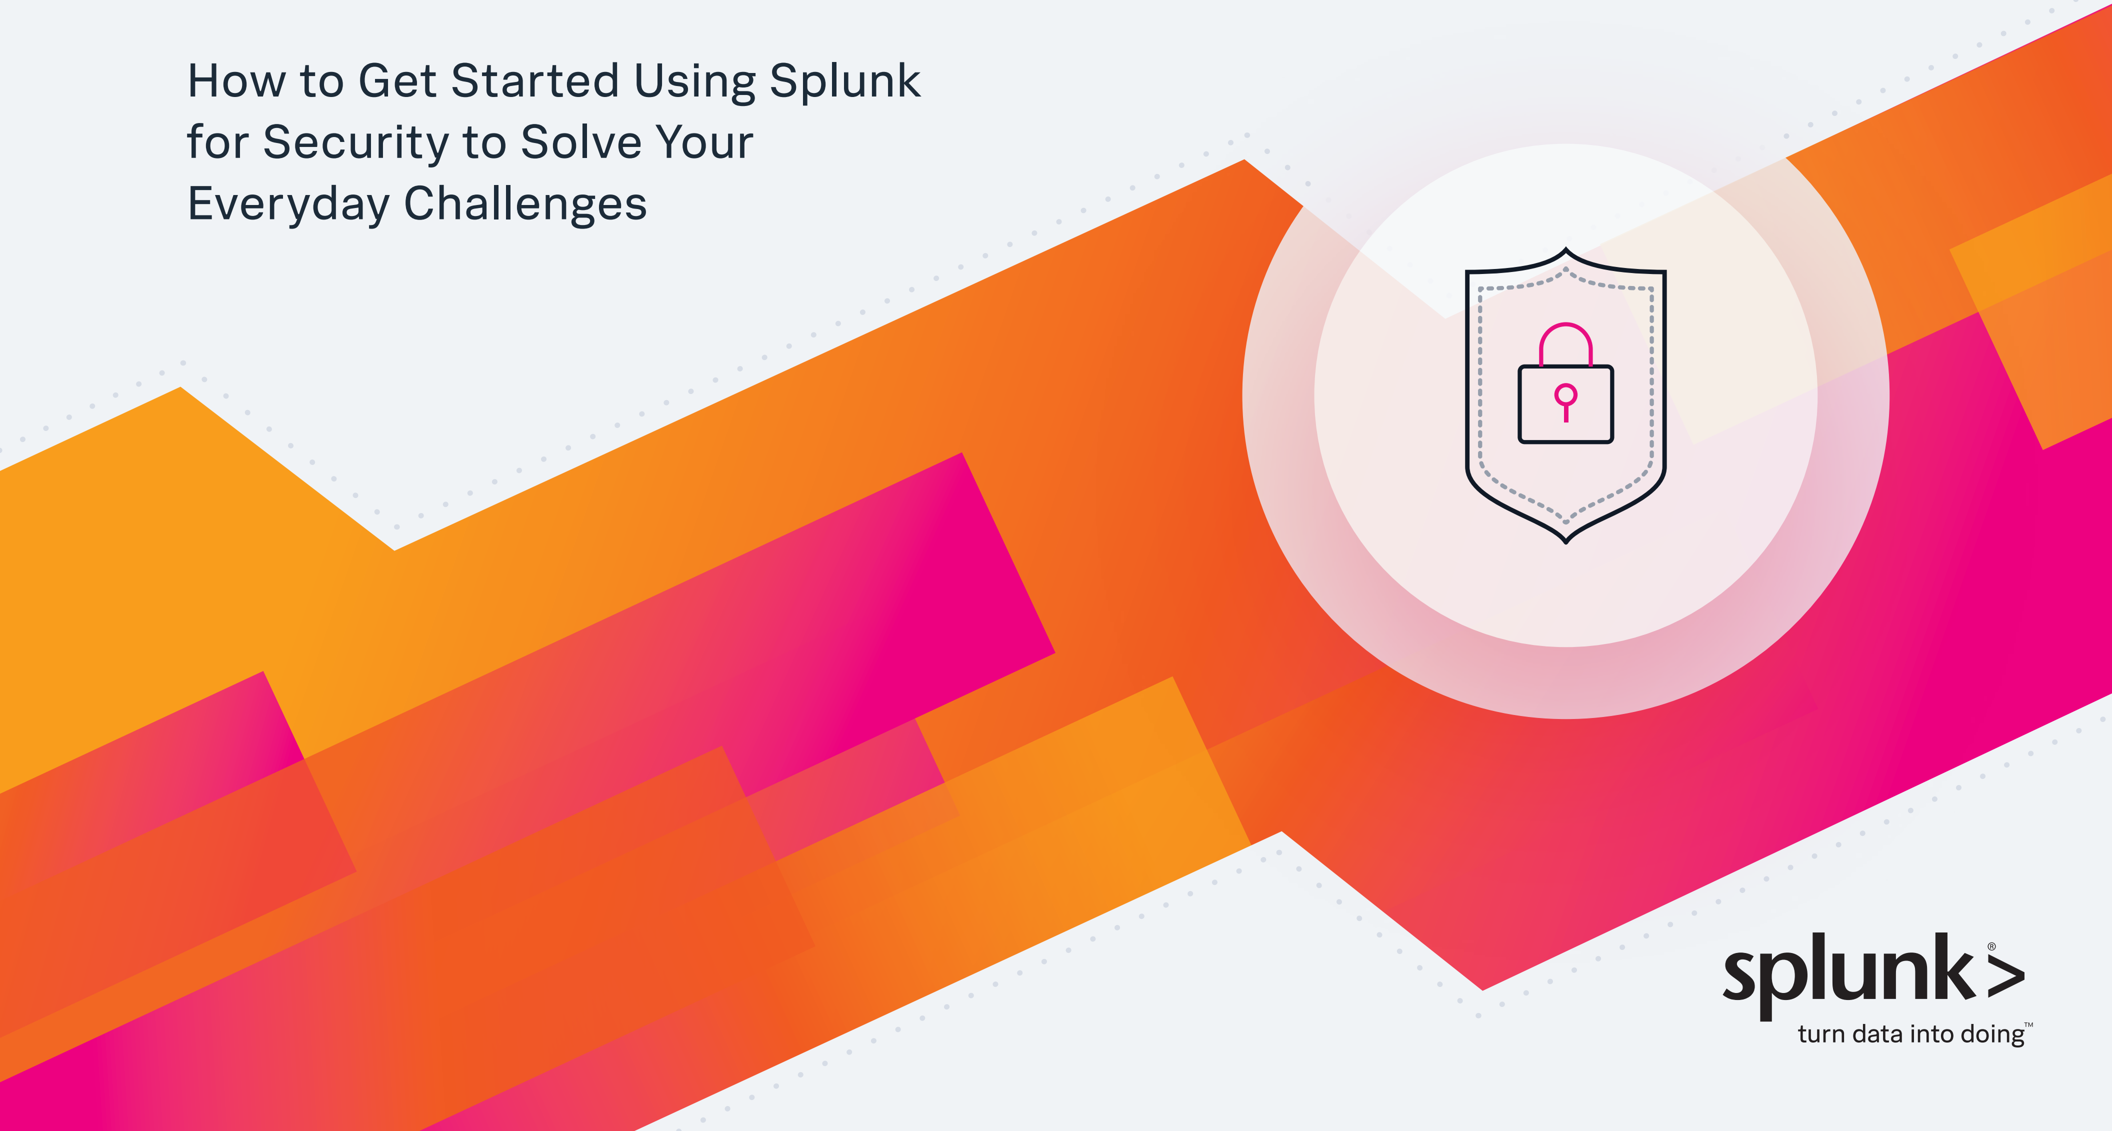

splunk>

turn data into doing™

What's Your Plan for Cybersecurity?

Are you simply "planning for the worst, but hoping for the best?"

Table of Contents

So how can you best defend your organization and hunt down new adversaries?

Ultimately, by taking a holistic approach to your defense system across the enterprise.

Introduction

What's your plan for cybersecurity? Are you simply "planning for the worst, but hoping for the best?" With digital technology touching every part of our lives and new threats popping up daily, it's imperative that your organization is precise, informed and prepared when it comes to defending your assets and hunting your adversaries.

High-profile breaches, global ransomware attacks and the scourge of cryptomining are good enough reasons why your organization needs to collect, leverage and understand the right data. You'll also need to implement the right processes and procedures, often alongside new technologies, methods and requirements–all with an ever-increasing velocity and variety of machine data.

So how can you best defend your organization and hunt down new adversaries? Ultimately, by taking a holistic approach to your defense system across the enterprise. This is why Splunk believes every organization needs a security nerve center, implemented by following a six-stage security journey that we will describe for you.

Let's break down what that means.

Splunk in the Security Operations Center (SOC)

Data-driven businesses take advantage of the investigate, monitor, analyze and act (IMAA) model to advance their security by optimizing their people, processes and technology. It includes using all the data from the security technology stack, which can help you investigate, detect and take rapid, coordinated action against threats in a manual, semi-automated or fully-automated fashion. When security teams invest in their security infrastructure, their security ecosystem and skills become stronger, making it possible to expand security practices into new areas and proactively deal with threats.

The Splunk Data-to-Everything Platform and Splunk's security portfolio brings together multiple cybersecurity areas, as well as others outside of security, to foster collaboration and implement best practices for interacting with your data. Security teams can use Splunk solutions to drive statistical, visual, behavioral and exploratory analytics that inform decisions and actions. From there, the platform allows for a modern workflow, from collecting data all the way to invoking actions to address cyberthreats and challenges.

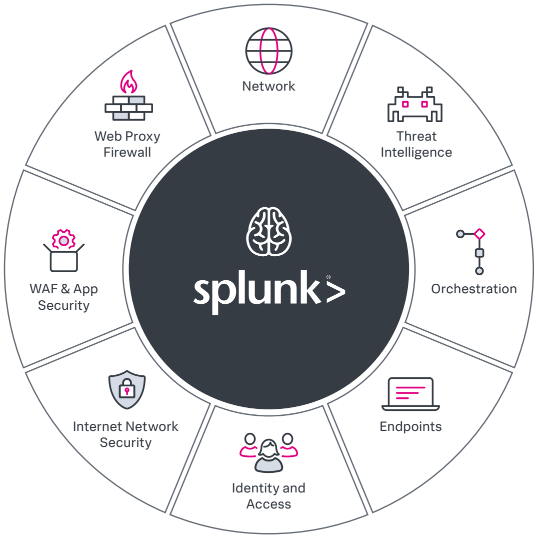

Figure 1: Splunk Enterprise Security includes a common framework for interacting with data and invoking actions. The Adaptive Operations Framework enables security teams to quickly and confidently apply changes to the environment. Splunk Enterprise Security can automate the response as well, enabling the security infrastructure to adapt to the attacker using a range of actions appropriate to each domain.

Sound good?

Great. So how do I make all of this happen in the real world, you ask?

To get you started, we put together this short guide to introduce you to the top security use cases organizations face and to show you how Splunk's analytics-driven platform can help you solve your security challenges. This guide is divided into three sections:

1. **Understanding the Fundamentals.** Here you will find an introduction to the security journey and a quick primer on security use cases with each use case mapped to relevant Splunk solutions.
2. **Embarking on Your Analytics-Driven Security Journey.** Here we explain the six stages of the data-driven security journey–and what you should be able to do, and how well, at each stage.
3. **Solving Common Security Challenges With Splunk.** Here we walk through examples of how to solve common security challenges associated with:
   - Incident investigation and forensics
   - Security monitoring
   - Advanced threat detection
   - SOC automation
   - Incident response, compliance
   - Fraud and analytics detection
   - Insider threat

Ready to create a kick-ass security practice? We thought so.

Understanding the Fundamentals

Cyber criminals never rest, which means that you should constantly be looking for new security use cases and insights to maintain high levels of protection in your environment.

We're here to help.

Splunk's Analytics-Driven Security Journey

Anyone who has been asked the question, "Are we secure?" knows that cybersecurity is a journey, not a destination. While the expedition may not have an end point, and there will always be challenges, there are things you can do to make the trip successful.

First, you must understand your environment and find a place to begin. Ask yourself: What am I trying to protect? What is my critical data? How will I respond to the threats?

The six-stage analytics-driven security journey, shown in Figure 2, will help you answer these questions and create a kick-ass security practice that enables you to understand the gaps in your defenses, see the next challenge and take action to meet it head on.

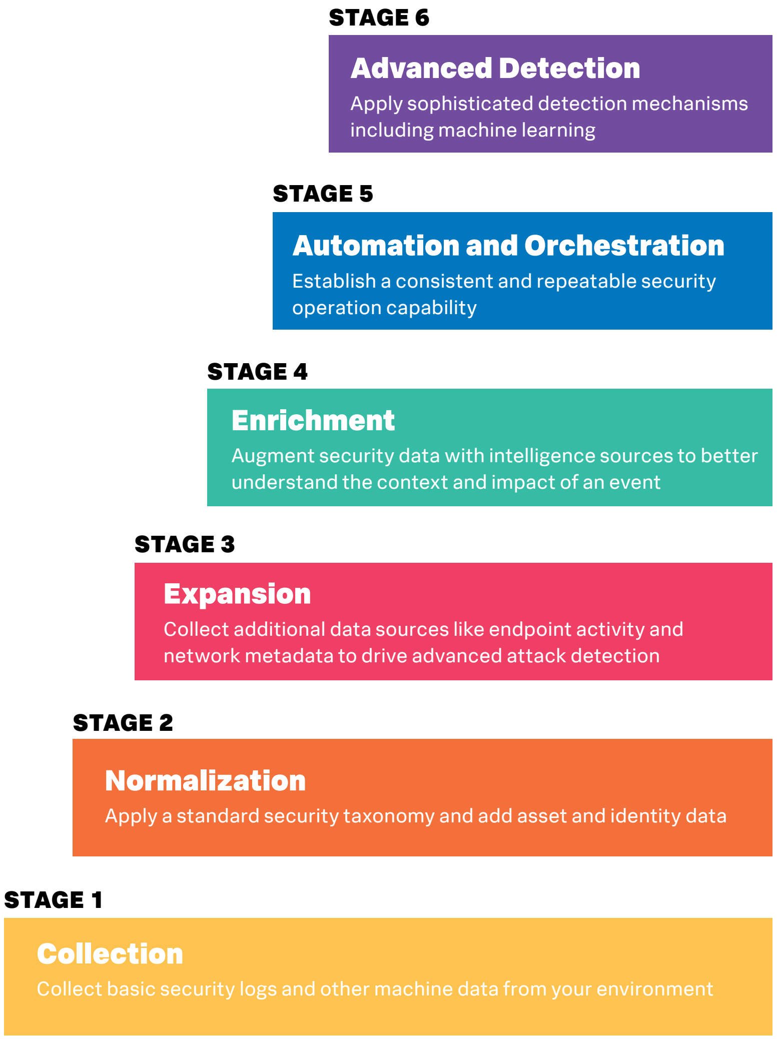

STAGE 6

Advanced Detection

Apply sophisticated detection mechanisms including machine learning

STAGE 5

Automation and Orchestration

Establish a consistent and repeatable security operation capability

STAGE 4

Enrichment

Augment security data with intelligence sources to better understand the context and impact of an event

STAGE 3

Expansion

Collect additional data sources like endpoint activity and network metadata to drive advanced attack detection

STAGE 2

Normalization

Apply a standard security taxonomy and add asset and identity data

STAGE 1

Collection

Collect basic security logs and other machine data from your environment

Figure 2: Splunk's Analytics-Driven Security Journey

Splunk's Security Suite

Would you embark on a hiking expedition without a trail map and a backpack full of provisions and the right gear? Of course not. Just as no trip can be successful without the proper equipment, no security journey can be successful without the right technology.

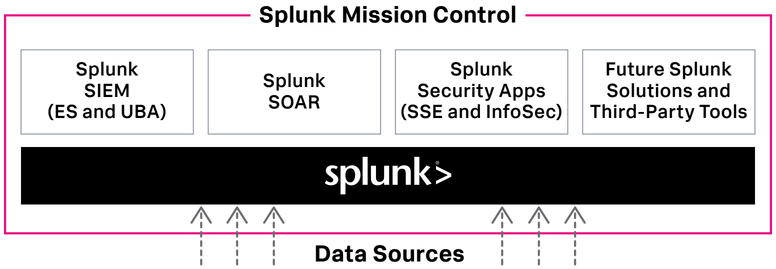

Splunk Mission Control
Splunk SIEM (ES and UBA)
Splunk SOAR
Splunk Security Apps (SSE and InfoSec)
Future Splunk Solutions and Third-Party Tools
splunk>
Data Sources

Splunk's Security Suite helps security teams navigate uncharted waters and quickly identify, investigate, respond to and adapt to threats in dynamic, digital business environments. Splunk solutions can be used by a tier-1 analyst to do basic research on a time period, keyword, IP address or machine name. The same products enable advanced tier-2 and tier-3 analysts to perform advanced correlations, build analytical models or perform advanced forensics.

Splunk's Security Suite	
Splunk Enterprise	Is a flexible platform addressing an array of security use cases, enabling you to monitor and analyze machine data quickly from any source to deliver insights to act and the analytics-driven foundation to strengthen your overall security. Available in the cloud.
Splunk Enterprise Security	A security information and event management (SIEM) solution that provides insights into machine data generated from security technologies such as network, endpoint and access; as well as malware, vulnerability and identity information. Available in the cloud.
Splunk User Behavior Analytics	A machine-learning-powered solution that delivers answers organizations need to find unknown threats and anomalous behavior across users, endpoint devices and applications.
Splunk SOAR	A security orchestration, automation and response (SOAR) platform that integrates with your existing security technologies to provide a layer of "connective tissue" between them, making them smarter, faster and stronger.
Applications	Apps developed by Splunk, partners and our community to enhance and extend the power of the Splunk platform. The Splunk App for Payment Card Industry (PCI) Compliance is an example. Available in the cloud.
Splunk Security Essentials	Explore new use cases and deploy security detections from Splunk Security Essentials to Splunk Enterprise and Splunk Cloud and Splunk SIEM and SOAR offerings. Now a fully-supported app with an active Splunk Cloud license, start strengthening your security posture and quicken your time-to-value with Splunk.
Splunk Enterprise Security Content Updates	For customers with Splunk Enterprise Security (ES), this delivers security analysis guides, called "Analytic Stories," that explain how to best use Splunk ES to investigate and take action on new threats detected in your environment, what searches to implement and what you should be able to achieve.

The Security Use Cases

Next, you will find the specific security use cases we've mapped to the journey. Go ahead. Choose your own adventure, or security challenge. The purpose of this book is to teach you how Splunk's analytics-driven platform can help solve your security challenges and advance your security journey, including:

Mapping Splunk Solutions to Security Use Cases

Use Case	Splunk Solution
Incident Investigation and Forensics	Splunk Enterprise, Splunk Enterprise Security, Splunk SOAR
Security Monitoring	Splunk Enterprise, Splunk Security Essentials App, Splunk Enterprise Security, Splunk SOAR
Advanced Threat Detection	Splunk Enterprise, Splunk Security Essentials App, Splunk Enterprise Security, Splunk User Behavior Analytics
SOC Automation	Splunk Enterprise, Splunk Enterprise Security, Splunk SOAR
Incident Response	Splunk Enterprise, Splunk Enterprise Security, Splunk SOAR
Compliance	Splunk Enterprise, Splunk Security Essentials App, PCI, Splunk Enterprise Security
Fraud Analytics and Detection	Splunk Enterprise, Splunk Security Essentials App, Splunk Enterprise Security
Insider Threat Detection	Splunk Enterprise, Splunk Security Essentials App, Splunk User Behavior Analytics

The Security Use Cases Defined

Lastly, we include a quick primer on the use cases so we are all on the same page.

Incident Investigation and Forensics

Security incidents can occur without warning and can often go undetected long enough to pose a serious threat to an organization. Usually by the time security teams are aware of an issue, there's a good chance the damage has been done. Splunk provides security teams with a "single source of truth" for all time-stamped machine data in a computing environment. This helps them drive better, faster security investigations, reducing the chance of a threat going undetected for extended periods.

Security Monitoring

Security monitoring enables you to analyze a continuous stream of near-real-time data for threats and other potential security issues. Data sources for monitoring include network and endpoint systems–as well as cloud devices, data center systems and applications. The Splunk Data-to-Everything Platform enables security teams to detect and prioritize threats found in the stream of data from these sources.

Advanced Threat Detection

An advanced persistent threat (APT) is a set of stealthy and continuous computer-hacking processes, often orchestrated by a person or persons targeting a specific entity. APTs usually target private organizations and/or states for business or political motives. Splunk Enterprise enables organizations to search and correlate data to track advanced threats. Splunk Enterprise Security and Splunk User Behavior Analytics elevate existing capabilities to apply a kill chain methodology through statistical analysis, anomaly detection and machine learning techniques to detect unknown and advanced threats.

SOC Automation

Security operations teams adopt Splunk software for orchestration and automation of enrichment and response actions, as well as for case management (i.e. incidents). They use Splunk's SOC automation solutions to scale operations, accelerate response and remediate threats and other security issues. Splunk solutions also help organizations operationalize analytics-driven security practices and empower security teams to collaborate across the extended team.

Incident Response

Incident Response (IR) involves the monitoring and detection of security events on IT systems, and the execution of response plans to those events. IR Teams are sometimes called blue teams. Blue teams defend an organization’s infrastructure when threats are detected, whereas red teams attempt to discover weaknesses in the existing configuration of those same systems. Splunk offers a variety of IR capabilities in the security portfolio, depending on the offerings you choose. Each offers mechanisms to perform investigations for detected events. Splunk solutions may also include capabilities to guide an incidenct responder through standardized response procedures.

Compliance

In nearly all environments, there are regulatory requirements in one form or another—especially when dealing with the likes of GDPR, HIPAA, PCI, SOX and even common guidelines that aren’t considered true compliance, such as the 20 CIS Critical Security Controls. There are many ways of solving compliance challenges using Splunk solutions. One example is creating correlation rules and reports that identify threats to sensitive data or key employees, as well as to automatically demonstrate compliance.

Fraud Analytics and Detection

Machine data plays a pivotal role in and is at the heart of detecting fraudulent activities in the digital age. Splunk can onboard new data so that fraud teams are able to better detect and investigate anomalies. As a result, companies are positioned to reduce financial losses, protect reputation and maintain organizational efficiencies.

Insider Threat Detection

Insider threats come from current or former employees, contractors or partners who have access to the corporate network and intentionally or accidentally exfiltrate, misuse or destroy sensitive data. They often have legitimate access to networks and permission to download sensitive material, easily evading traditional security products. By using Splunk solutions, security teams have the ability to detect and prioritize threats posed by insiders and compromised insiders that would have otherwise gone undiscovered.

Embarking on Your Analytics-Driven Security Journey

To be effective, a cybersecurity program must continually evolve. The problem is that many organizations lack a clear sense of where they are and how to improve. Knowing where you are in your journey will help you manage your time and resources more effectively. And with a sense of what’s to come, you can better plan for success in the later stages.

Here’s a breakdown of the six stages of the analytics-driven security journey that uses data to stay ahead of attacks. For each stage we take a look at:

- Use case applicability
- Data sources
- Milestones
- Challenges

Stage 1: **Collection**

Collect basic security logs and other machine data from your environment.

Security Use Case Applicability

Incident Investigation & Forensics

●○○○

Security Monitoring

●○○○

Advanced Threat Detection

●○○○

SOC Automation

○○○○

Incident Response

●○○○

Compliance

●○○○

Fraud Analytics and Detection

●○○○

Insider Threat

●○○○

Description

Stage 1 focuses on obtaining the raw materials necessary to start gaining a deeper understanding of the environment you must defend.

Data Sources

The best practice for Stage 1 is to capture machine data generated by the four foundational components of your security infrastructure:

1. **Network.** Visibility into network traffic is critical for any security team. At this early level, the priority is to see what types of traffic are entering and exiting your network. It's critical to see the traffic that's permitted as well as communication attempts that have been blocked.

   Sources include:

   - Firewall traffic logs from vendors such as:
     - Palo Alto Networks
     - Cisco
     - Checkpoint
     - Fortinet

2. **Endpoint (host-based).** Endpoint logs complement network visibility to give insight into malicious activities such as malware execution, an insider performing unauthorized activity or an attacker dwelling in your network. It's important to capture this data from servers, workstations and all operating systems.

   Sources include:
   - Windows event logs
   - Linux system logs
   - Linux Auditd logs
   - MacOS system logs

3. **Authentication.** Authentication logs can tell you when and from where users are accessing systems and applications. Since most successful attacks eventually include the use of valid credentials, this data is critical in helping to tell the difference between a valid login and an account takeover.

   Sources include:
   - Windows Active Directory
   - Local authentication
   - Cloud identity and access management (IAM)
   - Linux auditd logs
   - MacOS system logs

4. **Web activity.** Many attacks start with a user visiting a malicious website or end with valuable data being exfiltrated to a site that the attacker controls. Visibility into who's accessing what sites and when is critical for investigation.

Sources include:
- Next-generation firewall (NGFW) traffic filters or proxy logs from vendors such as:
  - Palo Alto Networks
  - Cisco
  - Checkpoint
  - Fortinet
  - Bluecoat
  - Websense

Milestones

After successfully onboarding data from these four categories, you should have achieved the following milestones:

- Critical activity logs reside in a separate system where they can't easily be tampered with by an attacker; and
- Data from the four categories is available to perform basic investigations.

Challenges

Collecting the different data sources can be a hassle and making sure the data is onboarded correctly can be tedious. It's often done incorrectly and insufficient information is captured, causing lost time and incomplete investigations.

Stage 2: **Normalization**

Apply a standard security taxonomy and add asset and identity data.

Security Use Case Applicability

Incident Investigation & Forensics

▰▱▱▱

Security Monitoring

▰▱▱▱

Advanced Threat Detection

▰▱▱▱

SOC Automation

▰▱▱▱

Incident Response

▰▱▱▱

Compliance

▰▱▱▱

Fraud Analytics and Detection

▰▱▱▱

Insider Threat

▰▱▱▱

Description

In Stage 2, you're ensuring your data is compliant with a standard security taxonomy. This means that fields representing common values, such as the source IP address, port, username and so on, now have common names, regardless of which device created the event. This critical investment in normalizing data allows you to:

- Consume a larger selection of detection mechanisms from vendors and the community;
- Start implementing a security operations center to track systems and users on your network; and
- Begin to scale the capabilities of your security team.

Even if you don't plan to stand up a formal SOC, normalized data will:

- Facilitate cross-source correlation;
- Streamline investigations; and
- Improve the effectiveness of an analyst.

Data Sources

In Stage 2 you should collect reference information about:

- IT assets (systems, networks, devices, applications); and
- User identities from Active Directory, LDAP and other IAM/SSO systems.

Milestones

Milestones for Stage 2 include:

- Data is mapped properly to the Common Information Model (CIM);
- Search performance is improved dramatically through the use of accelerated data models associated with CIM; and
- Asset and user details are correlated to events in your security log platform.

Challenges

Although you have basic data that is searchable, you lack the insights or understanding you need for deeper security detections and endpoint visibility.

Stage 3: **Expansion**

Collect additional high-fidelity data sources, like endpoint activity and network metadata, to drive advanced attack detection.

Security Use Case Applicability

Incident Investigation & Forensics
●●○○

Security Monitoring
●●○○

Advanced Threat Detection
●●○○

SOC Automation
●●○○

Incident Response
●●○○

Compliance
●○○○

Fraud Analytics and Detection
●○○○

Insider Threat
●○○○

Description

Domain name system (DNS) and endpoint data will unlock a rich set of detection capabilities, empowering threat hunters to uncover and track adversaries dwelling in the network.

Data Sources

Data sources at this stage include:

1. **Network.** Most threat hunters and threat intelligence analysts will tell you that if they could only have one data source for analysis, it would be DNS.

   Sources include:

   - Protocol-specific wire data from sources like Splunk Stream or Bro;
   - DNS query-level data from debug-level logs or from wire data sources; and
   - DHCP activity.

2. **Endpoint.** Rich endpoint activity that captures process creation, file changes, registry modifications, network connections and so on provides an amazingly clear history of critical events occurring on an endpoint.

   Sources include:

   - sysmon
   - Osquery
   - Carbon Black Defense

Milestones

By collecting high-fidelity data sources, you will have:

- Laid the foundation for advanced detections; and
- Obtained the ability to match some common indicators of compromise.

Challenges

The network and endpoint data you're collecting is rich in detail; however, it lacks context and might contain indicators of compromise that are known to your peer organizations, but remain undetected in your environment.

Stage 4: **Enrichment**

Augment security data with intelligence sources to better understand the context and impact of an event.

Security Use Case Applicability

Incident Investigation & Forensics	🟩🟩🟩⬜
Security Monitoring	🟩🟩🟩⬜
Advanced Threat Detection	🟩🟩🟩⬜
SOC Automation	🟩🟩⬜⬜
Incident Response	🟩🟩🟩⬜
Compliance	🟩⬜⬜⬜
Fraud Analytics and Detection	🟩⬜⬜⬜
Insider Threat	🟩⬜⬜⬜

Description

In addition to collecting vital machine data, high-performing security teams enrich their data with intelligence from internal and external sources. Contextual and investigative knowledge, including threat-intelligence feeds, open-source intelligence (OSINT) sources and internally sourced information, allows security personnel to extract more value from the collected data to detect security events and incidents sooner.

Data Sources

Data sources include:

- Local IP/URL block lists
- Open-source threat intelligence feeds
- Commercial threat intelligence feeds

Milestones

By enriching data with intelligence that provides context, security personnel are able to:

- Understand the urgency of an alert based on the criticality of the asset; and
- Augment alerts by matching them against threat-intelligence feeds, pivoting to other systems and initiating additional context gathering activities.

Challenges

You have significant detection capabilities, but your team is operating in an ad-hoc fashion or failing to consider the context of what they are seeing by correlating with information from outside of the business. Also, requests are not tracked, performance is not measured, collaboration is ad hoc and lessons learned are neither stored nor leveraged for future use.

Stage 5: **Automation and Orchestration**

Establish a consistent and repeatable security operation capability.

Security Use Case Applicability

Incident Investigation & Forensics
■■■■

Security Monitoring
■■■■

Advanced Threat Detection
■■■□

SOC Automation
■■■■

Incident Response
■■■■

Compliance
■■□□

Fraud Analytics and Detection
□□□□

Insider Threat
■□□□

Description

Leveraging a security orchestration, automation and response (SOAR) solution allows organizations to reduce risk in a number of powerful ways. Some of the key benefits of implementing automation and orchestration is the ability to strengthen your defenses by integrating existing security tools and threat intelligence sources, speed response to security events, simplify the investigation process and minimize damage from attacks. Mature organizations are able to continuously triage and prioritize inbound alerts automatically, freeing their human resources to focus on the most critical issues that require their attention. Mature organizations also gain better consistency and repeatability through the execution of standardized automation playbooks vs. execution of a response plan manually.

Data Sources

Data sources in this stage include the high-fidelity events that are generated by data platforms like Splunk Enterprise. Correlation searches, notable events, and other high-fidelity events are ingested by an automation and orchestration system for further actioning.

Milestones

Milestones in Stage 5 include the ability to:

- Track incidents;
- Regularly measure analyst effectiveness;
- Take action according to prescribed playbooks; and
- Automate simple response actions and combine them together into more sophisticated orchestration.

Challenges

Security teams are usually hard at work on the front lines, identifying, analyzing and mitigating threats when and where possible. Yet despite their best efforts, security incident backlogs continue to grow with time spent on consuming investigations and already-known threats. (The reality is that there simply aren't enough skilled professionals to analyze the volume of incidents that most organizations face.)

Stage 6: **Advanced Detection**

Apply sophisticated detection mechanisms, including machine learning.

Security Use Case Applicability

Incident Investigation & Forensics
●●●●

Security Monitoring
●●●●

Advanced Threat Detection
●●●●

SOC Automation
●●●●

Incident Response
●●●●

Compliance
●●○○

Fraud Analytics and Detection
○○○○

Insider Threat
●●●●

Description

By applying machine learning, data science and advanced statistics to analyze the users, endpoint devices and applications in your environment, you're giving yourself a fighting chance to detect adversaries, unknown threats and insider threats, even when they leave only subtle traces of their activity.

Data Sources

Hunting adversaries requires more granular data collection from your endpoints. Rich endpoint activity that captures process creation, file changes, registry modifications and network connections provide an amazingly clear history of critical events occurring on an endpoint.

Sample sources include:

- Microsoft sysmon
- Osquery
- Carbon Black Defense

Milestones

In Stage 6 you're employing:

- The most advanced techniques available to identify unknown threats; and
- New detection mechanisms as they become available, leveraging both your team's expertise and outside research organizations.

Challenges

At this point, you'll be challenged to constantly improve your security organization and gain new capabilities. Your team will also likely be required to perform new research. But by following the journey and maturing your capabilities, you're at the top of your game. Although you will always be under attack, you have put yourself in the best position to detect and prevent many common and not-so-common threats to modern organizations.

Solve Common Security Challenges With the Splunk Security Operations Suite

The security journey can be a bumpy road. Wouldn't it be great if you had a handbook of the challenges you may encounter so that, when they happen, you have the tools to handle the situation and stay on course?

No worries. We have you covered.

Here we provide some examples to help you solve 16 common security challenges (you can find more in the Splunk Security Essentials app or Splunk Security Online Demo). Each example explains the challenge and provides information about data sources, use cases, Splunk sotlutions, programming difficulty, how to implement, alert volume, known false positives and how best to respond.

The examples include:

- Incident investigation and forensics
  - Detect lateral movement with WMI
  - Identify multiple unauthorized access attempts
- Security monitoring
  - Detect public S3 buckets in AWS
  - Find multiple infections on host
- Advanced threat detection
  - Detect connections to new domains
  - Find emails with lookalike domains
- SOC automation
  - Automate malware investigations
  - Automate phishing investigations
- Incident response
  - Detect web data exfil DLP alerts for user
  - Identify basic dynamic DNS detection
- Compliance
  - Detect new local admin account
  - Find user logged into in-scope systems they should not have
- Fraud analytics and detection
  - Detect compromised user accounts
  - Find anomalous healthcare providers
- Insider threat detection
  - Detect large web upload
  - Detect successful login of account for former employee

Incident Investigation and Forensics

Detect Lateral Movement With WMI

STAGE 3

MITRE ATT&CK Tactics

Lateral Movement | Execution

MITRE ATT&CK Techniques

Remote Services | Windows Management Instrumentation

Data Sources

Windows Security | Endpoint Detection and Response

Security Challenge

Windows management instrumentation (WMI) has been gaining popularity amongst attackers for its ability to perform system reconnaissance, antivirus and virtual machine detection, code execution, lateral movement, persistence and data theft.

Use Case

Advanced threat detection

Category

Lateral movement

Splunk Solutions Required

Simple search assistant

SPL Difficulty

Basic

How to Implement

This use case requires sysmon to be installed on the endpoints you wish to monitor and the sysmon add-on installed on your forwarders and search heads.

Alert Volume

Low

Known False Positives

No known false positives

How to Respond

When this search fires, you will want to start your incident response process and investigate the actions taken by this process.

Detect Lateral Movement with WMI Help

To detect lateral movement with WMI we first load our sysmon EDR data. Any other process launch logs with the full command line will also suffice. We look for any instances of Windows management instrumentation command-line (WMIC) being launched (EventCode 1 indicates a process launch), and filter to make sure our suspicious fields are in the CommandLine string.

```
index=* sourcetype=XmlWinEventLog:Microsoft-Windows-sysmon/
Operational EventCode=1 Image=*wmic* CommandLine=*node*
CommandLine="*process call create*"
| table _time host Image CommandLine
```

Identify Multiple Unauthorized Access Attempts

STAGE 1

MITRE ATT&CK Tactics

Credential Access

MITRE ATT&CK Techniques

Brute Force

Data Sources

Authentication

Windows Security

Security Challenge

Most login failures are due to failed passwords. However, multiple login failures to sensitive systems where the users simply aren't authorized can indicate malicious intent. In most organizations, it's rare for a user to get an unauthorized message, beyond low-risk scenarios such as proxy logs. When this occurs for higher-risk activities, such as system logins, file share access and so on, and when it occurs persistently for a user, there's usually a reason to investigate.

Use Case

Insider threat

Category

Insider threat

Splunk Solutions Required

Simple search assistant

SPL Difficulty

Medium

How to Implement

Ensure that you have data being ingested from the Universal Forwarder and the Splunk Technology Add-on present, and all will work automatically.

Alert Volume

Low

Known False Positives

The most likely scenario where this detection indicates a false positive is where the user's access has simply been messed up. For example, there was an AD group change last night, and the user was accidentally removed from the "dev_system_access" security group. Beyond that, there's no standard pattern that would be expected for false positives.

How to Respond

When this alert fires:

1. Evaluate whether the user has previously had access to the desired resources.
2. Look for recent job role changes.
3. Look for recent changes around AD groups.

In most organizations, the next escalation step would be to consult the resources' owner, and/or the user's manager, to determine whether this behavior is intended. Keep an eye out for indications of malicious intent alongside potential account compromise.

Identify Multiple Unauthorized Access Attempts Help

To find multiple unauthorized access attempts using live data, we use the simple searchand the below search processing language. Here we are bringing in Windows Security logs and specifically looking for the status code 0xC000015B, which indicates that the user hasn't been granted the requested logon type. We're looking for any user with many of these per day, which can indicate attempted access to sensitive resources. The screenshot below shows the results of a search on demo data.

```
index=* source=win*security user=* EventCode=*
action=failure Logon_Type=* Failure Reason Logon Type
Status=0xC000015B
```

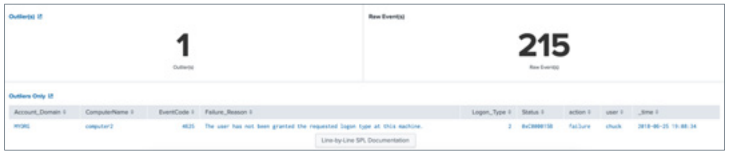

Account_Domain	ComputerName	EventCode	Failure_Reason	Logon_Type	Status	action	user	_time
MYORG	computer2	4625	The user has not been granted the requested logon type at this machine.	2	0xC000015B	failure	chuck	2018-06-25 19:08:34

Security Monitoring

Detect Public S3 Buckets in AWS

STAGE 3

Data Sources

Audit Trail **AWS**

Security Challenge

It's an all-too-common story. People host files in an AWS S3 bucket for quick transfer but forget to take them down, or use S3 buckets to backup sensitive data, but inadvertently mess up the permissions. Because misconfigured and public S3 buckets needlessly expose sensitive data to the risk of exploitation and are a common way for breaches to occur, it's important to detect when new or existing S3 buckets are set to public.

Use Cases

Security monitoring
Advanced threat detection

Category

Data exfiltration, SaaS

Splunk Solutions Required

Splunk security essentials
Splunk Add-On for Amazon Web Services
Splunk Simple Search Assistant

SPL Difficulty

Medium

How to Implement

The search for public S3 buckets is facilitated by normalized data, which is mapped to the Common Information Model. The Splunk Add-On for Amazon Web Services provides visibility of various AWS service components, including events from the CloudTrail Service and S3 buckets. Assuming you use the AWS Add-on for Splunk to pull these logs in, this search should work automatically for you without issue. While implementing, make sure you follow the best practice of specifying the index for your data.

Alert Volume

Very low

Known False Positives

There are two types of undesired alerts that can occur from this search. They occur when someone:

1. Intentionally creates a public bucket. You may wish to whitelist marketing employees who do this on a regular basis, or create a policy for how to create a public bucket so that you can exclude intentional public buckets from searches.
2. Creates a bucket that is public momentarily, but then switches it to private mode.

How to Respond

When this search fires, you will want to start your incident response process and investigate the actions taken by this process.

Detect Lateral Movement with WMI Help

To detect lateral movement with WMI we first load our sysmon EDR data. Any other process launch logs with the full command line will also suffice. We look for any instances of Windows management instrumentation command-line (WMIC) being launched (EventCode 1 indicates a process launch), and filter to make sure our suspicious fields are in the CommandLine string.

How to Respond

When an alert for a public S3 bucket fires, there are three questions that should be asked:

1. Is the S3 bucket still public?
2. Are the files public?
3. What is in the bucket?

The first question is easy to answer–just search your logs for the bucket name and "PutBucketACL." You will see any subsequent ACL changes. The second and third questions are trickier, and requires that server access logging is turned on for the S3 bucket (not done by default, and it is pretty inconvenient, so don't bet on it).

If you have a corporate AWS environment, prioritize analyzing any open S3 buckets. You may even wish to automate the remediation of them through AWS functions.

Detect Public S3 Bucket in AWS Help

To search for public S3 buckets using live data, we use the simple searchand the below search processing language. The live search operates on AWS Cloudtrail logs, filtering for the PutBucketAcl events that occur when bucket permissions are changed, and filtering for any that include AllUsers. The screenshot shows the results of a search on demo data.

```
index=* sourcetype=aws:cloudtrail AllUsers
eventName=PutBucketAcl
| spath output=userIdentityArn
path=userIdentity.arn
| spath output=bucketName
path="requestParameters.bucketName"
| spath output=aclControlList path="requestParameters.
AccessControlPolicy.AccessControlList"
| spath input=aclControlList output=grantee path=Grant{}
| mvexpand grantee
| spath input=grantee
| search "Grantee.URI"=*AllUsers
| table _time, Permission, Grantee.URI, bucketName,
userIdentityArn | sort - _time
```

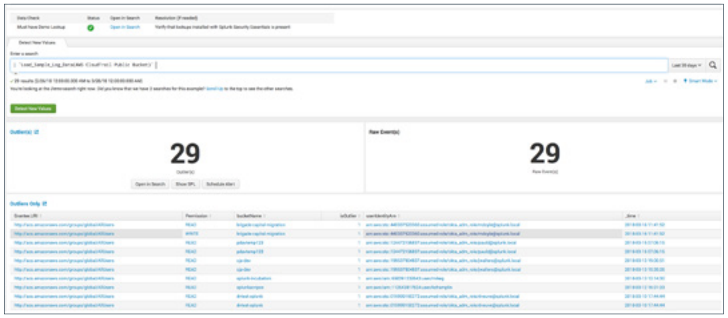

Find Multiple Infections on Host

STAGE 1

MITRE ATT&CK Tactics

Initial Access

Execution

MITRE ATT&CK Techniques

Drive-by Compromise

Spearphishing Attachment

Spearphishing Link

User Execution

Data Sources

Anti-virus

Anti-malware

Security Challenge

Viruses happen, but multiple viruses occurring on a single host all at once are a greater concern. Such activity could indicate an exploit kit that tries several techniques, where some might succeed, or a host with multiple unrelated viruses. Traditional anti-malware products can be effective in detecting known malware, but they can fail when faced with new or evolving malware types. Because malware variants can provide a backdoor to internal systems, allow for long-term persistence or exfiltrate data, you should immediately prioritize and investigateV malware-infected hosts to determine what else might not have been caught.

Use Case

Security monitoring

Category

Endpoint compromise

Splunk Solutions Required

Splunk Security Essentials
Splunk Common Information Model Add-on
Splunk Simple Search Assistant

SPL Difficulty

Basic

How to Implement

Detecting hosts with multiple infections requires the collection of logs from an antivirus solution. With Symantec logs onboard, for example, this search should work easily. If you have a different antivirus product, you can easily adapt the field names and sourcetypes for that product to the search criteria–especially if you use a Splunk Add-on that maps them to the Common Information Model (search on Splunkbase).

Alert Volume

Low

Known False Positives

No known false positives.

How to Respond

When multiple infections occur on the same host, your response plan should be the same as any malware event, just with greater urgency.

Find Multiple Infections on Host Help

To find hosts that have logged multiple infections in a short period of time using live data, our example uses the simple searchand the below search processing language. First, we bring in our basic dataset, Symantec Endpoint Protection Risks, over the last 24 hours. While there are several approaches to grouping events–and stats is the fastest–we're using transaction because it's the easiest. This will let us group all the events based on the Computer_Name. Finally, we can filter for whether there are at least three events that spanned at least a few minutes. The screenshot below shows the results of a search on demo data.

```
index=* sourcetype=symantec:* earliest=-24h
| transaction maxpause=1h Computer_Name
| where eventcount >=3 AND duration>240
```

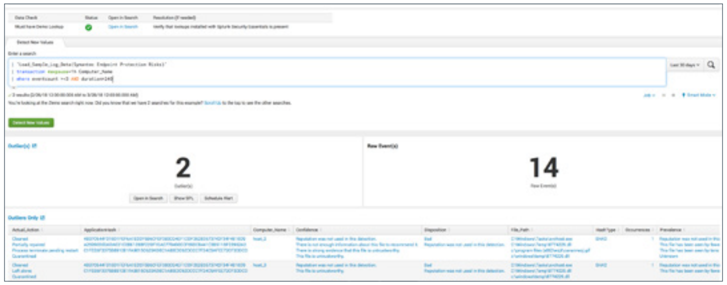

Advanced Threat Detection

Detect Connection to New Domain

STAGE 2

MITRE ATT&CK Tactics

Exfiltration

Command and Control

MITRE ATT&CK Techniques

Exfiltration Over Command and Control Channel

Exfiltration Over Alternative Protocol

Standard Application Layer Protocol

Data Sources

Web Proxy

NGFW

Security Challenge

In most organizations, the domains that users visit today have a tremendous amount of overlap with the ones they visited yesterday. But what about the small percentage of domains that were requested from your network today, but are not previous destinations for your systems? Sure, there's going to be some legitimate traffic going to a few domains today that haven't been seen on the network before, but it's likely to be a small percentage of the overall set of domains. The remainder of these new, never-before-seen domains represent a potential threat.

Knowing when users browse to new domains can be relevant in various scenarios, but the primary one is when your system connects to an attacker-controlled domain that is used as a hub for command and control communications or running a staging server for data exfiltration or delivery of malware. If you believe that a host is infected, investigating to see whether it hit new domains is a great indicator to check.

Use Case

Advanced threat detection

Category

Command and control, data exfiltration

Splunk Solutions Required

Splunk Enterprise
Splunk Security Essentials
Splunk URL Toolbox
Splunk Simple Search Assistant

SPL Difficulty

Medium

How to Implement

This method of anomaly detection tracks the earliest and latest time for any arbitrary set of values (such as the first logon per user + server combination, or first view per code repository + user combination or first Windows event ID indicating a USB key usage per system). With normal usage, you'd check to see if the latest value is within the last 24 hours and alert if that's the case. This is a major feature of many security data science tools on the market (though not Splunk UBA) that you can get easily with Splunk Enterprise.

Implementing this search is relatively straightforward, as it expects CIM-compliant data. Start by ingesting your proxy data (or other web browsing visibility, such as stream:http or Bro), and make sure there is a uri field. The only other step is to make sure that you have the URL Toolbox app installed, which allows Splunk to parse out the domains. When scaling this search to greater volumes of data (or more frequent runs), we recommend leveraging acceleration capabilities.

Alert Volume

Very high

Known False Positives

In most organizations, the percentage of new domains is small. However, if you sent all of these alerts to the analysts for investigation, it would be overwhelming because the majority of “new domain alerts” will be triggered by legitimate traffic. While there are no known false positives, per se, the value of any given “new domain” alert is so small that you want to treat these alerts differently from most correlation searches. These are mostly appropriate for contextual data or to correlate with other indicators.

How to Respond

New domain events are generally best viewed as contextual data for another event; for example, uncleaned malware, new services or unusual logins. The easiest way to accomplish this is to record the events in a summary index, and then include searching that index as a part of your investigative actions. Splunk Enterprise Security customers can do this easily with the Risk Management Framework. By creating a risk indicator adaptive response action when saving this search, it will then adjust the risk score of the assets involved and show up in Investigator Workbench when you analyze an asset. Ultimately, to analyze the efficacy of any given alert here, we recommend looking up the domains in an open source intelligence resource such as VirusTotal or ThreatCrowd.

Detect Connection to New Domain Help

To detect connections to new domains using live data we use the simple searchand the below search processing language. First, we bring in our proxy dataset, leveraging Common Information Model fields and filtering for just events that actually have a URI. Next we use URL Toolbox to parse out the domain from the URL. Finally, we exclude IP addresses from our search using the regex filtering command. This is an optional step, but we’ve found that the value to noise ratio when including IP addresses can be quite high given that some applications will connect to many ephemeral AWS instance IPs for normal operations. Lastly, we use the stats command to calculate what the earliest and the latest time is that we have seen this combination of fields and check to see if the earliest time we saw this event was within the last day (aka, brand new). The screenshot below shows the results of a search on demo data.

```
tag=web url=*
| eval list="mozilla" | `ut_parse_extended(url,list)`
| regex ut_domain!="^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}$"
| stats earliest(_time) as earliest latest(_time) as
latest by ut_domain, sourcetype
| where earliest >= relative_time(now(), "-1d@d")
```

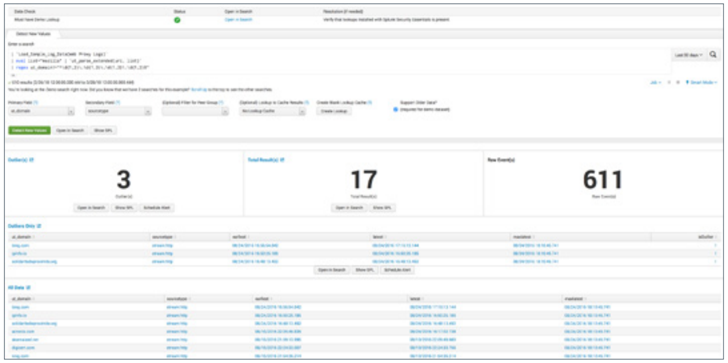

Find Emails With Lookalike Domains

STAGE 4

MITRE ATT&CK Tactics

Initial Access

MITRE ATT&CK Techniques

Spearphishing Link

Data Sources

Email

Security Challenge
Emails with lookalike domains are a common phishing tactic. Some attackers will switch easily mistakable letters, such as splunk.com receiving an email from spiunk.com. Or, they may use a believable subdomain (such as .help.com, .support, etc.) The problem is that people are more likely to open an email when they think it has been sent by a legitimate source. With the spoofed email, the difference is almost imperceptible.

Use Case
Advanced threat detection

Category
Endpoint compromise, SaaS

Splunk Solutions Required
Splunk Search Assistant
First Time Seen Assistant
UTL Toolbox app

SPL Difficulty
Advanced

How to Implement
Implementing this search is generally fairly straightforward. If you have CIM-compliant data onboarded, this search should work out of the box. You are, however, always better off specifying the index and sourcetype of your email data–particularly when you have multiple email log sources, such as a perimeter email security appliance and a core Exchange environment. It should work like a charm if you have installed the URL Toolbox and have the right index, sourcetype and src_user field.

Alert Volume
Very low

Known False Positives
This search will search through incoming emails for any domains similar to the domains usually requested by your organization, much like running dnstwist on a domain name. If there are any incoming emails with source domain names that are similar to–but not the same as–those generally seen on a daily basis, it's possible for the search to create alerts that could be false positives. One might imagine a scenario where a company that manufactures wooden planks for pirate ships, plank.com, emails their sales rep at splunk.com. That would create a Levenshtein distance of two (where the 'a' in plank becomes a 'u' and we get an extra 's') and it would be flagged (Arrrr!). To reduce the number of false alerts, you could filter known examples out of the search, or you could pipe this into a First Time Seen detection to automatically remove past examples.

How to Respond

When this search returns values, initiate your incident response process and capture the time of the event, the sender, recipient, subject of the mail and any attachments. Contact the sender. If it is authorized behavior, document that this is authorized and by whom. If not, the user credentials may have been used by another party and additional investigation is warranted.

Find Emails With Lookalike Domains Help

To find emails with lookalike domains in live data we use the Simple Search Assistant, URL Toolbox and the below search processing language. We start by pulling email logs, where we have a source address, and aggregate per source address. Next, we extract the domain and aggregate per actual domain that we will analyze. We also filter out any domains that we own and from which we expect to receive email. Using the free URL Toolbox app we parse out subdomains from the top level domains. Because the field we are going to pass to the Levenshtein algorithm is domain_detected, we add each subdomain to the multi-value field domain_detected. URL Toolbox is given two multi-value fields, and it does the cross checking to calculate the Levenshtein score for each combination. We pull out the lowest score from this group. Finally, we filter for a Levenshtein score less than three. The screenshot below shows the results of a search on demo data.

```
index=* sourcetype=cisco:esa* OR
sourcetype=ms:o365:*:messagetrace OR
sourcetype=MSExchange*:MessageTracking OR
tag=email src_user=*
| stats count by src_user
| rex field=src_user "\@(?.*)"
| stats sum(count) as count by domain_detected
| eval domain_detected=mvfilter(domain_detected!=
"mycompany.com" AND domain_detected!="company.com" AND
domain_detected!="mycompanylovestheenvironment.com")
| eval list="mozilla" | `ut_parse_extended(domain_
detected, list)`
| foreach ut_subdomain_level* [eval orig_domain=domain_
detected, domain_detected=mvappend(domain_detected, '<>'
. "." . ut_tld)]
| fields orig_domain domain_detected ut_domain count
| eval word1=mvappend(domain_detected, ut_domain),
word2 = mvappend("mycompany.com", "company.com",
"mycompanylovestheenvironment.com")
| lookup ut_levenshtein_lookup word1 word2 | eval ut_
levenshtein= min(ut_levenshtein)
| where ut_levenshtein < 3
| fields - domain_detected ut_domain | rename orig_
domain as top_level_domain_in_incoming_email word1 as
domain_names_analyzed word2 as company_domains_used
count as num_occurrences ut_levenshtein as Levenshtein_
Similarity_Score
```

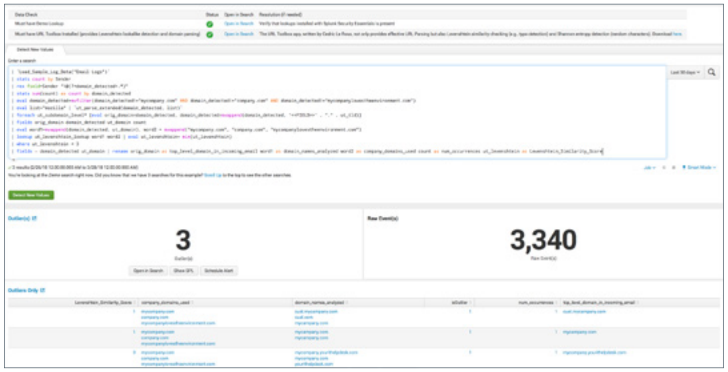

Data Check	Status	Open in Search	Resolution (if needed)
Must have Demo Lookup	✓	Open in Search	[illegible]
Must have URL Toolbox Installed (provides Levenshtein lookalike detection and domain parsing)	✓	Open in Search	The URL Toolbox app, written by [illegible], not only provides effective URL Parsing but also Levenshtein similarity checking [illegible] and Shannon entropy detection [illegible]. Download here.

Levenshtein_Similarity_Score	company_domains_used	domain_names_analyzed	isOutlier	num_occurrences	top_level_domain_in_incoming_email
1	mycompany.com company.com mycompanylovestheenvironment.com	[illegible]	1	1	[illegible]
1	mycompany.com company.com mycompanylovestheenvironment.com	[illegible]	1	1	[illegible]
2	mycompany.com company.com mycompanylovestheenvironment.com	[illegible]	1	1	[illegible]

SOC Automation

Automate Malware Investigations

STAGE 5

Data Sources

Authentication **Windows Security**

Security Challenge
When the same malware occurs on multiple systems, you may be on the brink of a major incident (seen frequently with worms, ransomware and broad phishing campaigns). Investigating and responding to malware alerts can take 30 minutes or more—each. By automating this investigation and response, Splunk SOAR validates that the process is malicious and takes immediate action to block the hash on the infected endpoints.

Use Cases
Security monitoring
Advanced threat detection
SOC automation

Category
Endpoint compromise, lateral movement

Splunk Solutions Required
Splunk SOAR

SPL Difficulty
Not applicable

How to Implement
Ingest malware events from your data source into your SOAR platform. Perform investigative actions, like retrieving reputation intelligence on applicable IPs, urls and files to help you make decisions faster. These context gather actions are great candidates for automation. Based on your decisions, execute containment and/or remediation steps, either manually or using automation playbooks.

Alert Volume
Very low

Known False Positives
Not applicable

How to Respond
The playbook investigates and remediates malware infections on the endpoint. By automating these responses you'll save time having to respond yourself and more immediate actions will be taken to block infected endpoints. You'll begin to automate the investigations and detections around use cases, such as, hiding files and directories, shim database file creations execution of a file with multiple extensions, single letter process on an endpoint and more.

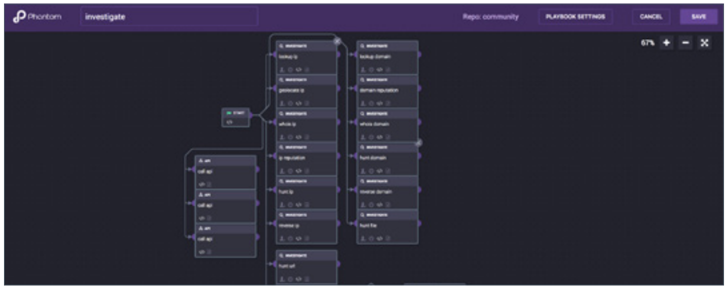

Alert Context Enrichment Help

Automate Phishing Investigations and Responses

STAGE 5

Data Sources

Audit Trail **AWS**

Security Challenge

Phishing emails can be detrimental to an organization if not detected. Investigating each email can be time consuming as an analyst may need to investigate what is in the body of the email, but also the attachments, as well as, any users that may have received the email. By automating the investigation, analysts can respond much faster to these attacks.

Use Cases

Security monitoring
Advanced threat detection
SOC automation

Category

Phishing, adversary tactics, account compromise

Splunk Solutions Required

Splunk SOAR

SPL Difficulty

Not applicable

How to Implement

Ingest suspicious emails into the SOAR platform. Perform investigative actions, like retrieving reputation intelligence on applicable IPs, urls, and files to help you make decisions faster. These context gather actions are great candidates for automation. Based on your decisions, take action to delete any copies of the phishing email from your mail system using an automation playbook.

Alert Volume

Very low

Known False Positives

Not applicable

How to Respond

Implementing and automating the correct investigations when building a playbook can enable you to limit the amount of human response necessary to carry out actions, since the response you'll want implemented in the playbook will help streamline phishing investigations enabling faster remediation and action to be taken.

Automate Phishing Investigations

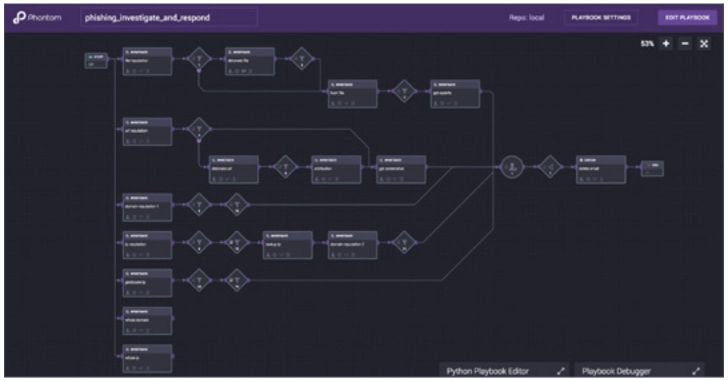

Incident Response

Detect New Data Exfil DLP Alerts for User

STAGE 3

MITRE ATT&CK Tactics

Exfiltration

MITRE ATT&CK Techniques

Exfiltration

Data Sources

DLP

Security Challenge

When a user who normally does not generate data exfil DLP alerts suddenly starts, it is more notable than a traditional alert. For crucial rules or high-privileged users, investigate these events to determine whether sensitive company intelligence is leaving the organization.

Use Case

Insider threat

Category

Insider threat

Splunk Solutions Required

Simple search assistant

SPL Difficulty

Medium

How to Implement

Implementation of this rule is straightforward–the only requirement is to be able to record which DLP alerts represent data exfiltration. That nomenclature or configuration can vary wildly from one organization to another, so this will require coordination with your DLP team. Beyond that, as long as you have a user field and a signature field defined, the search will work.

Alert Volume

High

Known False Positives

This is a strictly behavioral search, so we define “false positive” slightly differently. Every time this fires, it will accurately reflect the first occurrence in the time period you’re searching over (or for the lookup cache feature, the first occurrence over whatever time period you built the lookup). But while there are really no “false positives” in a traditional sense, there is definitely lots of noise.

How to Respond

Because this is a behavioral alert, you should generally not use this in isolation unless:

- The severity of the alert or the priority of the user dictate that it is so crucial it must be looked at in isolation, or
- Your DLP is so carefully tuned that alerts are rare.

For everyone else, most alerts should only be considered when in conjunction with other alerts, via a risk aggregation mechanism in Splunk ES or the threat models in Splunk UBA.

Detect New Data Exfil DLP Alerts for User Help

This example uses the Simple Search Assistant. Our dataset is a base dataset of DLP events. For this analysis, we are filtering for data exfiltration alerts. The screenshot below shows the results of a search on demo data.

```
index=* tag=dlp tag=incident
| stats earliest(_time) as earliest latest(_time) as latest
by user, signature
| where earliest >= relative_time(now(), "-1d@d")
```

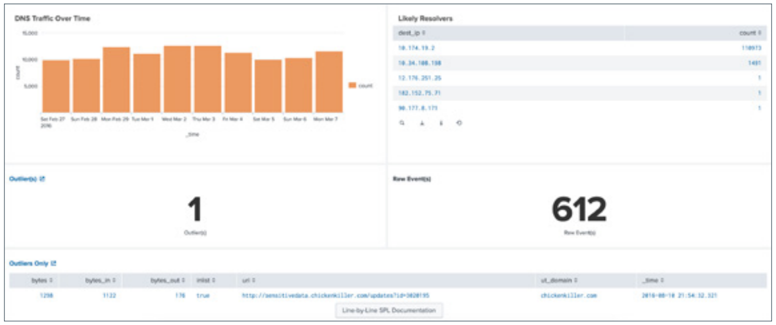

Likely Resolvers

dest_ip	count
10.174.19.2	110973
10.34.188.198	1491
12.176.251.25	1
182.152.75.71	1
98.177.8.171	1

Outliers Only

bytes	bytes_in	bytes_out	inlist	url	ut_domain	_time
1298	1122	176	true	http://sensitivedata.chickenkiller.com/updates?id=3020195	chickenkiller.com	2016-08-18 21:54:32.321

Identify Basic Dynamic DNS Detection

STAGE 1

MITRE ATT&CK Tactics

Command and Control

Adversary OPSEC

Establish & Maintain Infrastructure

MITRE ATT&CK Techniques

Dynamic DNS

Standard Application Layer Protocol

Data Sources

Web Proxy

NGFW

DNS

Security Challenge

Attackers desire flexibility in their command and control capabilities, and dynamic DNS can provide that flexibility. While there are legitimate uses of dynamic DNS (many IT professionals use it to access home networks), the risks of not monitoring the practice can be significant. Fortunately, between Splunk and a list provided by Malware Domains, finding dynamic DNS in your environment is easy.

Use Cases

Security monitoring, advanced threat detection

Category

Command and control

Splunk Solutions Required

Simple Search Assistant
URL Toolbox

SPL Difficulty

Basic

How to Implement

The first step in implementing this detection is to acquire a list of dyndns providers. Once you download a list, you will need to format it to fit the Splunk lookup format. Once you have the file in place, the rest should move on smoothly.

Alert Volume

Medium

Known False Positives

Production services that use dynamic DNS, while rare, do happen. Those will cause some base level of false positives, although they should never be business-critical services. The most common scenario for dynamic DNS is users reaching out to their homes to see their dogs via webcam, and the like. Whether to allow–and thus tune out–these users or prohibit that activity is ultimately a policy decision.

How to Respond

When this alert fires, look for the common allowable scenarios, particularly that of users who are accessing their home networks. If that does not seem to be the case:

1. Consult data from Splunk Stream or from your packet capture to determine what type of data was sent
2. Review the DNS name and IP in open source intelligence to see if there is anything of note (although that is often hard for this scenario).
3. If this is a critical host, consider endpoint logging via Microsoft sysmon or other endpoint response mechanisms to identify the process that is creating these connections.

Identify Basic Dynamic DNS Detection Help

This example uses the simple searchand the below search processing language to detect outbound communication to dynamic DNS servers on live data. First, we bring in our dataset of proxy logs. To locate dynamic DNS providers, we separate out subdomains from the registered domain using URL Toolbox. Next, we can use our lookup of ddns domains. This will add a field called inlist with the value “true” for any matches. Lastly, we can look for those records that are matches. The screenshot below shows the results of a search on demo data.

```
index=* sourcetype=pan:threat OR (tag=web tag=proxy)
earliest=-20m@m earliest=-5m@m
| eval list="mozilla" | `ut_parse_extended(url,list)`
| lookup dynamic_dns_lookup domain as ut_domain OUTPUT
inlist
| search inlist=true
| table _time ut_domain inlist bytes* uri
```

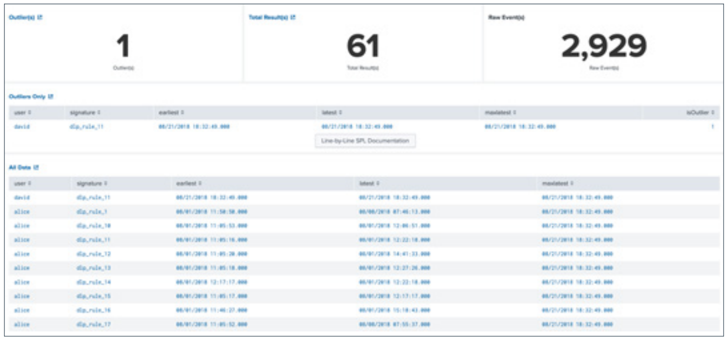

Outlier(s)
1
Outlier(s)

Total Result(s)
61
Total Result(s)

Raw Event(s)
2,929
Raw Event(s)

Outliers Only

user	signature	earliest	latest	maxlatest	isOutlier
david	dlp_rule_11	08/21/2018 18:32:49.000	08/21/2018 18:32:49.000	08/21/2018 18:32:49.000	1

All Data

user	signature	earliest	latest	maxlatest
david	dlp_rule_11	08/21/2018 18:32:49.000	08/21/2018 18:32:49.000	08/21/2018 18:32:49.000
alice	dlp_rule_1	08/01/2018 11:50:50.000	08/08/2018 07:46:13.000	08/21/2018 18:32:49.000
alice	dlp_rule_10	08/01/2018 11:05:53.000	08/01/2018 12:06:51.000	08/21/2018 18:32:49.000
alice	dlp_rule_11	08/01/2018 11:05:16.000	08/01/2018 12:22:18.000	08/21/2018 18:32:49.000
alice	dlp_rule_12	08/01/2018 11:05:20.000	08/01/2018 14:41:33.000	08/21/2018 18:32:49.000
alice	dlp_rule_13	08/01/2018 11:05:18.000	08/01/2018 12:27:26.000	08/21/2018 18:32:49.000
alice	dlp_rule_14	08/01/2018 12:17:17.000	08/01/2018 12:22:18.000	08/21/2018 18:32:49.000
alice	dlp_rule_15	08/01/2018 11:05:17.000	08/01/2018 12:17:17.000	08/21/2018 18:32:49.000
alice	dlp_rule_16	08/01/2018 11:46:27.000	08/01/2018 15:18:43.000	08/21/2018 18:32:49.000
alice	dlp_rule_17	08/01/2018 11:05:52.000	08/08/2018 07:55:37.000	08/21/2018 18:32:49.000

Compliance

Detect New Data Exfil DLP Alerts for User

STAGE 1

MITRE ATT&CK Tactics

Defense Evasion   Persistence

MITRE ATT&CK Techniques

Valid Accounts   Create Account

Data Sources

Audit Trail   Windows Security

Security Challenge
Local admin accounts are used by legitimate technicians, but they're also the Holy Grail for attackers. Once an attacker gets inside the network, he or she will most likely want admin privileges to gain undetected and unfettered access to desired accounts and assets. One easy way to achieve this is to compromise an existing account and then elevate the privileges.

Use Cases
Advanced threat detection, security monitoring, compliance

Category
Endpoint compromise

Splunk Solutions Required
Splunk Security Essentials
Splunk Enterprise

SPL Difficulty
Medium

How to Implement
First, verify that you have Windows Security logs coming in and that you have implemented account change auditing. Reference the Windows Security data source documentation if you need assistance.

Once your logs are coming in, you should be able to search for "sourcetype="WinEventLog:Security" EventCode=4720 OR EventCode=4732" to see account creation or change events. Finally, make sure that your local admin group name is "administrators" so that we are looking for the right group membership changes.

Alert Volume
Medium

Known False Positives
The only real source of false positives for this search would be help desk admins who create local admin accounts. If this is a common practice in your environment, you should filter out their admin account creation messages by excluding their usernames from the base search. If your local admin group doesn't include the term "administrators" then it would potentially generate false negatives.

How to Respond
When this search returns values, initiate your incident response process and capture:

- New account name
- Time of creation
- User accounts that created the account
- System that initiated the request
- Any other pertinent information

Contact the owner of the system. If the event is authorized behavior, document that this is authorized and by whom. If not, the user credentials may have been used by another party and additional investigation is warranted. In addition to the investigation, now is a good time to ensure that the legitimate admin accounts truly need the assigned permissions and are protected by complex and long passwords.

Detect New Local Admin Account Help

This example uses the simple searchand the below search processing language to look for newly created accounts that are elevated to local admin status. Our dataset is a collection of Windows Security logs with account creation events or account changes with group membership events. The screenshot below shows the results of a search on demo data.

```
index=* source="winEventLog:Security" EventCode=4720 OR
(EventCode=4732 Administrators)
| transaction Security_ID maxspan=180m
| search EventCode=4720 (EventCode=4732 Administrators)
| table _time EventCode Account_Name Target_Account_Name
Message
```

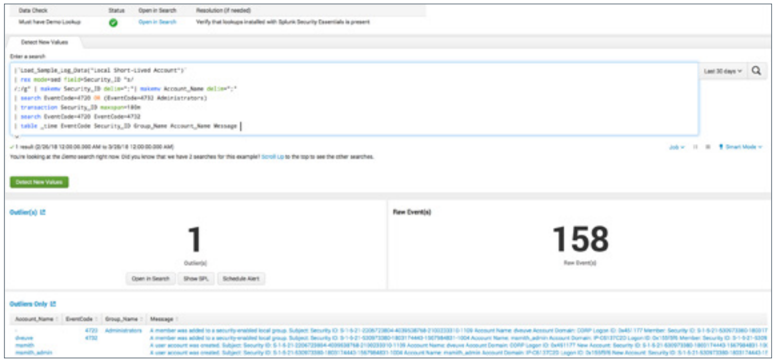

Data Check	Status	Open in Search	Resolution (if needed)
Must have Demo Lookup	✓	Open in Search	Verify that lookups installed with Splunk Security Essentials is present

```
|`Load_Sample_Log_Data("Local Short-Lived Account")`
| rex mode=sed field=Security_ID "s/
/:/g" | makemv Security_ID delim=":"| makemv Account_Name delim=":"
| search EventCode=4720 OR (EventCode=4732 Administrators)
| transaction Security_ID maxspan=180m
| search EventCode=4720 EventCode=4732
| table _time EventCode Security_ID Group_Name Account_Name Message
```

✓ 1 result (2/26/18 12:00:00.000 AM to 3/28/18 12:00:00.000 AM)

You're looking at the *Demo* search right now. Did you know that we have 2 searches for this example? Scroll Up to the top to see the other searches.

Account_Name	EventCode	Group_Name	Message
- dveuve msmith msmith_admin	4720 4732	Administrators	A member was added to a security-enabled local group. Subject: Security ID: S-1-5-21-2206723804-4039538768-2100233310-1109 Account Name: dveuve Account Domain: CORP Logon ID: 0x451177 Member: Security ID: S-1-5-21-6309733380-180317 A member was added to a security-enabled local group. Subject: Security ID: S-1-5-21-6309733380-1803174443-1567984831-1004 Account Name: msmith_admin Account Domain: IP-C6137C20 Logon ID: 0x1555F6 Member: Security ID: S-1-5-21-6309 A user account was created. Subject: Security ID: S-1-5-21-2206723804-4039538768-2100233310-1109 Account Name: dveuve Account Domain: CORP Logon ID: 0x451177 New Account: Security ID: S-1-5-21-6309733380-1803174443-1567984831-10 A user account was created. Subject: Security ID: S-1-5-21-6309733380-1803174443-1567984831-1004 Account Name: msmith_admin Account Domain: IP-C6137C20 Logon ID: 0x1555F6 New Account: Security ID: S-1-5-21-6309733380-1803174443-1

Find User Logged Into In-Scope System They Should Not Have

STAGE 4

MITRE ATT&CK Tactics

Credential Access | Privilege Escalation | Collection

MITRE ATT&CK Techniques

Valid Accounts | Data from Information Repositories | Account Manipulation

Data Sources

Authentication | Windows Security

Security Challenge

Under the General Data Protection Regulation (GDPR), organizations are required to maintain a complete audit trail about the authorized access of employees, vendors and/or processors to systems and applications that process personal data. GDPR gives individual citizens of the European Union and the European Economic Area the right to ask an organization where their data is stored and what entities are accessing the data.

To fulfill such a request, an organization will need to identify which employees, vendors and processors have accessed the personal data in question; and also identify and report on which other services regularly process said data. When processing personal data on behalf of a controller, there will also be a requirement to prove that only authorized individuals have accessed the data in question. If there is an audit trail that shows unauthorized access, then this will need to be documented and reported to the data privacy authorities.

By using data mapping that is reinforced by controls to detect violations, an organization can identify:

- Which employees, vendors and/or processors accessed the data;
- Where the data might be stored; and
- Which other services regularly process the data.

If you're processing data on behalf of a controller, this search can prove that only authorized individuals have accessed the data.

Use Cases

Insider threat, compliance

Category

GDPR, IAM analytics, lateral movement, operations

Splunk Solutions Required

Splunk Enterprise

SPL Difficulty

Basic

How to Implement

Start by using your data mapping results to build a lookup that associates systems to their GDPR category. Then do the same for users. At that point, as long as you have on-boarded CIM-compliant data, everything should go smoothly!

Alert Volume

High

Known False Positives

This search will fire when someone who is not in the documented list accesses the data. The most likely scenario for false positives is that the documented list of authorized users is out-of-date.

How to Respond

Look for indications that someone should be added to the documentation, but validate with your data protection officer (DPO) or their team before making any changes. Consider automating the update of the authorized user list and pull it from the source where your DPO holds the definitive record of authorized users. Another option is to generalize and enrich the information to departments that are allowed access by enriching the username with the department names.

Find User Logged Into In-Scope System Help

This example uses the simple searchand the below search processing language to find unauthorized users logging into in-scope systems in live data. The dataset is a collection of Windows Authentication logs that includes logins from Windows Security logs. Our search looks for the host in the GDPR categorization lookup and filters for only hosts that are in scope for GDPR. Next we look up the user in the GDPR categorization lookup, and finally look for users who don't have a matching GDPR category or who aren't authorized for any GDPR information. The screenshot below shows the results of a search on demo data.

```
index=* source=win*security user=* dest=* action=success
| bucket _time span=1d
| stats count by user, dest
| lookup gdpr_system_category.csv host as dest OUTPUT
category as dest_category | search dest_category=*
| lookup gdpr_user_category user OUTPUT category as
user_category
| makemv delim="|" dest_category | makemv delim="|"
user_category
| where isnull(user_category) OR user_category !=
dest_category
```

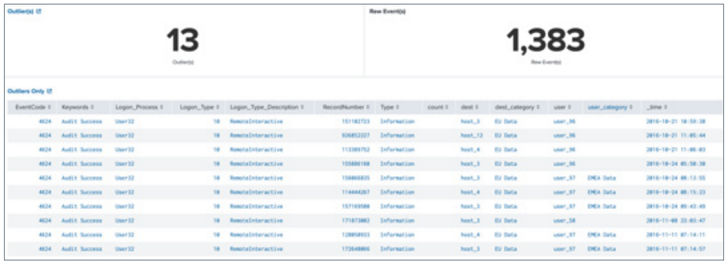

Outlier(s)

13

Outlier(s)

Raw Event(s)

1,383

Raw Event(s)

Outliers Only

EventCode	Keywords	Logon_Process	Logon_Type	Logon_Type_Description	RecordNumber	Type	count	dest	dest_category	user	user_category	_time
4624	Audit Success	User32	10	RemoteInteractive	151182723	Information		host_3	EU Data	user_96		2016-10-21 10:59:38
4624	Audit Success	User32	10	RemoteInteractive	926852227	Information		host_12	EU Data	user_96		2016-10-21 11:05:44
4624	Audit Success	User32	10	RemoteInteractive	113389752	Information		host_4	EU Data	user_96		2016-10-21 11:06:03
4624	Audit Success	User32	10	RemoteInteractive	155886168	Information		host_3	EU Data	user_96		2016-10-24 05:58:38
4624	Audit Success	User32	10	RemoteInteractive	156066835	Information		host_3	EU Data	user_97	EMEA Data	2016-10-24 08:13:55
4624	Audit Success	User32	10	RemoteInteractive	114444267	Information		host_4	EU Data	user_97	EMEA Data	2016-10-24 08:15:23
4624	Audit Success	User32	10	RemoteInteractive	157169588	Information		host_3	EU Data	user_97	EMEA Data	2016-10-24 09:43:49
4624	Audit Success	User32	10	RemoteInteractive	171873082	Information		host_3	EU Data	user_58		2016-11-08 23:03:47
4624	Audit Success	User32	10	RemoteInteractive	128858933	Information		host_4	EU Data	user_97	EMEA Data	2016-11-11 07:14:11
4624	Audit Success	User32	10	RemoteInteractive	173648866	Information		host_3	EU Data	user_97	EMEA Data	2016-11-11 07:14:57

Fraud Analytics and Detection

Detect Compromised User Accounts

STAGE 1

Data Sources

Application Logs

Web Access Logs

Security Challenge

Be it banking, credit card, email, medical or any number of accounts from myriad service providers, fraudsters can take over online accounts without your recognizing it. By leveraging phishing, spyware and/or malware scams, attackers obtain vital credential information to gain access to accounts. Some main uses of account takeovers are around credit card fraud, use of account entitlements and use of account subscriptions. By posing as the real customer, fraudsters can change account details, make purchases, withdraw funds and leverage the stolen information to access other accounts and even more sensitive data. Depending on the activity hackers may not make changes to an account and potentially will use it at the same time as the owner.

Use Cases

Fraud analytics and detection

Category

Account takeover, password list attack (credential stuffing)

Splunk Solutions Required

Splunk Enterprise

SPL Difficulty

Medium to high

How to Implement

Determine critical user accounts data and that the fields are properly extracted. You'll also want to implement security enhancements, such as blocking bad authentications, two-factor authentications, use of Captcha on all authentications, machine learning or biometrics detections are some things you can implement. By having well thought-out enhancements it makes it more difficult to create a work around and adding any amount of friction can help prevent unauthorized access. Things like rate limiting, IP blocking and blocking bad requests can also help to limit the scale of attack.

Alert Volume

Medium

Known False Positives

None

How to Respond

These are mostly brute force attempts to take over user accounts. Investigate attacking IPs and subnets and adjust firewall rules accordingly to minimize the potential of account takeover. Notice spikes on a time chart and investigate accounts that are experiencing a high volume of attacks.

Detect Compromised User Accounts Help

Use web logs to show the behavior of a user or IP addresses, and authentication logs to help you know what accounts were actually compromised. This provides useful information in looking at high failure rates. Other account logs can also help understand if any other changes have been made, such as email changes. The data must contain information about login attempts and flag whether the attempt succeeded or failed. The search processing language is shown below. The screenshot below shows the results of a search on demo data.

```
index=web-logs action=login result=failure
| stats count, sparkline as trend by src_ip | where count>5
| sort - count
| table _time src_ip trend count
```

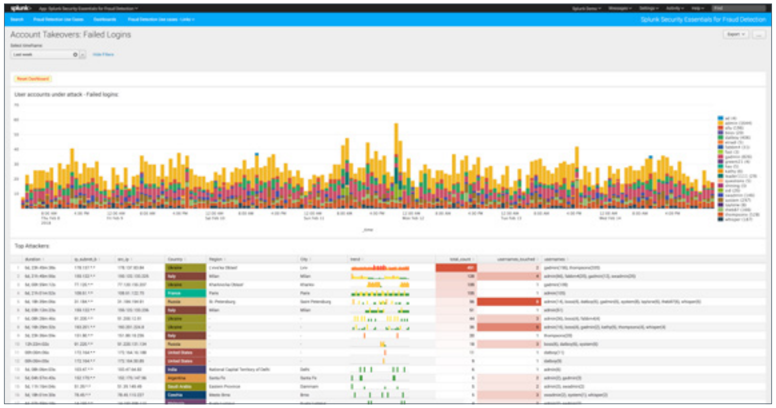
Splunk Security Essentials for Fraud Detection
Account Takeovers: Failed Logins
User accounts under attack - Failed logins
Top Attackers

Find Anomalous Healthcare Transactions

STAGE 1

Data Sources

Application Logs

Security Challenge

More than 400 people across the country have been charged with participating in prescription drug claim scams. Such fraud can impact regulations and requirements, making it more difficult for providers to conduct daily business and for customers to obtain prescriptions that are actually needed. This search finds nationwide and statewide anomalies in prescription drug claims.

Use Cases

Fraud analytics and detection

Category

Account takeover

Splunk Solutions Required

Splunk Enterprise
Splunk Machine Learning Toolkit (MLTK),
Splunk Stream

SPL Difficulty

Medium

How to Implement

Datasets may be downloaded from https://data.cms.gov/. Data comes in CSV format, making it easily ingested. You can download the app in order to view the dashboard and drilldown for source SPL. However, the app already comes with CMS datasets packaged.

Alert Volume

Medium

Known False Positives

Results are shown as anomalies and outliers. There is no definitive indication of whether the providers shown are fraudulent or not. However, upon further research, we found that in many cases anomalous providers (especially those prescribing opioids in large quantities) were involved in questionable business practices, sometimes years after datasets were published.

How to Respond

Clicking on a provider name will open the detailed profile analysis dashboard. This allows you to investigate detailed prescription data and confirm that a provider's prescribing behavior is not matching the peer group behavior within a "nationwide provider profile versus this provider profile" charts.

Find Anomalous Healthcare Providers Help

Anomalies are shown on a map. Clicking on the yellow circle will show summary data on a specific anomaly. Clicking on a provider name will open a detailed profile analysis dashboard with specific data related to a given provider.

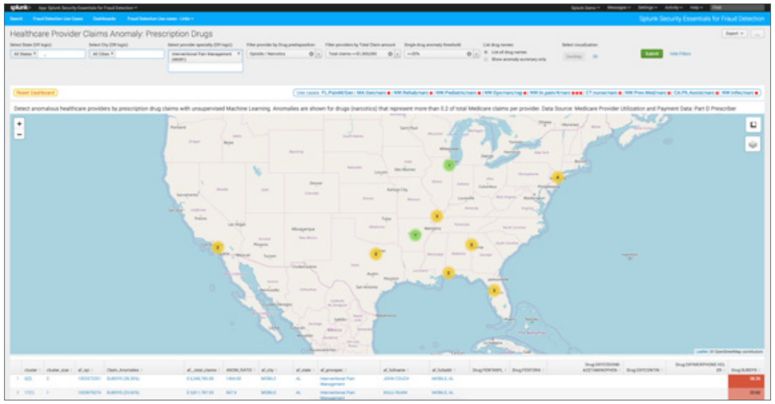

Healthcare Provider Claims Anomaly: Prescription Drugs

Detect anomalous healthcare providers by prescription drug claims with unsupervised Machine Learning. Anomalies are shown for drugs (narcotics) that represent more than 0.2 of total Medicare claims per provider. Data Source: Medicare Provider Utilization and Payment Data: Part D Prescriber

Insider Threat Detection

Detect Large Web Upload

STAGE 1

MITRE ATT&CK Tactics

Exfiltration

MITRE ATT&CK Techniques

Exfiltration Over Command and Control Channel

Exfiltration Over Alternative Protocol

Data Sources

Web Proxy

NGFW

Security Challenge

Data exfiltration usually occurs over standard channels these days, with insiders uploading data to Google, Dropbox, Box, smaller file sharing sites or even unlisted drop sites. Because HTTPS is always allowed out, exfiltration becomes relatively easy in most organizations.

Use Cases

Security monitoring, insider threat

Category

Data exfiltration

Splunk Solutions Required

Splunk Enterprise
Splunk UBA
Splunk simple search

SPL Difficulty

Basic

How to Implement

This search should work immediately for any Palo Alto Networks environment, and can be easily adapted to apply to any other source of proxy visibility. This includes dedicated proxies, along with network visibility tools such as Splunk Stream or Bro. Just adjust the sourcetype and fields to match, and you will be good to go.

Alert Volume

Medium

Known False Positives

This search will fire for many innocent occurrences, such as uploading vacation photos and so on. Many organizations will try to filter this down by focusing on users who are on a watchlist, either because they have access to sensitive data (executives, scientists, etc.) or because of employment reasons (performance plan, notice given, contract ending, etc.). These watchlists can be implemented by using lookups.

How to Respond

When this alert fires, it will usually do so for perfectly legitimate reasons (uploading vacation photos, etc.). In response, many analysts will look to see where the data was sent and whether the user has previously uploaded data to that site. Often, analysts will call the user to confirm the activity, preferably with knowledge regarding the employee's status in the organization. For example, is the employee on a performance plan or reaching the end of a contract. Both scenarios would indicate a greater risk of data exfiltration. If you have SSL inspection turned on via your NGFW or DLP system for the destination site, you can sometimes see the actual files that were transferred, which can help provide context.

Detect Large Web Upload Help

This example uses the simple searchand the below search processing language. The live search uses a dataset of proxy logs and looks for any events that are larger than 35 MB. The screenshot below shows the results of a search on demo data.

```
index=* sourcetype=pan:traffic OR (tag=web
tag=proxy) OR (sourcetype=opsec URL Filtering) OR
sourcetype=bluecoat:proxysg* OR sourcetype=websense*
earliest=-10m
| where bytes_out>35000000
| table _time src_ip user bytes* app uri
```

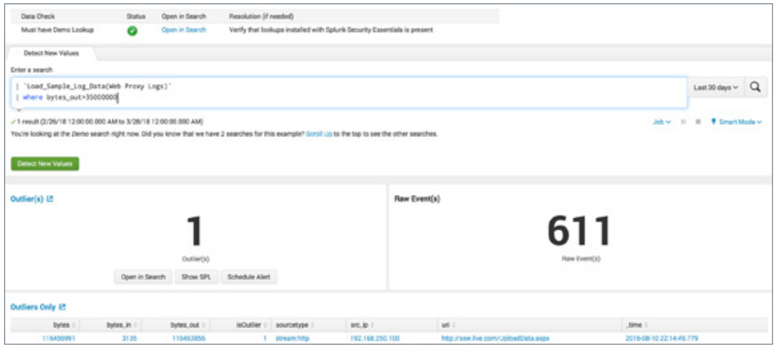

Data Check	Status	Open in Search	Resolution (if needed)
Must have Demo Lookup	✓	Open in Search	Verify that lookups installed with Splunk Security Essentials is present

Enter a search

```
| `Load_Sample_Log_Data(Web Proxy Logs)`
| where bytes_out>35000000
```

✓ 1 result (2/26/18 12:00:00.000 AM to 3/28/18 12:00:00.000 AM)

You're looking at the *Demo* search right now. Did you know that we have 2 searches for this example? Scroll Up to the top to see the other searches.

bytes	bytes_in	bytes_out	isOutlier	sourcetype	src_ip	uri	_time
116456991	3135	116453856	1	stream:http	192.168.250.100	http://sow.live.com/UploadData.aspx	2016-08-10 22:14:49.779

Detect Successful Login of Account for Former Employee

STAGE 4

MITRE ATT&CK Tactics

Privilege Escalation

Credential Access

MITRE ATT&CK Techniques

Valid Accounts

Account Manipulation

Data Sources

Authentication

Windows Security

Security Challenge

Users who have left your organization should generally not be logging in. It could mean that their credentials were compromised earlier, or it could mean that they are trying to log in to take some inappropriate actions. Either way, this is something you want to detect.

Use Cases

Security monitoring, insider threat

Category

Account compromise, insider threat

Splunk Solutions Required

Splunk Simple Search Assistant

SPL Difficulty

Basic

How to Implement

If you have followed the data onboarding guides in the Splunk Security Essentials app, this search will work immediately for you. You should generally specify the index where you are storing Windows Security logs (e.g., index=oswinsec). If you use a mechanism other than the Splunk Universal Forwarder to onboard that data, verify the sourcetype and fields that are used. The rest is simple!

Alert Volume

Low

Known False Positives

If your organization doesn't disable or remove accounts, then this search may not be actionable. If this is you, consider creating some boundaries around this behavior by specifying systems where acceptable post-termination activity can be expected to occur, such as the email environment. Also, put a detective control in place to ensure that passwords are changed when an employee goes from active to inactive. Further, try to limit the usage of accounts after the employee leaves.

How to Respond

When this alert occurs, the first thing to understand is whether it was a continuation of normal system operations (e.g., the desktop under the desk was still logged in or iPhone account still active) versus a deliberate action. Obviously, success or failure also carries weight. Finally, particularly for sysadmin-type employees in less structured organizations, make sure that there are no services or scheduled jobs running under that account where disabling the account outright might impact operations.

Detect Successful Login of Account for Former Employee Help

This example uses the simple searchand the below search processing language to detect successful authentication activity on the accounts of former employees on live data. Our dataset is a collection of anonymized Windows Authentication logs with successful logins. A lookup shows the user status, allowing us to filter for users where either the expiration is at least a day ago or they are disabled. The screenshot below shows the results of a search on demo data.

```
index=* (source=win*security OR sourcetype=linux_secure OR
tag=authentication) user=* user!="" action=success
| lookup user_account_status.csv user
| where _time > relative_time(terminationDate, "+1d")
```

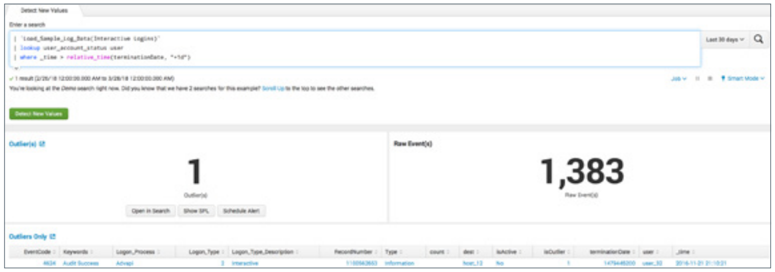

```
| `Load_Sample_Log_Data(Interactive Logins)`
| lookup user_account_status user
| where _time > relative_time(terminationDate, "+1d")
```

✓ 1 result (2/26/18 12:00:00.000 AM to 3/28/18 12:00:00.000 AM)

You're looking at the Demo search right now. Did you know that we have 2 searches for this example? Scroll Up to the top to see the other searches.

1
Outlier(s)

1,383
Raw Event(s)

EventCode	Keywords	Logon_Process	Logon_Type	Logon_Type_Description	RecordNumber	Type	count	dest	isActive	isOutlier	terminationDate	user	_time
4624	Audit Success	Advapi	2	Interactive	1150643653	Information		host_12	No	1	1479446200	user_32	2016-11-21 21:10:21

Learn **More.**

Ready to learn more about how to improve your security posture using Splunk’s analytics-driven security? Learn to solve more than 300 different security challenges for free by downloading the Splunk Security Essentials app from Splunkbase. Then work with Splunk’s security professionals and partners to implement the use cases there within your environment. Contact us to get started today!

21-13315-Splunk-Essential Guide to Security-EB-125