


Security Use Cases Enhanced by AI and ML



splunk>

This e-book is designed to help readers looking for ways to get value from implementing artificial intelligence (AI) or machine learning (ML) in Splunk, outlining example use cases that have been successfully implemented elsewhere. The business challenge, recommended approach to Splunk and value will be presented for each use case. Additionally, links to supporting information — such as customer case studies or documentation to help reproduce the use case in the Splunk environment — will also be included.

Table of contents

What is artificial intelligence and machine learning?	3
Why do organizations invest in artificial intelligence?	3
How is ML/AI used across Splunk Security?	3
Foundational elements for security	5
Considerations at the start of your AI and ML journey for security	5
 Use cases	
1. Identifying user access anomalies	9
2. Spotting potential insider threats	11
3. Detecting domain generation algorithms	12
4. Finding command line anomalies	13
5. Using ML for threat hunting	14
6. Detecting malicious patterns of network traffic	16
7. Detecting fraudulent activity.	17
8. Predicting data downtime in Splunk	19
9. Demystifying security searches with the Splunk AI assistant	20
 Get started today	
Explore more resources	21

What is artificial intelligence and machine learning?

The term machine learning (ML) is often used interchangeably with the term artificial intelligence (AI), but ML is a subfield of AI. ML is a field of computer science that develops computer systems that can autonomously learn from experience by processing the data they receive and improving the performance of specific tasks.



Artificial intelligence is the ability of a system to handle representations, both explicit and implicit, and to perform tasks that would be considered intelligent if performed by a human.



Machine learning is the ability for computer systems to use algorithms and statistical models to continuously improve the performance of specific tasks.



Deep learning is a specialized type of an ML algorithm designed to mimic a human brain's neural network, allowing machines to use massive amounts of data to learn from their own actions and improve future outcomes.



Generative AI, also known as GenAI, broadly falls under the category of machine learning. It simply refers to algorithms that can create content, including text, imagery, video, simulations, code, audio and more. Examples of generative AI include tools such as ChatGPT, DALL-E and Google Bard.

On the whole, the AI and ML space is constantly evolving. The important thing is understanding that these techniques can be applied to solve business problems, as long as there is data to train them.

Why do organizations invest in artificial intelligence?

The past few years have seen organizations having to cope with disruption on a global scale, with business resilience being tested like never before. As noted in our [Digital Resilience Pays Off](#) report, being able to prepare for change is a key factor in building resilience and thriving during uncertain times. One subject that is often associated with change and innovation is AI and ML. In terms of cybersecurity, the ability to predict and prevent incidents before they occur is one of the key areas for driving value with ML; companies that can prevent downtime have much greater resilience than those who are reactive to downtime. Organizations that adopt ML and auto remediation across all their products and services are twice as likely (66%) to be prepared for the demands of a recession, compared to those that do not (34%).

According to Forrester's Total Economic Impact report, organizations with Splunk Observability report the following:

- 70% decrease in system outages
- 75% decrease in MTTR
- 250 hours more uptime and 243% ROI

How is ML/AI used across Splunk Security?

Splunk provides a number of ways to utilize AI/ML across the product portfolio. Broadly speaking, there are two ways to use AI/ML: either through out-of-the-box features that are integrated into existing product workflows or through customization.

ML is embedded into the Splunk platform within [Splunk Cloud Platform](#) and [Splunk Enterprise](#), and available with a [Splunk Enterprise Security](#) subscription allowing users to:

- Detect anomalies, such as identifying outliers in the number of logon failures.
- Generate forecasts, for example forecasting VPN usage to identify deviations from normal activity.
- Make predictions, like predicting potential botnet activity from network activity.
- Cluster data into groups, for instance, clustering windows event logs to spot potentially malicious outliers.

These techniques can be applied via assistants that guide the user through a series of steps to train, assess and operationalize ML models. Alternatively, ML-based analytics can be created directly using Splunk's search language — Search Processing Language (SPL) — with a number of ML search commands embedded into core search and reporting, such as *predict* and *cluster*. The patterns tab in the search and reporting app also presents embedded machine learning to help identify groups of similar events in search results.

In addition to the core platform, Splunk also provides ML-powered experiences across the following products:


- Out-of-the-box ML analytics in [Splunk Enterprise Security](#) (ES), a market leading security information and event management (SIEM) solution.
- Pre-defined threat detection modeling in [Splunk User Behavior Analytics](#) (UBA), designed to identify advanced persistent threats and insider threats.
- “Mission guidance” in [Splunk SOAR](#) recommends actions and playbooks that an analyst should run and other humans can consult about an incident.
- [Splunk Attack Analyzer](#) uses EMBER models to determine maliciousness of malware samples.
- Workflows in [Splunk IT Service Intelligence](#) (ITSI) — an AIOps solution — for creating adaptive thresholds for key metrics, as well as predicting potential outages.
- Assistive wizards in [Splunk Infrastructure Monitoring](#) to detect outliers in metrics or predict when resource utilization thresholds will be crossed.

We also offer assistive intelligence experiences to provide personalized guidance.

- The [Splunk AI Assistant](#) uses generative AI to provide a chat experience that helps users author and learn SPL by interacting in plain English and providing query suggestions, explanations and detailed breakdowns.

- The [Splunk App for Anomaly Detection](#) enables Splunk users to detect anomalies in their time series data sets and metrics using powerful machine learning algorithms in just a few clicks, while providing an end-to-end operationalization workflow to streamline creating and running anomaly detection jobs.
- The Splunk platform can be extended with add-ons that are designed specifically for running ML workloads, namely the Machine Learning Toolkit and the [Splunk App for Data Science and Deep Learning](#).

- The [Splunk Machine Learning Toolkit](#) (MLTK) provides SPL commands, custom visualizations, assistants, and examples to explore a variety of ML concepts all inside the Splunk platform. Extending beyond MLTK, the Splunk App for Data Science and Deep Learning (DSDL) provides the ability to integrate advanced custom ML and deep learning systems with the Splunk platform.



Foundational elements for security

The Splunk platform is widely used for security use cases. [Gartner](#), [IDC](#) and [Forrester](#) have named Splunk Enterprise Security a leader in SIEM and security analytics. The core data platform underpins this security industry leadership by allowing the user to query and visualize security data in near real time.

Using a data schema — the [Common Information Model](#) (CIM) or the [Open Cybersecurity Schema Framework](#) (OCSF) — machine generated data can be normalized to provide a holistic view of activity from across a range of different data sources. Additionally, the user can enrich data with contextual information such as user business unit information, vulnerability data or incident management feeds.

Splunk Enterprise Security ships with a number of [out-of-the-box correlation searches](#) using ML to detect potentially risky behavior. Alongside this out-of-the-box content, [Risk-Based Alerting](#) (RBA) can be used by ES to detect unusual changes in user or system behavior that might indicate a compromise.

The [Splunk Threat Research Team](#) (STRT) publishes [Machine Learning for Security](#) content including detections and analytics stories that can easily be deployed in ES to detect sophisticated attacks.

Furthermore, Splunk UBA is designed to use unsupervised ML to detect advanced persistent threats and insider threats.

These capabilities and more are available to assist you on your journey to the modern security operations center (SOC), with both custom ML and new generative AI assistance to strengthen your security posture and bridge the security skills gap.

Considerations at the start of your AI and ML journey for security

Before starting a new AI or custom machine learning project within your security practice, Splunk recommends running a brief impact assessment to help prepare for successful results. At a minimum, the scope of the impact assessment should consider objectives, risk and execution capacity. Some of the main things to consider as part of this assessment are highlighted in the following sections.

Assess the objectives

Surveys suggest that [87% of data science projects](#) fail to make it to production, highlighting the importance of defining clear outcomes to make an ML project successful. At a high level, these outcomes often fall into the following categories:

- Increasing detection efficiency
- Reducing manual processing or minimizing human error
- Identifying previously unknown scenarios

The most successful ML projects are often tied to granular outcomes, such as increasing detection accuracy by 70% for alerts related to access anomalies or reducing manual triage time for security operations center (SOC) analysts by 50% when assessing alert storms.

When creating an ML project, developing business outcomes and success metrics are important items to consider. Typical questions to consider when determining an objective for an ML project are:

Increasing detection efficiency

- a. Are current true positive and true negative rates known for existing detections? If so, are there understood business benefits from improving these rates? For example, by improving detection accuracy, analysts will not have to spend as much time triaging false positives.

- b. Are there certain alerts that trigger frequently, generating a lot of noise for the SOC? If so, improving the accuracy of these alerts could improve the overall efficiency of the SOC.
- c. Can target benchmarks be set to reduce false positives or improve alert accuracy?

Reducing manual processing or minimizing human error

- a. Is there case management data available detailing the amount of time SOC analysts spend triaging alerts? This information can help identify alerts that could benefit from reduced triage times.
- b. Are there alerts that are ignored or closed in high volumes without triage? This may indicate situations where potential threats are being missed due to noisy alerts.
- c. Are there business objectives set for the amount of time analysts should spend improving SOC capabilities? Automating routine tasks to free up time for higher value activity is one mechanism for achieving this.

Identifying previously unknown scenarios

- a. Have there been historic security incidents that were not identified by existing triggering alerts?
- b. Are there emerging security threats that are not well understood?
- c. Have there been targeted attacks previously from malicious actors that haven't used traditional tactics, techniques or procedures (TTPs)?

Assess the risk

Guidance published by the World Economic Forum notes that risks can be associated with using AI/ML techniques. Often, these risks are related to the difficulty of explaining the processing of data, with users of AI systems often unsure how a particular output has been generated, making it difficult to determine the correct action to take based on the output.

Three areas that are important when considering risk of an AI system are:

- Visibility: How much detail is required on how the data is being processed by an AI system? ML algorithms use complex mathematics to process input data and create an output, which can often make it difficult to understand why a particular output has been generated.
- Control: What requirements for the hosting of data are in place in a given organization? In the public sector and in highly regulated industries many organizations have requirements to host certain types of data on authority infrastructure, making use of cloud services challenging.
- Failure Tolerance: What will be the consequence of a failed AI project? Investment in AI is no guarantee for success, therefore consider how much flexibility there is in the business objectives and the downstream impact on the SOC team while working on the AI project.


Each of these areas are expanded below.

The matrix to the right illustrates how some of the considerations about visibility and control of data might impact the choice of Splunk products (please refer to section A of AI Procurement in a Box Workbook for further questions to assess risk against visibility and control). Note that this is suggested guidance, and where appropriate, MLTK or an out-of-the-box use case may be preferable depending on risk appetite or the need for customization.

Additionally, the risk of failure should be considered. Investment in an AI system is not guaranteed to deliver the required outcomes. Whether it's reducing alert volumes or predicting the presence of malware, it might not be possible to generate the perfect results every time.

Therefore, consider upfront how much tolerance exists for outcomes that are below the user's definition of success. Additionally, think about current workloads and how much additional capacity exists to handle false positives or whether an investment is needed in new

processes to handle new types of alerts. For example, if investigating a system to predict potential insider threat activity, what would need to be done if the model is right only 70% of the time? Good practice is to think about the current workloads in the SOC and whether it would be able to handle triaging false positives from this type of alert — which could be handled through the use of RBA in Splunk Enterprise Security. Additionally, consider what actions the SOC can take if insider activity is predicted. Are there clear next steps that can be taken by the SOC if something is flagged?



Assess execution capacity

Provided a clear objective can be set for an AI or machine learning project, the requirements for executing the project should also be evaluated. The list below is not exhaustive or prescriptive, but presents some of the areas to think about before embarking on an experience with ML.

People:

- a. Are subject matter experts on the data available to provide guidance on the meaning and quality of the data?
- b. What ML expertise is available to help develop and support these use cases in the future?
- c. If there is no ML expertise, can partners like Splunk provide the right resources to upskill or guide the project?

Process:

- a. Who will use the analysis and how will they use it in their daily functions?
- b. How will the use case be operationalized?

Information:

- a. Is the data required for analysis already indexed in Splunk or are there plans in place for getting it indexed?
- b. Are there any special handling requirements for the data, for example does personal data need to be obfuscated?
- c. Are there blogs and content available about the use case on which the project is based?
- d. Does the output of the ML processing need to be explainable to a non-technical end user?

Technology:

- a. Is ML necessary to solve the use case or can a correlation rule or basic statistics suffice? To determine this, performance analysis comparing workloads to accuracy using different methodologies may be required; also looking at how users will consume the outputs (simpler methods are often more explainable).
- b. Will the current Splunk architecture need updating to handle special ML processing, such as introducing a dedicated search head for training ML models?

Use cases


This section outlines introductory use cases for using AI for security. Although this document focuses on use cases, typically ML-based analytics in Splunk use one of four common techniques: **anomaly detection, predictive analytics, clustering and graph analytics.**

The table to the right shows how the following use cases map to the different techniques.

For each use case example in the table, we take a closer look at the:


- Business challenge
- Splunk’s approach
- Value
- Case studies and further information

Anomaly Detection




- Deviation from past behaviors
- Deviation from peers
- Unusual change in features

Predictive Analytics




- Predict health score
- Predicting events
- Trend forecasting
- Early warning of failure

Clustering



- Identify peer groups
- Event correlation
- Reduce alert noise
- Behavioral analytics

Graph Analytics



- Most influential nodes
- Link analysis
- Community detection
- Impact analysis

Use Case	Anomaly Detection	Predictive Analytics	Clustering	Graph Analytics	Generative AI
1. Identifying User Access Anomalies	✓				
2. Spotting Potential Insider Threats	✓		✓		
3. Detecting Domain Generation Algorithms (DGA)s		✓			
4. Finding Command Line Anomalies	✓	✓			
5. Using ML for Threat Hunting	✓		✓	✓	
6. Detecting Malicious Patterns of Network Traffic	✓				
7. Detecting Fraudulent Activity	✓		✓	✓	
8. Predicting Data Downtime in Splunk	✓	✓			
9. Demystifying Security Searches with the Splunk AI Assistant					✓

1. Identifying user access anomalies



Business challenge

Research suggests that compromised credentials provide the entry point for **71% of cyber attacks against businesses**. Without recognizing what normal user behavior looks like in an organization, catching potential compromises can be difficult. Furthermore, baselining user behavior can be challenging when different departments, geographies, teams or facilities have different working patterns. If organizations are unable to quickly spot compromises and baseline user behavior, then threat actors will leverage those handicaps to easily gain access to systems undetected.

Splunk's approach

We have out-of-the-box capabilities in Splunk Enterprise Security that are designed to detect potential user compromises using ML, such as detecting potential brute force attempts by baselining failed logon attempts. Additionally, our UBA product profiles and baselines user access data to detect potential compromises.

For customers who wish to design custom rules for detecting unusual account behavior, the MLTK provides a range of algorithms for generating baselines from historic data that can be used to detect deviations from the baseline.



Value

Increase detection efficiency: Use of ML to identify user access anomalies can help to identify potential compromise to the business and mitigate risk.

Reduce manual processing: Use of ML can increase efficiency by programmatically identifying anomalies, reducing time spent analyzing data manually.

Case studies and further information

Ministry of Energy, Israel

The Israeli Ministry of Energy used the MLTK to help baseline user behavior and spot anomalies. Working in an operational technology (OT) environment, the Israeli Ministry of Energy faces challenges surrounding the length of time software is used, often being operational for five-plus years, without easy mechanisms to patch. Additionally, the software is often deployed in environments where security cannot be applied as uniformly as with enterprise IT.

Operating with a variety of current and older technologies posed a challenge to fully understanding user activity across the Israeli Ministry of Energy's systems. To spot user access anomalies in the past, security analysts would manually analyze user activity to detect the presence of anomalies. Using the MLTK, the Israeli Ministry of Energy automated the detection of these anomalies, thereby increasing team efficiency.

The Israeli Ministry of Energy followed the process of normalizing their data to the Splunk CIM to enrich their assets with contextual data such as device location and model type. The Ministry of Energy was then able to risk score assets using NIST's vulnerability database and information from a summary index detailing the device history. From this enriched data, baseline models were built using MLTK's Density Function algorithm, describing what normal authentication looked like across their OT environments. Using these models, the Israeli Ministry of Energy could identify anomalous user logon behavior, automating what used to be a manual process for their analysts.

Find out more [here](#).

Further information

Additionally, Splunk has a simple guide that describes how to create a use case for detecting user access anomalies [here](#). Provided that all correct apps installed and some authentication data in a Splunk instance is present, this how-to guide should take a few hours to implement and test.



Because of its inconsistent nature, it was the task of human analysts to spot anomalies. With the Density Function our analysts now receive more meaningful alerts and spend less time on tedious, manual work.

Efi Kaufman, Head of Big Data and Analytics,
Israel's Ministry of Energy CERT Team

2. Spotting potential insider threats

Business challenge

Insider threat is a broad topic that also examines unintentional data loss or espionage and reconnaissance. This broadness increases the difficulty of identifying risk, and pressures businesses to monitor a wide range of indicators to understand their exposure. Occasionally, it's possible that no single identifier exists for detecting insider threats. Therefore, protecting against insider threats is a challenging task, with analysts needing an understanding of what normal looks like for multiple indicators to spot potential compromise.

Splunk's approach

Splunk Enterprise Security ships with [searches](#) that use ML to identify unusual changes in RBA generated risk scores, which can be used to identify changes in user behavior that potentially indicate compromise.

Additionally, Splunk UBA utilizes multiple ML-enabled analytics to identify potential insiders, such as suspicious data exfiltration, suspicious data collection or suspicious data movement from a given user.

Many customers also use the MLTK to identify potential insider threats by creating custom analytics in their environment. Usually based on anomaly detection techniques, the custom analytics help identify users who are using unapproved software, sending suspicious communications or displaying potential flight risk behavior, such as uploading data to job websites.

Value

Increase detection efficiency: Insider threats can be costly to organizations, exposing intellectual property or sensitive data or causing reputational damage. Being able to efficiently spot potential insider threats is an important capability for many organizations to minimize risk and prevent disasters.

Reduce manual processing: Using ML-based analytics to detect insiders by directing busy security analysts toward the most risky threats can help prioritize scarce and expensive resources.

Identifying previously unknown scenarios: Insider attacks can often be “low and slow,” making detection difficult using traditional techniques. ML-based behavioral detections are able to identify less obvious attacks, flagging previously unknown indicators of compromise.

Case studies and further information

Lockheed Martin

Lockheed Martin operates across the globe with hundreds of thousands of users and serves many defense customers. The ability to identify high-risk user behavior is critical for their business operations. Lockheed Martin used MLTK to develop a custom, risk scoring framework for identifying insider threats.

Specifically, Lockheed Martin's process involves first centralizing and normalizing their data in Splunk and using correlation and atomic searches to extract high-value information. This high value data is either accelerated or sent to summary indexes to make subsequent processing more efficient. Lockheed Martin then built baselines on

top of this data using a range of ML techniques, feeding detections from these baselines into the risk score framework for alerting and security analyst triage.

More information on Lockheed Martin's use case can be found [here](#).

Science Applications International Corporation

In a similar fashion to Lockheed Martin, Science Applications International Corporation (SAIC) developed analytics using Splunk's MLTK to detect potential insider threats on behalf of their customers. By integrating their insider threat analytics more closely with Splunk Enterprise Security, SAIC was able to provide insights to SOC analysts through an interface that the analysts were familiar with already, providing MLTK-enriched insights to security analysts. More details can be found [here](#).

Further information

The Splunk App for Behavioral Profiling provides a set of workflows for operationalizing ML-driven detection and scoring of behavioral anomalies at scale in complex environments, correlated to profile and highlighting the entities which require investigation. These workflows are well suited for insider threat detection systems and can help organizations set up sophisticated methods for identifying malicious insiders.

3. Detecting domain generation algorithms

Business challenge

Recent threat research suggests that [88% of organizations experienced DNS attacks](#). DNS attacks often establish command and control (C2) with dynamic resolution via use of domain generation algorithms (DGAs). As malware families evolve, challenges will increase for defenders to detect, block and track these threats in real time. Despite these challenges, patterns in domain names are detectable using ML, making the detection of DGAs an optimal problem for ML to solve.

Splunk's approach

Classification algorithms in MLTK can be used to identify domain names created by a DGA. Furthermore, DSDL offers several more advanced methodologies for detecting complex behaviors like domains created by a DGA.

Value

Increase detection efficiency: Attackers often use C2 infrastructure for ransomware, data theft or other malicious activity, thus spotting the signs of threat actors using these types of techniques early can reduce the risk of compromise. Traditional methods such as managing IP-based or domain name-based deny lists to protect against malicious domains can be extremely time consuming when some DGAs generate upwards of 50,000 domains a day. Therefore, having effective detection methods can improve the efficiency of security operations.

Case studies and further information

Viasat

The Viasat security team faced a challenge — the generation of too many notables — with a handful of analysts unable to triage the thousands of alerts received every day. While improving their correlation search logic helped, Viasat began to see real success when they applied the MLTK to fine-tune their correlations. Moreover, Viasat developed a workflow for their security analysts to intervene with these correlation searches and used ML to help fine-tune the detections.

One particular use case, detecting DGAs, saw real benefit from using ML. By using the Shannon entropy of domain names (a method of measuring randomness of a string) and the frequency of tri-grams in the domain name (sequences that are three characters in length) compared to historic observations, Viasat trained a classification algorithm to predict whether a given domain name was potentially generated by a DGA or not. This algorithm was trained using a labeled data set, where domain names identified as generated by a DGA were tagged accordingly.

Labeling data can be a labor-intensive manual process, so the team at Viasat also implemented a workflow action in ES that allowed analysts to record whether domain names identified in a notable event were generated by a DGA or not as part of their triage. Viasat then could continue to refine and update their labeled DGA data set as part of their analyst workflows.

Read more about Viasat's experience [here](#).

Further information

Using a large pre-labelled, open source data set, the DGA app for Splunk provides an end-to-end workflow for training, testing, and deploying an ML model for detecting DGAs. The app utilizes out-of-the-box algorithms that ship with the MLTK including term frequency-inverse document frequency (TF-IDF) to break down domain names in short character sequences, principal component analysis (PCA) to reduce the number of dimensions in the data set, and a range of different classification algorithms for making the predictions. The app can be found [here](#).

In addition to the DGA app, our Splunk Machine Learning for Security team has produced an analytic for detection of potential DGAs using the DSDL. This detection relies on a pre-trained model that can be run using the DSDL app, which can offer a quick time to value for this use case depending on the need for visibility into the ML processing. More details are available [here](#).

4. Finding command line anomalies

Business challenge

Attackers looking to “live off the land” often will use common, pre-existing tools like PowerShell or the command line, with [one security vendor blocking 480,000 potentially malicious PowerShell commands](#) in one month alone. Use of native tools can make prevention and detection of attackers using these tactics difficult, with security analysts often needing to examine the specific command and context to determine if the command is malicious.

Splunk’s approach

Through use of the MLTK, users can train models for detecting potentially malicious commands. By creating benchmarks for what non-malicious commands look like in an environment, deviations from these benchmarks can indicate potentially malicious commands in the form either of normal commands run by unusual users or unusual commands being run.

Value

Increase detection efficiency: Detecting attackers attempting to live off the land early in time reduces potential damage to an organization.

Reduce manual processing: The ability to identify potentially malicious commands can improve analyst efficiency, reducing the amount of time that analysts have to spend manually investigating command line logs.

Case studies and further information

Siemens

As a Splunk security customer for over a decade, Siemens has a mature SOC team operating across multiple locations and ingesting data from thousands of sources. Siemens partnered with Splunk to explore the MLTK to extend their security monitoring with ML and baseline their data centers, identifying security relevant patterns.

Siemens successfully developed a detection method for potentially malicious commands being run in a web shell. Given that single command lines have a high degree of variance and may not be indicative of malicious activity alone, Siemens determined that not only did they have to analyze the entire command line, but they also had to look at sequences of command lines. To reduce noise, Siemens first created a list of potentially suspicious command line words to find matches. From these matches, Siemens generated sequences that were then classified via MLTK’s TF-IDF algorithm for extracting important terms in a corpus of command line sequences followed by logistic regression to determine if the command line sequences were potentially risky. For this classifier, the team created a labeled data set of known malicious command line sequences.

This use case allowed Siemens to classify over 20 million commands a day, providing improved day-to-day operations for their security analysts.

Read more about Siemens’ experiences with MLTK [here](#).

Further information

The Splunk Machine Learning for Security research team has also created detections for potentially malicious command line activity. The first of these is relatively simple, using the MLTK to generate a baseline of the typical length of a command line, helping to identify situations when command lines are unusually long. Read about this analytic [here](#).

In addition to this use case, this team has produced an analytic that uses a pre-trained model to detect potentially malicious code on the command line. The model in this analytic identifies unusual combinations of keywords found in samples of command lines where adversaries executed PowerShell code, primarily for C2 communication. For example, adversaries will leverage IO capabilities such as *streamreader* and *webclient*, threading capabilities such as *mutex* locks, programmatic constructs like *function* and *catch*, and cryptographic operations like *computehash*. The model will output a score, where anything above zero can be considered potentially malicious, where the numeric output may need to be converted into a meaningful indicator for the SOC to handle depending on their existing processes. More details on this example can be found [here](#).

5. Using ML for threat hunting

Business challenge

Attackers have diversified their tactics, techniques and procedures (TTPs) over time to evade existing security defenses. As a result, traditional detection rules, such as using threat intelligence, correlation rules or simple heuristics are often not enough to identify unusual behavior from a threat actor. If a SOC team or security analyst cannot hunt through security information, malicious actors can go undetected if using novel or unusual tactics.

Splunk's approach

At its core, Splunk is a data analytics platform for searching through machine data, which offers threat hunters the ability to search through their security log data to detect potentially suspicious events.

In addition to the core platform, Splunk UBA provides advanced and insider threat detection using unsupervised ML, helping organizations to find unknown threats and anomalous user behavior across devices and applications.

Furthermore, the MLTK can support threat hunting by applying a number of different algorithms and guided workflows to identify anomalies.

Value

Increase efficiency: SOC teams gain efficiency when they have good threat hunting practices and a better understanding of what baseline looks like for their environment in certain situations, reducing triage time in some scenarios.

Identify previously unknown scenarios: Threat hunting in Splunk can help improve defenses by identifying sophisticated attacks and creating new indicators of compromise.

Case studies and further information

Ministry of Energy, Israel


The [Ministry of Energy, Israel](#), operates in an OT environment, and faces challenges matching OT data sources to threat intelligence. Additionally, many of the attacks they face are 'low and slow' to avoid detection, with attackers using legitimate administration tools to hide within the environment for longer.

Despite the challenges they face, the environments that the Ministry of Energy, Israel monitors are usually fairly static, often with changes occurring seasonally. Through experimentation with the MLTK, the

Israeli Ministry of Energy quickly developed a proof-of-concept for threat hunting in their OT environments, using dimensionality reduction and clustering techniques on Windows event logs to identify unusual patterns of behavior in some of their environments. By partnering with Splunk, the Israeli Ministry of Energy was able to develop their proof of concept into a threat hunting workflow, where Windows event logs were fed into a ML pipeline to visualize potentially malicious behavior.

This approach enabled the Israeli Ministry of Energy to identify suspicious behavior in some of their environments, as well as identify periods of time when maintenance occurred outside of the usual cycle.

More details of the Israeli Ministry of Energy's threat hunting experience can be found [here](#).



Saudi Aramco

Saudi Aramco is another customer monitoring an OT environment who has seen success from ML-based threat hunting. One of the largest companies in the world, Saudi Aramco operates a huge number of assets across their estate, generating large volumes of data with different behavior profiles.

Viewing security as a big data problem, Saudi Aramco used the MLTK to codify a number of threat detection rules, such as DLL injection anomalies or PowerShell anomalies. Furthermore, Saudi Aramco fed the outputs of these anomaly detection rules into a threat scoring model to help their analysts to understand and prioritize threats across the environment.

Further information about their use cases can be seen in the following [webinar](#).

Siemens

Siemens used the MLTK and graph analytics to understand the relationships between their use cases. For example, Siemens identified patterns in C2 activity and system compromises by analyzing the connected components across triggered alerts.

Read more about their use of graph analytics [here](#).

Further information

The Splunk [SURGe](#) security research team has published guidance on the Prepare, Execute and Act with Knowledge (PEAK) [threat hunting framework](#). This includes information on model assisted threat hunting, where hunters use ML techniques to create models of known good or known malicious behavior and look for activity that deviates from or aligns with these models. Think of this as almost like a hybrid of the hypothesis-driven and baseline types, but with substantial automation from the ML.

Hunting through log data to detect use of unwanted protocols, such as DNS over HTTPS (DoH), can be difficult. In the talk linked to below, a methodology for hunting through HTTPS data for DoH traffic using ML is presented [here](#) by members of Splunk's data science team. Additionally, DSDL contains an example technique for hunting through JA3 hashes to identify patterns

6. Detecting malicious patterns of network traffic

Business challenge

Network security monitoring is an important function within a SOC and is used to detect potential infiltration, exfiltration and lateral movement by malicious actors. However, understanding and baselining what normal network behavior looks like in an organization can be challenging, with potentially thousands of enterprise endpoints and API calls creating noise that attackers can use as a hideout.

Splunk's approach

Splunk Enterprise Security contains out-of-the-box correlation searches that utilize the MLTK to detect unusual volumes of network activity or substantial increases in port activity. Additionally, the Splunk UBA product ships with many ML-enabled use cases for network security monitoring, such as potential data exfiltration identification.

In addition to these products that provide ready-to-use content, the MLTK can be used to produce analytics that detect unusual patterns of network traffic.

Value

Increase detection efficiency: The ability to baseline typical network behavior and identify times when there is anomalous activity on the network can help detect a range of potentially malicious behavior, such as DDoS attacks, botnet activity or the presence of malware.

Reduce manual processing: Furthermore, behavioral-based detections can improve analyst efficiency, reducing the amount of time that security analysts spend manually triaging network behavior after an alert triggers.

Case studies and further information

Siemens

As described [above](#), Siemens partnered with Splunk to develop a set of ML-based security detections to augment their SOC. Siemens was able to baseline proxy communications to identify outliers in outbound web proxy communications. This baseline allowed Siemens to detect potential malware communications, in particular flagging HTTP tunneling of SSH traffic.

To develop this use case, Siemens identified eight features from their proxy logs in which they wanted to identify outliers, including: bytes in, bytes out, and the number of distinct IP addresses visited by a given source IP over hourly intervals. For each of these eight features, a model was trained using the MLTK's DensityFunction to baseline expected behavior for each feature. From these baselines,

outliers are detected for each feature. Every hour, an anomaly score is created for each source IP based on the number of outliers across all eight features. Creating a set of drill downs from the anomaly detection allowed analysts quickly to triage potentially malicious web proxy traffic.

Find out more about Siemens experience with the MLTK [here](#).

Further information

Additionally, Splunk has a simple guide that describes how to detect network traffic anomalies based on the amount of data being transferred between source and destination IP pairs [here](#). Provided that all the correct apps are installed and some network traffic data exists in a Splunk instance, this how-to guide should take a few hours to implement and test.

Additionally, a number of Splunkbase apps are available that provide workflows for detecting unusual network behavior. Examples include the Botnet App for Splunk, which uses an open source data set and a guided workflow to help users to train a set of classification models for detecting potential botnet activity on their network. This app can be found [here](#).

7. Detecting fraudulent activity

Business challenge

With over half of organizations reporting themselves as victims of fraud with losses of over [\\$42 billion](#), the ability to detect fraudulent activity is important for minimizing losses as well as maintaining brand reputation. However, spotting fraudsters can be difficult when fraudsters have insider knowledge on how to subvert defenses or use rapidly pivoting tactics to evade detection. Therefore, the ability to baseline normal activity is critical to detecting sophisticated fraudsters.

Splunk's approach

With Splunk's MLTK, users can create a wide range of detections for fraud using algorithms and guided workflows for creating behavioral detections. These behavioral detections include the ability to identify unusual patterns in customers transactions such as unusually higher transaction values, spotting outliers in multi-dimensional data sets using clustering such as identifying account takeover activity based on historic user behavior, transaction activity, and types of activity.

Value

Increase detection efficiency: The ability to detect fraud early can save money as well as protect an organization's brand reputation.

Reduce manual processing: Additionally, by baselining typical behaviors in an organization, fraud analyst efficiency can be improved, alerting analysts to suspected fraud instances only and reducing triage time by presenting behavioral indicators that might previously have been identified manually.

Identify previously unknown scenarios: Fraud activity often exploits unrealised weaknesses in an organization's defenses, which are not being monitored. Using ML-based approaches that can baseline typical activity can uncover unknown weaknesses.

Case studies and further information

Aflac

Aflac is a supplemental insurer, providing wraparound cover for healthcare emergencies. As a victim of account takeover fraud in 2016, Aflac recognised the need to improve their defenses against this type of malicious activity.

Working with over 30 features to be used to indicate potential fraudulent activity, Aflac developed a risk scoring methodology to identify risk for their customer accounts. Since Aflac did not have a labeled data set, they attempted to determine risk scores manually for each feature, often leading to accounts with high activity having high risk scores. Turning to the MLTK, Aflac used clustering to profile their risk index, identifying clusters of similar policy holders. This clustering analytic helped to find outlier accounts committing potentially fraudulent activity.

Read more about Aflac's experience [here](#).

NewYork-Presbyterian Hospital

Originally using Splunk for security use cases, the NewYork-Presbyterian Hospital realized they also could use Splunk to detect potentially fraudulent activity around the use of controlled substances and other medicines.

Using the core Splunk platform and the MLTK, the NewYork-Presbyterian Hospital partnered with Splunk to create a controlled substance monitoring platform. Using techniques like clustering to identify unusual activity including unauthorized authoring of prescriptions, the NewYork-Presbyterian Hospital is helping to safeguard members of the public.

Read more about this partnership [here](#).

BlockFi

BlockFi provides financial products for exchanging cryptocurrency. Due to early stage regulations for cryptocurrency, malicious actors are able to exploit immature or non-standardised defenses across the industry. Specifically, BlockFi faced a number of risks from targeted fraud or ransomware.

Using the MLTK, BlockFi investigated potential account takeover activity by identifying anomalous user agent strings or anomalies in account password resets. Furthermore, through the use of graph-based analytics, BlockFi also visualized fraud, botnet, and malware behavior.

Read about BlockFi's use of the MLTK to combat fraud [here](#).

Further information

Gemini Data is a solution provider that works with a wide range of customers across the globe. Tackling common industry problems, Gemini Data used Splunk to develop a solution for monitoring chargeback risk. Read about Gemini Data's use of Splunk and the MLTK [here](#).

Splunk has also produced a number of Splunkbase apps for detecting fraudulent activity. The [Splunk App for Fraud Analytics](#) (SFA) is a comprehensive fraud detection solution built on the existing development frameworks of Splunk Enterprise Security. SFA offers your fraud team a standardized workflow, extensive interactive visual investigation capabilities, and a robust risk-based alerting framework, which is completely customizable and extensible. Additionally the [Splunk Security Essentials for Fraud Detection](#) highlights a number of ways that Splunk can be used to detect different types of fraud.



New York-Presbyterian is taking a leading role in protecting the public by implementing highly effective controls to avoid the illegitimate use of controlled substances. Ultimately, we hope that other hospitals benefit from this new platform as well.

Jennings Aske, Senior Vice President and Chief Information Security Officer

8. Predicting data downtime in Splunk

Business challenge

Maintaining operational resilience becomes challenging when data feeds are interrupted, subsequently reducing visibility into how services are performing. Data source owners do not generally grant full access to Splunk administrators, so identifying interrupted data feeds into Splunk is problematic. Furthermore, understanding how event feeds operate can be difficult when many thousands of hosts and many hundreds of data source types are being monitored, with each host and source type combination potentially having different behavior patterns.

Splunk's approach

Splunk can collect data from most systems using forwarders, database connectors or by using data manager, for example. Once data has been collected for a period of time, ML models can be created to describe the expected number of events for a given host and source combination. The models continuously monitor the data feeds and detect anomalies when data streams start to deviate from the expected throughput.

Value

Increase detection efficiency: By finding and addressing abnormalities during data ingestion (before an ingest pipeline gets disrupted or broken), data uptime can be maintained in Splunk. As Splunk is used for cyber security monitoring or the monitoring of critical systems, maintaining data uptime will provide continued visibility of potential threats or service degradation, helping to maintain operational resilience.

Case studies and further information

Saudi Aramco


[Saudi Aramco](#) monitors their data feeds coming into Splunk to improve the effectiveness of their security monitoring. Using anomaly detection techniques, Saudi Aramco produced reports on input data feeds to determine if logs were being ingested as expected.

Read more about this use case and others from Saudi Aramco [here](#).

Further information

Several ways exist where anomaly detection can be used to monitor data input feeds into Splunk, but for a prescriptive approach on how to set up these monitoring capabilities, we have an in-depth webinar that walks through implementing data feed anomaly detection [here](#).

Furthermore, we also have a step-by-step deep dive on implementation, including sample xml for creating an anomaly detection dashboard [here](#). This deep dive should take a few hours to implement from end-to-end.



9. Demystifying security searches with the Splunk AI Assistant

Business challenge

Security teams will often operate with hundreds of detection rules, making it challenging to ensure awareness of the purpose and logic behind each individual rule. This is compounded by patchy or inconsistent documentation for the rules, with teams often relying on tribal knowledge or a few experts to maintain the effective operation and triage of their detections.

Splunk's approach

Splunk publishes comprehensive [documentation](#) on over 1,450 use cases that have been created by the Splunk Threat Research Team. We also provide [documentation](#) on all of the Splunk search commands, which can be used to understand what detection rules are doing.

With the release of the [Splunk AI Assistant](#), customers can now use ML — or more specifically a large language model (LLM) — to explain in plain English what a particular Splunk search is doing. Using this app users can simply copy and paste an SPL search and have the app convert this search to plain English. The Splunk AI Assistant also empowers users to search their data using plain English. Now you can write a prompt of what you want in plain English, and Splunk AI Assistant translates the request into query ideas that you can execute or build on, all within a familiar Splunk interface.

Under the hood, the Splunk AI Assistant uses generative AI to lower the barrier to using SPL. It provides a chat experience which is intuitive and simple. For new users, it reduces the learning curve for using SPL. Advanced users will find the assistant helpful in unlocking the power of SPL and understanding complex SPL queries.

Value

Reduce manual processing: Using the Splunk AI Assistant security teams can more easily document and describe their security detection rules. This can support faster onboarding of new team members and improve knowledge transfer across the security operations team.

Case studies and further information

Further information

More information can be found about the Splunk AI Assistant in our [documentation](#) and on [Splunkbase](#).

Get started today

- Download [MLTK](#) and [DSDL](#) to get started with ML.
- Explore the [Splunk AI Assistant](#) to access the power of Generative AI in your workflows.
- Check out our MLTK [deep dives](#) for detailed implementation guides for some popular uses of ML in Splunk, including a video [overview](#) and detailed [walkthrough](#).
- Explore solutions like [Enterprise Security](#), [Enterprise Security Content Updates](#), or [User Behavior Analytics](#) for out-of-the-box ML use cases.
- Review [Splunk Blogs](#) to find out more about given techniques and approaches to ML in Splunk (mapping to ML techniques and some blogs in the table below).
- The Splunk account team also can help explore support available from Splunk while exploring ML use cases.

Explore more resources

Anomaly Detection

[Cyclical Statistical Forecasts and Anomalies part 1](#)
[Cyclical Statistical Forecasts and Anomalies part 4](#)
[Cyclical Statistical Forecasts and Anomalies part 5](#)
[Cyclical Statistical Forecasts and Anomalies part 6](#)
[A Splunk Approach to Baselines, Statistics and Likelihoods on Big Data](#)
[Anomalies Are Like a Gallon of Neapolitan Ice Cream part 1](#)
[Anomalies Are Like a Gallon of Neapolitan Ice Cream part 2](#)
[Understanding and Baselining Network Behaviour Using Machine Learning part 2](#)

Predictive Analytics

[Cyclical Statistical Forecasts and Anomalies part 2](#)
[Cyclical Statistical Forecasts and Anomalies part 3](#)
[Anomalies Are Like a Gallon of Neapolitan Ice Cream part 1](#)
[Anomalies Are Like a Gallon of Neapolitan Ice Cream part 2](#)
[ITSI and Sophisticated Machine Learning](#)
[Predicting Resource Exhaustion with Double Exponential Smoothing](#)

Clustering

[Anomalies Are Like a Gallon of Neapolitan Ice Cream part 2](#)
[Visualizing a Space of JA3 Signatures with Splunk](#)

Graphs

[Understanding and Baselining Network Behaviour Using Machine Learning part 1](#)
[Chasing a Hidden Gem: Graph Analytics with Splunk's Machine Learning Toolkit](#)



Splunk, Splunk> and Turn Data Into Doing are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. © 2024 Splunk Inc. All rights reserved.