



With the continuous onslaught of successful cyber security attacks, many organizations are finding the need to load more and more information into their centralized log management solution. When organizations use a classic SIEM or Splunk for log management, they often find the cost of scaling the environment extraordinary. This data sheet describes how a license-free open source solution can boost your SIEM's performance while enabling greater economic scalability for multiple log management use-cases.

Your SIEM is Doing Too Much!

Originally, Security Information and Event Management (SIEM) systems were sold as a means to reduce the overload of false alerts from from Intrusion Detection Systems (IDS). However, since SIEMs need to look at many different log sources to accomplish this first use-case, many products expanded to include additional uses for the same data, becoming a complete centralized log management solution. Today, cyber security organizations need several uses from their log management system:

- Reduced Alerting (via Real-time Correlation)
- Compliance Reporting
- Advanced Security Analytics
- Forensics Investigations and Threat Hunting
- Threat Intelligence Management

Alerting and real-time correlation may only need a day or two of data retention. However, the newer use-cases need more, in particular forensics investigations can require years or decades of data retention. Forcing a classic SIEM to store this information – license costs aside – can slow down the system for its primary use-case.

Wrong Architecture

Most SIEMs are very well-suited for real-time correlation and alerting. However, these solutions are simply the wrong architecture to accomplish all the different use cases required of a full-featured centralized log management solution.

For data storage, classic SIEMs generally use commercial Relational Database Management Systems (RDBMS) that were originally designed to function on a single computer. Scaling these RDBMS incurs a large license fee – often in addition to the license fee of the SIEM itself – not to mention they were not originally designed for clusters of computers.

Splunk stores its data in a single, fully indexed repository without adding any managed structure to the data (no rows or columns). While this is great for forensics investigations, it was not designed from the beginning for compliance reporting or advanced analytics. Even more of a stretch is using Splunk for real-time correlation. While the data can move very quickly from point of origin to the results screen in front of an end-user, setting up hundreds of automated correlation rules means executing hundreds of additional queries every few seconds (or minutes or hours, depending on how “real-time” alerts are needed). The result is a drain on system resources so severe many customers have to reduce the frequency to four hours or disable the functionality altogether.

Security Data Lake to the Rescue

If you have invested significant resources into your SIEM, there is no reason to throw it away; keep it for its original primary use. However, instead of investing premium amounts to



leverage the same system for after-thought use-cases, it will be more economical to place a strategic investment in a multi-purpose security data lake.

Solutions such as System Soft Technologies' (SSTech) Security Intelligence and Analytics (SIA) solution are built from the ground up, with this variety of use-cases in mind. Some of the features and functionality of the solution include:

Real-Time and Batch - All data loading parsing routines accommodate both batch and real-time processing

Structured and Unstructured - To accommodate all scenarios, SSTech deploys the solution with all data feeds and enrichment processes going to both batch and real-time subsystems.

Spot and Metron - SSTech is committed to deploying solutions that will support analytics intended for either product, all in one environment.

Compliance Reporting - Included are a complete set of compliance reports ready to fit into any organization's larger compliance initiative for regulations such as PCI, HIPAA, FISMA, NIST, etc.

Threat Intelligence - Adding OpenTAXI Server to the solution means all data sources can easily be enriched with the latest threat intelligence information for known bad IP addresses, domain names, and URLs.

Insider Threat Detection - SSTech included a comprehensive implementation of user behavior analysis including advanced statistical and machine-learning techniques with a human-assisted feedback loop.

Incident Response and Forensics Investigation - An augmentation to SIEM capabilities inclusive of packet replay utilities, evidence store and hunting services commonly used by SOC analysts.

Application Framework - The solution is an application framework that enables a single view of diverse, streaming and batch processing security data at scale to aid security operations centers in rapidly detecting and responding to threats.

Pluggable Framework - An open pluggable framework enables easy creation of new parsers to integrate into batch and streaming feeds, enrichment processes, and the extensible analytics schema.

Standard BI Interface - All queries and reports are available in Zeppelin and Tableau, or use your own Enterprise standard!

How it Helps

Expanding to a security data lake instead of investing more in your existing SIEM not only helps achieve scalability in a more economical fashion, but more importantly, by reducing the data sent to your existing SIEM it increases the performance of your existing SIEM with no investment at all.