Trusted Cyber Technologies

THALES
Building a future we can all trust

Industry Standard Protection for Your Log Data

A joint Thales and Splunk Solution

White Paper

Contents

Data security in the age for your log data

According to the 2021 Thales Data Threat Report, the COVID-19 pandemic forced many changes to enterprises that caused a ripple effect throughout the security community, which was evident in the over 2600 respondents surveyed from more than 10 countries across the globe. The shift to remote work and the subsequent accelerated use of cloud-based infrastructure have profound impacts on security teams. At the same time, breaches continue at an increasing rate. Sixty percent of organizations report they have been breached (41% in the past year alone) with threats now emerging from a wide variety of external as well as internal sources.

In today's information-oriented economy, the crown jewels of an organization are its data. Whatever you call it, "log data" is one of the most underused and undervalued assets of any organization. Log data comprises of activity logs associated with changes to system configurations, data from APIs, message queues, change events, output of diagnostic commands, call detail records, and more. It is valuable because it contains a definitive record of all the activity and behavior of your customers, users, transactions, applications, servers, networks, and mobile devices. Some of the most important insights enterprises can gain— where things went wrong, how to optimize the customer experience, the fingerprints of fraud—are hidden in the log data that is generated by the normal operations of your organization.

The challenge with leveraging log data is it comes in an array of unpredictable formats, and traditional monitoring and analysis tools were not designed for the variety, velocity, volume or variability of this data. This is where a security information and event management (SIEM) vendor like Splunk comes in. Splunk's industry-leading collect, categorize and correlate event data coming from various devices, systems and applications throughout the enterprise. It is an enterprise ready, fully integrated solution for log management, data collection, storage, and visualization.

In this paper, we will examine how CipherTrust Transparent Encryption from Thales secures the Splunk log repositories and databases. Then, we will show how Live Data Transformation, an extension of CipherTrust Transparent Encryption, can encrypt Splunk buckets seamlessly with zero-downtime. We will close with how the CipherTrust Security Intelligence solution for Splunk extends the reach of security information and event management (SIEM) capabilities to detect and counter attacks on sensitive data.

Centralized log analysis and visualization with Splunk

Enterprises recognize that event logs play an important role in their organization. Logs are used for various purposes such as IT operations, system and application monitoring, business analytics, security and compliance, and much more. Having a centralized logging system makes life easier for developers and administrators, especially when there is a need to correlate abnormal events in system activity and applications, detect security threats, secure the application due to unexpected hits on services, or review the performance of the application. Some of the great features of a centralized log aggregation and analysis system are its low-cost maintenance, easy log correlation and searching, meaningful alerting, and enhanced dashboard.

Splunk is a centralized log analysis tool for machine-generated log data, which is unstructured/structured and complex multi-line data-that provides features, such as Easy Search/Navigate, Real-Time Visibility, Historical Analytics, Reports, Alerts, Dashboards and Visualization.

With Splunk, enterprises can:

- Find and detect security issues faster and investigate security alerts in minutes, not hours or days
- Monitor end-to-end infrastructure to avoid service degradation or outages
- Gain operational intelligence with real-time end-to-end visibility and critical insights into customer experience, transactions and other key business metrics

Analyzing risks of using analytics-driven security platforms

What enterprise IT managers require today is a simple way to correlate all security relevant data, so they can manage the organization's security posture. Instead of merely watching events after they occur, an IT organization should anticipate their occurrence and implement measures to limit the enterprise's vulnerability in real time. For that, organizations need an analytics-driven SIEM platform.

In addition to typical data-oriented concerns about security and privacy, the collection of log data can introduce additional risks, including:

- **Insider Threats.** Organizations typically establish strong controls around the data center and have protected many of the sources from which Splunk pulls data. It is important, however that the Splunk data warehouse itself be strongly protected, as the information is arguably more valuable in this state.
- **Role-Based Access Control.** It is essential that the valuable information contained in a Splunk database and its logs be accessible only to the authorized users. In addition, Splunk administrators can themselves provide a new threat vector. These users have access to a wealth of sensitive information and may expose it inadvertently. It is therefore critical that access to Splunk data be trackable.
- **Data Volume.** Splunk Enterprise provides a vital means to correlate huge volumes events coming from various sources at the device, network, system/OS and application levels, and convert them to a limited set of actionable alerts and security incidents.
- **Data Integrity.** Splunk Enterprise also provides message integrity measures that show whether an event has been inserted or deleted from the original stream.

The Splunk and CipherTrust Solution

With advanced persistent threats (APTs) occuring more commonly, hackers are actively seeking to steal credit card data, personally identifiable information (PII), critical intellectual property (IP), and other sensitive data to sell to the highest bidder. Some of the most effective tools for fighting these attacks are the security intelligence and threat detection capabilities of SIEM solutions, such as Splunk. Splunk monitors both real-time events and tracks long-term data to find anomalous usage patterns, qualify possible threats to reduce false positives, and alert organizations when risks are detected. The CipherTrust Data Security Platform from Thales is the leader in advanced data security solutions and services delivering trust wherever information is created, shared, or stored. Security solutions ensure that critical data is both discovered and protected in any deployment–on-premises, in the cloud, in data centers, or in big data environments–without sacrificing business agility. CipherTrust Transparent Encryption, an integral part of the CipherTrust Data Security Platform, encrypts all data Splunk writes to disk (index, raw data, metadata, and everything else).

The CipherTrust Transparent Encryption software agent that enables data-at-rest encryption, privileged user access control and the collection of security intelligence logs without change to applications, databases or infrastructure. In joint solution provided by Splunk and Thales, the CipherTrust Transparent Encryption agents run at the file system level or volume level on a server that has the Splunk software installed. This data encryption solution stops root, system, cloud, and other administrators from accessing Splunk data in any form or location, while preserving their ability to do their jobs. Since 2001, CipherTrust Transparent Encryption (previously known as Vormetric Transparent Encryption) has been used by the world's largest companies and governments, protecting hundreds of thousands of servers and untold amounts of data. Now, more than ever, the combination of aggregated data compiled by Splunk and easy-to-deploy encryption from Thales is compelling.

Essential elements within the Data Security Platform include:

- **CipherTrust Transparent Encryption** – File system agents provide protection of the Splunk data warehouse, as well as files, folders, documents, logs and more.
  - **Live Data Transformation** – Zero-downtime encryption deployments allows for encrypting and re-keying data without taking applications offline. This allows deployment of data security controls to applications while maintaining business continuity and high availability.
- **CipherTrust Manager** – The brains of the CipherTrust Platform provides centralized administration of encryption keys and data security policies. CipherTrust Manager is available as Common Criteria and FIPS 140-2 certified virtual (Level1) and physical appliances (Level 2 and 3)
- **CipherTrust Security Intelligence Application for Splunk** – This application produces detailed security event logs that are themselves integrated into Splunk.

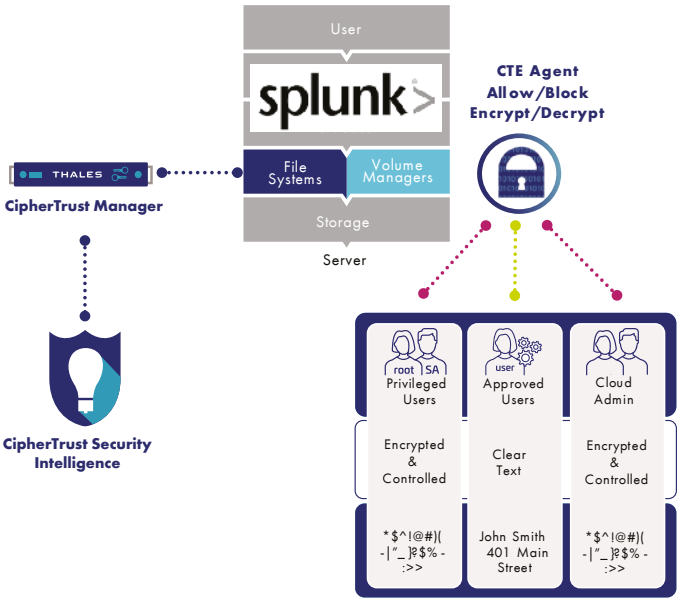

CipherTrust Transparent Encryption for Splunk Repositories

CipherTrust Transparent Encryption delivers data-at-rest encryption with centralized key management, privileged user access control, and detailed data access audit logging that helps organizations meet compliance and best practice requirements for protecting data, wherever it resides. The CipherTrust Transparent Encryption agent resides at the operating file-system or device layer on a server that has Splunk software installed, and encryption and decryption is transparent to all applications that run above it. It provides rich access controls, which allow organizations to determine who can access data, when they can access it, and what type of access they have.

Key advantages of CipherTrust Transparent Encryption:

- **Transparent data protection.** Continuously enforces file-level encryption that protects against unauthorized access by users and processes and creates detailed data access audit logs of all activities without requiring changes to applications, infrastructure, systems management tasks, or business practices.
- **Seamless and easy to deploy.** CipherTrust Transparent Encryption agents are deployed on servers at the file system or volume-level and support both local disks as well as cloud storage environments, such as Amazon S3 and Azure Files.
- **Define granular access controls.** Apply granular, least-privileged user access policies that protect data from external attacks and misuse by privileged users. Access control policies can target local server users and groups, as well as integrate with LDAP/Active Directory.
  Controls also include access by process, file type, time of day, and other options.
- **High-performance hardware accelerated encryption.** CipherTrust Transparent Encryption only employs strong, standard-based encryption protocols, such as the Advanced Encryption Standard (AES) for data encryption and elliptic curve cryptography (ECC) for key exchange. Encryption overhead is minimized using the AES hardware encryption capabilities available in modern CPUs.
- **Broadest system and environment support.** The agent is available for a broad selection of Windows, Linux, and AIX platforms, and can be used in Physical, virtual, cloud, container, and big data environments, regardless of the underlying storage technology.

CipherTrust Encryption agents are distributed and optimized for specific file system and encryption acceleration hardware across servers, resulting in very low latency and overhead. Agents employ logic and fine-grained policies defined by the CipherTrust Manager to evaluate attempts to access protected data, and then grant or deny access; all activities taking place around the protected data are logged. The agents have been deployed in tens of thousands of servers, making them the right solution for Splunk Enterprise Big Data requirements.

How it works?

- A CipherTrust Transparent Encryption agent runs on each Splunk indexer. This provides an overlay security file system. The example below shows a Linux directory, but the process works exactly the same for Windows servers.

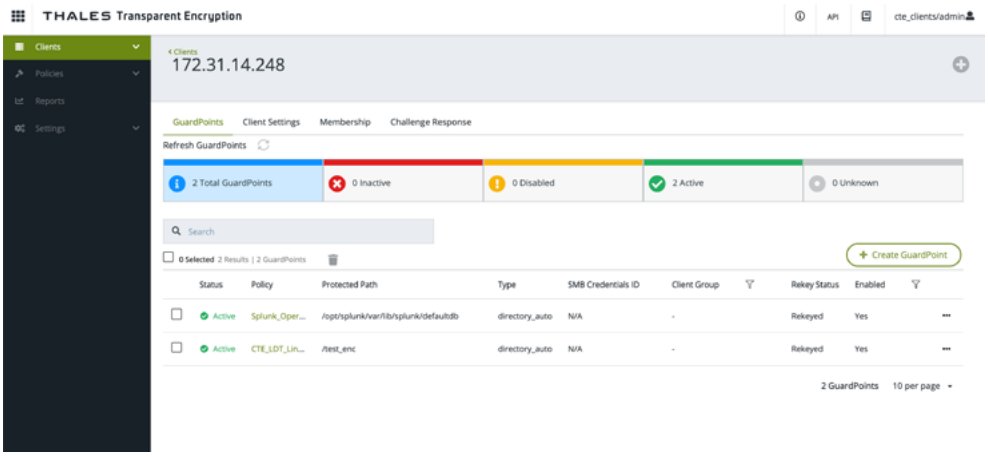

Multiple policies and guardpoints can be used on the same server, greatly enhancing the granularity of access control rules.

- The CipherTrust Transparent Encryption agent policy allows the Splunk binary (splunkd) and other processes permission to the guard point.

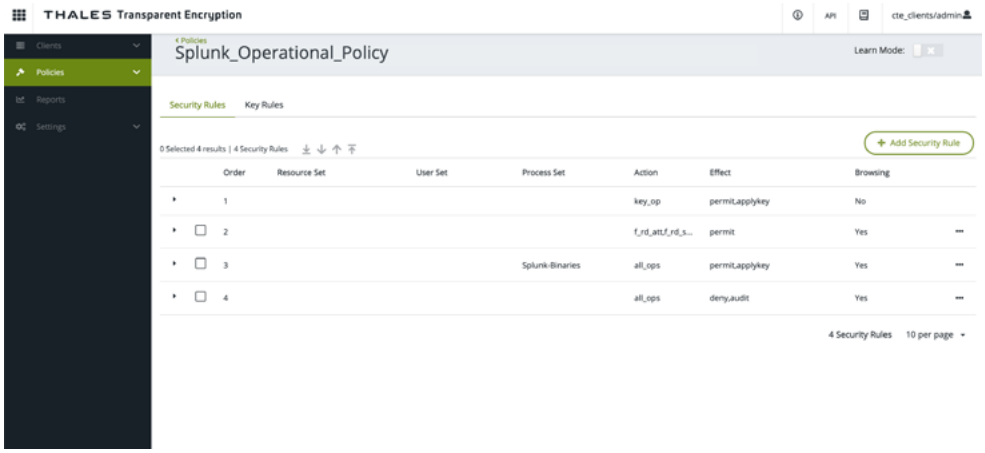

Access control policies follow the same logic as a firewall; rules are evaluated from top to bottom, as soon as a rule matches, the policy stops evaluation and triggers the "effect." This example would flow as follows:

1. The top line ensures that the agent is able to perform initial encryption operations and subsequent key rotations
2. This is known as the "metadata rule," This allows all users to see things like file size, type, name, etc. but does NOT permit actual reading of the file.
3. By including a Process Set, called "Splunk-Binaries" here, users can define individual executables which should have access to the encrypted data
4. The last line in most production policies will be a catch-all for any access that doesn't meet the earlier requirements. The DENY effect prevents encrypted data from being accessed, and the AUDIT effect logs the attempt for review.

- When data is written to the hot directory (hot dir), it is encrypted on the way in. When the splunkd process reads the data back into memory it's decrypted on the way out.

CipherTrust Transparent Encryption Live Data Transformation Advantage

Deployment and management of data-at-rest encryption can present challenges when transforming clear-text to cipher-text, or when rekeying data that has already been encrypted. Traditionally, these efforts required planned downtime, or they required labor-intensive data cloning and synchronization efforts. CipherTrust Transparent Encryption Live Data Transformation, with zero-downtime encryption deployments, eliminates these hurdles, enabling encrypt and rekey with unprecedented uptime and efficiency.

- **Improve Security and Data Availability.** CipherTrust Transparent Encryption Live Data Transformation, with zero-downtime encryption deployments, allows for encrypting and re-keying data without taking applications offline. This enables deployment of data security controls to applications while maintaining business continuity and high availability.
- **Reduce the Operational Costs of Encryption.** In the past, critical applications had to be taken offline for initial encryption of data and encryption maintenance, with substantial operational costs. Zero-downtime encryption eliminates this need.
- **Versioned Backups and Archives.** With key version management, Live Data Transformation offers efficient backup and archive recovery that enables more immediate access. In a data recovery operation, archived encryption keys recovered from the CipherTrust Manager are automatically applied to an older data set. Restored data is encrypted with the current cryptographic keys.

How it works?

- With CipherTrust Live Data Transformation, you can apply the CTE Guardpoint on cleartext buckets with only a momentary downtime of the indexer to mount the guardpoint. At that point, the data being written is immediately encrypted with the encryption key defined in the policy. The existing data will be encrypted in the background based on the QOS (Quality of Service) schedule that is set.

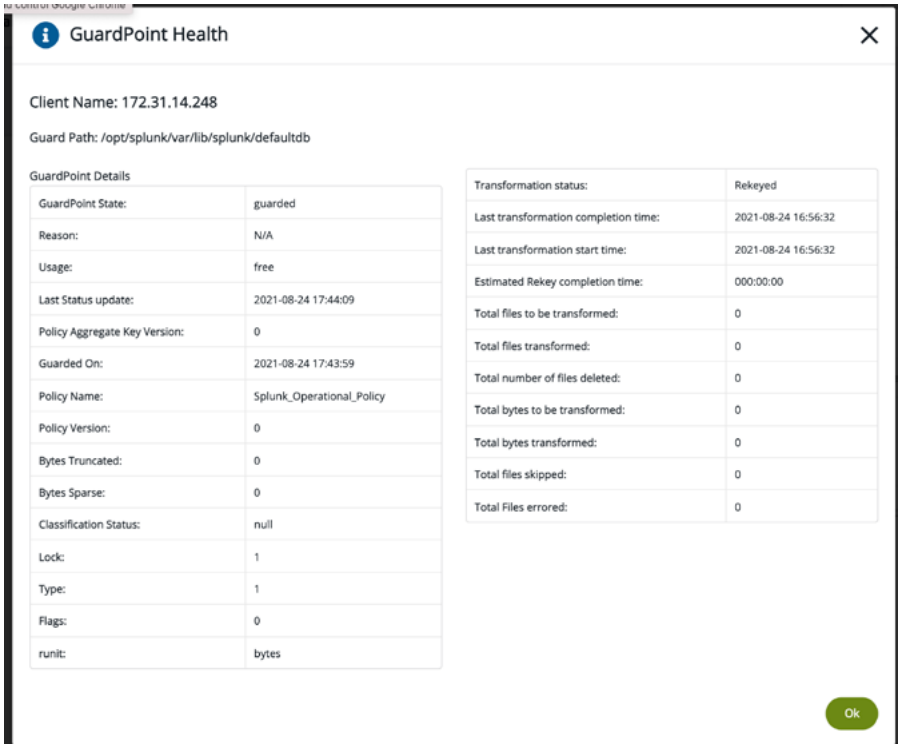

Live Data Transformation Status Window shows useful information for Security Admins, including what security policy is in place on the guard point, when the data was last re-keyed, and whether or not there is any data currently in use within the guard point.

Going the extra mile with CipherTrust Security Intelligence

Detailed data access audit logs delivered by CipherTrust Transparent Encryption are useful to identify unauthorized access attempts, as well as to build baselines of authorized user access patterns. CipherTrust Security Intelligence captures granular logs of all file access attempts and completes the picture with pre-built integration to the Splunk platform making this information actionable. The solution allows immediate automated escalation and response to unauthorized access attempts and captures all the data needed to build behavioral patterns required for identification of suspicious usage by authorized users.

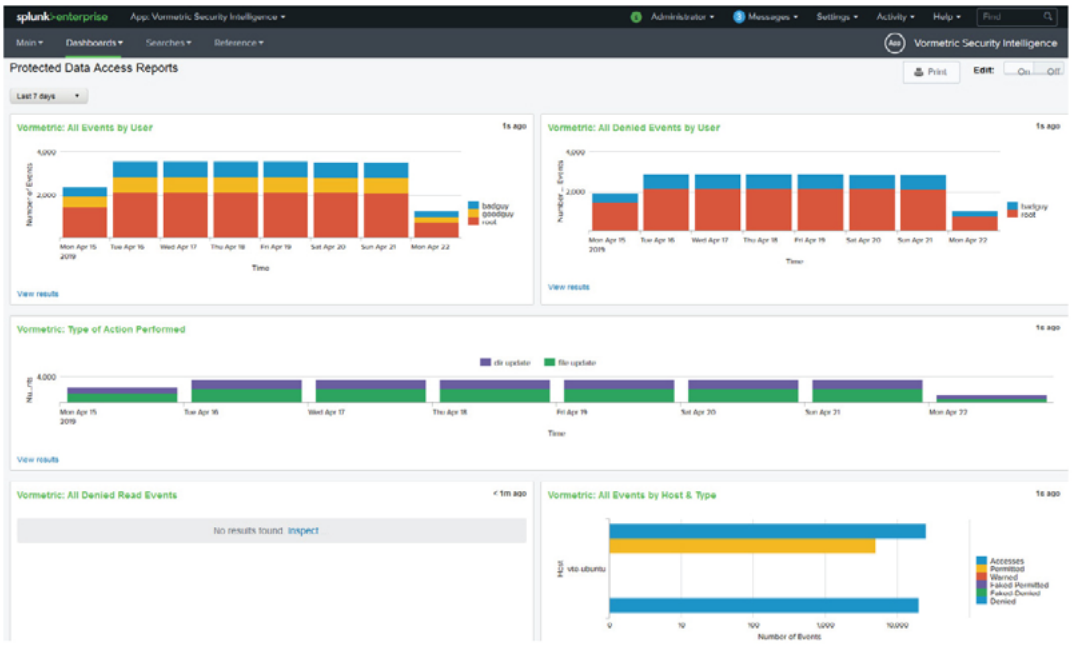

Conclusion

The number of data security threats to enterprises is growing quickly, so collecting and coordinating actionable intelligence from a variety of IT systems is essential to safeguard the reputations and business models of enterprises around the world.

But now, even the intelligence to identify and defend against these attacks is under siege. Such information is itself a prime target for insider attacks, and the number of company insiders with access to the data is growing. This makes the case for combining security intelligence data aggregated by Splunk with easy-to-deploy encryption from Thales more compelling than ever. This combined Thales/ Splunk solution identifies internal and external threats, and keeps sensitive data secure and fully compliant with regulatory statutes.

About Thales

The people you rely on to protect your privacy rely on Thales to protect their data. When it comes to data security, organizations are faced with an increasing number of decisive moments. Whether the moment is building an encryption strategy, moving to the cloud, or meeting compliance mandates, you can rely on Thales to secure your digital transformation.

Decisive technology for decisive moments.

THALES

Building a future we can all trust

Contact us

info@thalestct.com

> thalestct.com<