splunk>

TECH BRIEF

Event Correlation

Building Relationships Across Machine Data

Index Any Event. Ask Any Question.

Your IT systems and technology infrastructure generate massive amounts of machine data in varying formats inside and outside of your organization. This machine data includes data streaming from packaged and custom applications, app servers, web servers, databases, wire data from networks, virtual machines, mobile devices, telecoms equipment, operating systems, sensors, mainframes and much more. Correlating the events generated by all of these systems and understanding which are most relevant is a challenge.

Splunk Enterprise is the industry-leading platform for collecting and analyzing this machine data. It provides a unique approach for correlating complex events to deliver operational visibility across all of your IT and remote systems.

Finding the answers to questions like, "How far apart in time did a specific set of events occur?" or "What's the total amount of time it took for a transaction to take place?" requires complex correlation of multiple data sources beyond the capability of most traditional data collection systems. What's more, machine-generated data formats vary widely from structured syslog to SNMP to unstructured multi-line application data. To discover the relationships in data, most data collection systems allow you to perform simple text searches and apply Boolean operators such as "AND," "OR," "less than," "greater than" or "equal to" to search the data. While this action may create a filtered view, the user still has to export the search results into another tool, reformat the data and perform other manual work to find correlations.

Correlating data from large and widely varied data sources to discover meaningful information and relationships requires much more than Boolean operators.

Splunk Enterprise scales to ingest hundreds of terabytes of data per day, features a command language with over 100 analytical commands and can support structured and unstructured data in any environment—physical, virtual and cloud. These analytical commands support correlation in ways not possible in other data collection systems, for example, searching for patterns of activity across multiple sources of machine-generated data. Users can easily automate the results of correlations to generate alerts or support business metrics, leading to better business decisions and operational intelligence (see *Figure 1*).

Splunk Supports Five Correlation Types

- **Time and geolocation based** – Identify relationships based on time proximity or geographic location
- **Transaction based** – Track a series of related events together and display a single event and produce a "duration" and/or "event count"
- **Sub-searches** – Take the results of one search and use them in another
- **Lookups** – Correlate data to external sources
- **Joins** – Support for "SQL-like" inner and outer joins

Time and geolocation based – allows you to see all or any subset of events that take place over a given time period and also pinpoint their location. This basic correlation allows you to view events that have taken place in time periods ranging from the previous second to the last year, and is key to any security or operations investigation. Time-based and geo-spatial correlation capabilities are available to the user directly from the Splunk Enterprise UI.

Transaction based – track a series of related events as a single transaction. These events can come from any number of separate IT systems and data sources. For example, a key metric for credit card clearing organizations is the time it takes for a credit card purchase transaction to be authorized. The credit card transaction time represents how long it takes a transaction to work through the IT infrastructure and the anxiety a credit card holder feels while waiting for the transaction to be authorized. Metrics around the speed of business processes are also important for capacity planning. Reviewing these metrics can help the business understand where additional resources may be needed to improve performance.

Sub-searches – take the results of one search and use them in another to create if/then conditions. Using a sub-search allows users to only see the results of a search if certain conditions

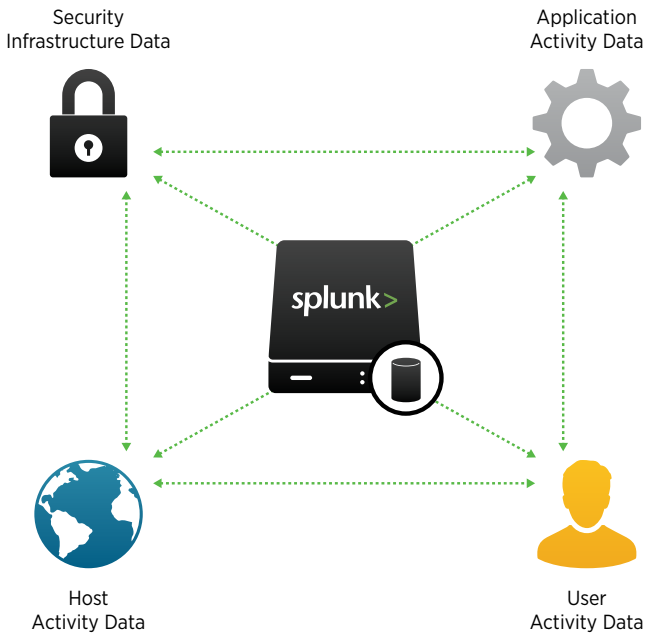

Figure 1: Splunk Enterprise correlates data from a wide variety of sources.

are met. Security information and event management systems operate on this premise. For example, a user may be only interested in viewing one event if the threshold for another event is met in a given time period.

Lookups – can be used to enhance, enrich, validate or add context to data collected in Splunk Enterprise. For example, correlating intrusion detection data (IDS) with data from an asset management system can reduce IDS false positives. An attack based on a Windows OS vulnerability seen by an IDS can be correlated with asset management data that indicates the host being attacked runs the AIX OS.

Joins – are similar in concept to "joins' in a SQL database. Inner and outer joins are supported. "Join" as part of a search string can link one data set to another based on one or more common fields. Two completely different data sets can be linked together based on a user name or event ID field presenting the results in a single view.

Summary

The analytical commands in Splunk Enterprise can be used in combinations that create operational metrics and business insights from machine-generated data. Results from any correlation can be displayed visually in a report or dashboard to support better decision-making. Splunk correlation commands can work together in the same search command to provide functionality similar to sophisticated event management or correlation systems.

Free Download

Splunk Enterprise. Download Splunk Enterprise for free. You'll get a Splunk Enterprise 6.2 license for 60 days and you can index up to 500 megabytes of data per day. After 60 days, or anytime before then, you can convert to a perpetual free license or purchase an enterprise license by contacting sales@splunk.com.

Splunk Cloud. Sign up for Splunk Cloud, which delivers Splunk Enterprise as a service. Currently available in the U.S.A. and Canada.

Features	Splunk Free	Splunk Enterprise	Splunk Cloud
Indexing Volume	500MB/day	Unlimited According to license	5GB/day to TB/day According to license
Data Onboarding	•	•	•
Universal Indexing	•	•	•
Search	•	•	•
Distributed Search		•	
Monitoring and Alerting		•	•
Reporting	•	•	•
Knowledge Mapping	•	•	•
Dashboards	•	•	•
Data Model	•	•	•
Pivot	•	•	•
Event Pattern Detection	•	•	•
High Performance Analytics Store	•	•	•
Report Acceleration	•	•	•
Embedded Reports	•	•	•
PDF Delivery		•	•
Access Control & Single Sign-On		•	•
Single-Site Clustering		•	
Multisite Clustering		•	
Distributed Management Console		•	
Universal Forwarder	•	•	•
Forwarder Management	•	•	•
Rich Developer Environment	•	•	•
Apps	•	•	•
Premium Apps		•	•
Standard Support	•		
Enterprise Support		•	•

*Splunk Cloud is currently available in the U.S.A. and Canada

250 Brannan St., San Francisco, CA 94107 info@splunk.com | sales@splunk.com 866-438-7758 | 415-848-8400 apps.splunk.com

splunk> listen to your data™

www.splunk.com

Item # TB-Splunk-EventCorrel-108

Splunk Product Features & Descriptions

Features	Definitions
Indexing Volume	Maximum indexing volume per day
Data Onboarding	Wizard-based workflow to simplify onboarding of any data source
Universal Indexing	Universal real-time indexing of machine data
Search	Ad hoc search across real-time and historical data
Distributed Search	Search across multiple Splunk deployments; supports load balancing and failover
Monitoring and Alerting	Monitor and alert for individual and correlated real-time events
Reporting	Ad hoc reports across real-time and historical data
Knowledge Mapping	Knowledge mapped to machine data artifacts
Dashboards	Highly customizable and interactive dashboards integrating real-time machine data and charts, reports and tables
Data Model	Used to define consistent relationships in machine data
Pivot	Drag-and-drop UI to explore, manipulate and visualize machine data
Event Pattern Detection	Automatically discovers patterns and commonalities in your data with a single click
High Performance Analytics Store	High performance analytics technology
Report Acceleration	Transparent data summarization technology
Embedded Reports	Embed charts and reports in other third-party business applications external to Splunk Enterprise
PDF Delivery	Scheduled and automated PDF generation and delivery of reports and dashboards
Access Control & Single Sign-On	Integrated role-based access control and user authentication with LDAP directory and single sign-on integration
Single-Site Clustering	High availability architecture for machine data availability in a single site deployment
Multisite Clustering	High availability architecture for disaster recovery in a multisite deployment
Distributed Management Console	Centrally manage the health and performance of distributed Splunk deployments
Universal Forwarder	Forwarding of data securely and reliably from remote systems in real time
Forwarder Management	UI for monitoring and deploying forwarder configurations
Rich Developer Environment	Developer platform for building enterprise apps that leverage Splunk software with modern web languages
Apps	Access to hundreds of partner, community and Splunk Apps from apps.splunk.com
Premium Apps	Access to premium Splunk Apps
Standard Support	Access full product documentation, Splunk Apps, Splunk Answers and IRC channel
Enterprise Support	Direct access to Splunk customer support, ability to manage cases online, tailored support levels

250 Brannan St., San Francisco, CA 94107 info@splunk.com | sales@splunk.com 866-438-7758 | 415-848-8400 apps.splunk.com

Item # TB-Splunk-EventCorrel-108