Splunk Data Fabric Search

Fast analysis across your entire data ecosystem

- Analysis at an **unprecedented scale**
- **Seamless analysis** across all Splunk **deployments**
- Expand analysis to **third-party data stores** (PRE-RELEASE)

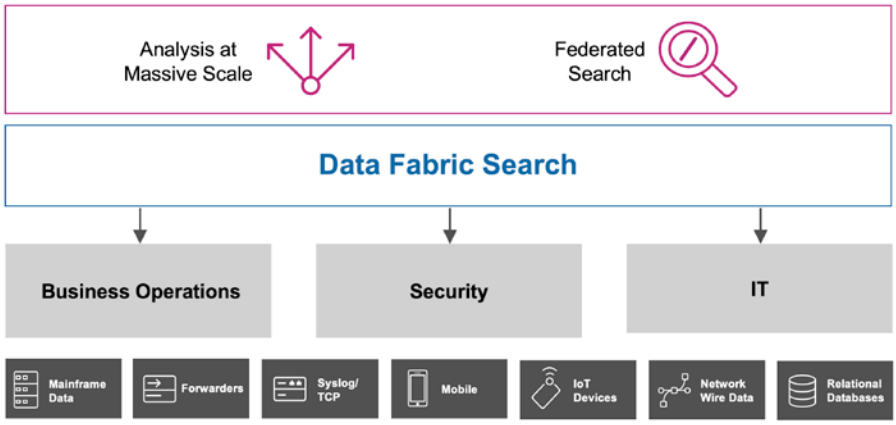

As the data and technology landscape continues to evolve to unprecedented levels of scale, with new devices, data types and sources, organizations are realizing the need for new and more efficient ways to store, analyze, and correlate diverse datasets.

Splunk Data Fabric Search (DFS) is a highly scalable and extensible solution that quickly weaves together insights from massive datasets across multiple diverse data stores. With DFS, organizations can quickly identify patterns in customer behavior across multiple products or services, analyze multiple hours to years of data, or correlate events across billions of individual user accounts or device IDs.

Unlike solutions that can't scale and aren't built to solve IT, security, IOT, and business analytics use cases, Splunk Data Fabric Search can analyze massive datasets, across any number of diverse data stores and is built on the Splunk Platform to solve security, IT, IOT and business analytics use cases.

Why Splunk Data Fabric Search?

For large organizations with multiple terabytes of data living across various systems whether it be across multiple Splunk deployments or across data stores outside of Splunk, Data Fabric Search connects massive datasets from across any data store and allows you to perform analysis on complex datasets containing several billions of unique values, including device IDs and user IDs, and datasets that span multiple hours or even years of data.

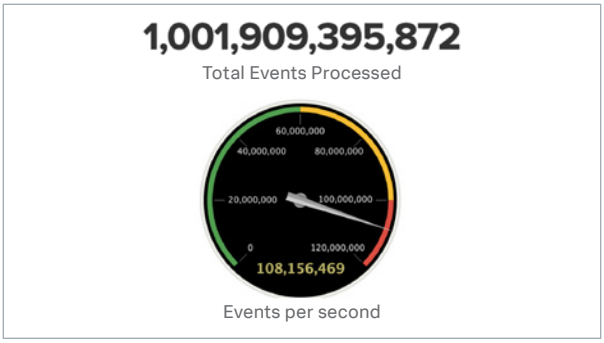

Analysis at an unprecedented scale

Analyze datasets in Splunk with long timeframes spanning hours to years of data and perform analysis on high-cardinality data containing up to billions of unique values such as user IDs or IP addresses.

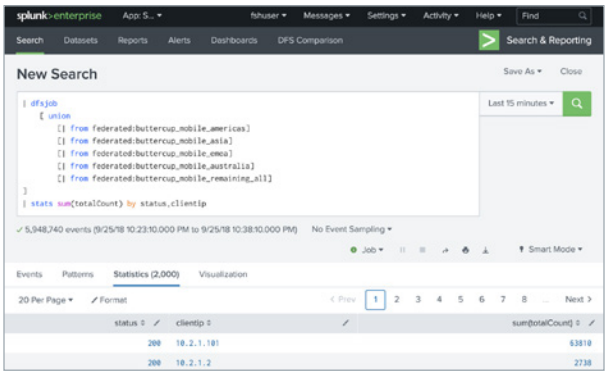

Seamless analysis across all Splunk deployments

Use a single query to analyze data across any number of Splunk deployments while maintaining user access permissions, regulating system resources and leveraging existing knowledge objects (e.g. saved searches, event tags etc.) across deployments.

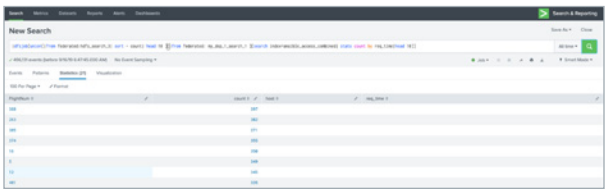

Expand analysis to third-party data stores

Use a single query to analyze data across Hadoop, Amazon S3 plus more to come (sign-up to learn more and participate in the pre-release program).

Technical Requirements:

Deployments with at least 1TB of data

Upgrade to Enterprise 8.0

Spark Requirements

- Dedicated Spark Cluster for Splunk
  - 4:1 ratio for Splunk indexers : spark nodes
  - Minimum 5 Spark nodes (with <=20 indexers)
- Minimum Spark Node Requirement
  - CPU: 8 core (16 recommended)
  - Memory: 64GB (128GB recommended)
  - Network: 10Gbps
  - Storage: 500GB+ (1200+ IOPS)

Recommended Model for Spark Configuration:

Download the Splunk DFS Manager app on Splunkbase. The app provides a user interface that allows customers to manage and monitor the Spark master and Spark workers by providing a high-level view of your DFS deployment and the available resources.

Think **Splunk Data Fabric Search** is a good fit for your organization? Learn more about Data Fabric Search Pricing. Contact our sales experts.

splunk>

Learn more: www.splunk.com/asksales

www.splunk.com

© 2019 Splunk Inc. All rights reserved. Splunk, Splunk>, Turn Data Into Doing, The Engine for Machine Data, Splunk Cloud, Splunk Light and SPL are trademarks and/or registered trademarks of Splunk Inc. in the United States and other countries.

2019-Splunk-SEC-Data Fabric Search-101-PB