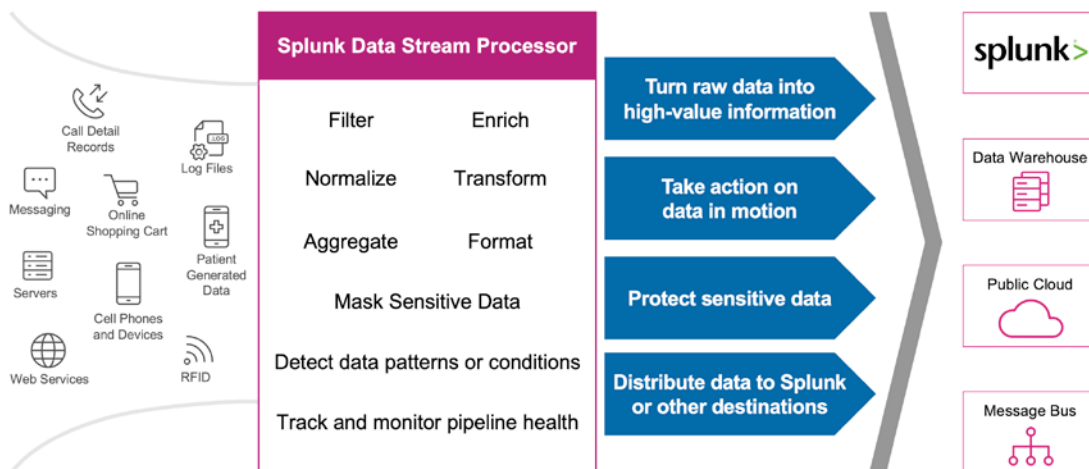


Splunk Data Stream Processor

A real-time stream processing solution that collects, processes and delivers data to Splunk and other destinations in milliseconds

- **Quickly turn raw data into high-value information** before data is sent to Splunk or another destination
- **Take action on data in motion** by quickly identifying abnormal data patterns or specific conditions that occur on the stream
- **Protect sensitive information** by masking sensitive data before it reaches downstream users
- **Distribute data across the enterprise** to Splunk or other third-party destinations



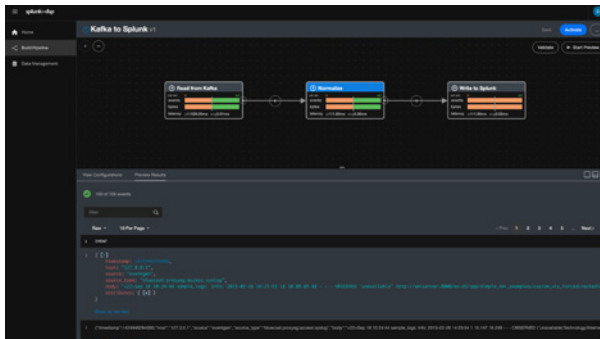
As the data and technology landscape continues to evolve to unprecedented levels of scale, new types of devices and more data sources, organizations need a robust solution that guarantees efficient data delivery across the enterprise.

Splunk Data Stream Processor (DSP) is a scalable stream processing solution built to guarantee delivery of high-volume, high-velocity data across the enterprise. As events occur, DSP continuously collects, formats, and organizes high-velocity, high-volume data based on specified conditions, masks sensitive or private information, detects abnormal data patterns, and then distributes results to Splunk or other destinations in milliseconds.

Unlike standalone stream processing solutions that aren't built on the Splunk platform or piecemeal solutions that lack full stream processing capabilities, DSP is a complete stream processing solution built on the Splunk platform to complement our entire portfolio of products that solve IT, security, IOT and business analytics use cases. Further, DSP is an extensible solution that can integrate with platforms beyond Splunk such as Apache Kafka®, Amazon S3, CloudTrail, Amazon Kinesis, Event Hubs plus more, to solve any use case.

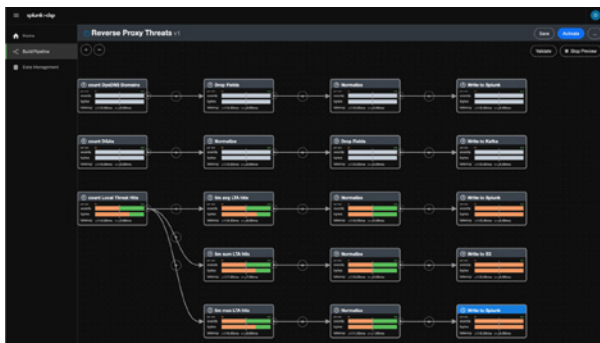
What is Stream Processing?

Stream processing is a technology that allows users to optimize data flow to get faster, more immediate data insights. Unlike batch processing which first collects large batches of data then processes that data, stream processing is designed to instantly process and analyze continuous streams of data in a short span of time, allowing organizations to immediately respond to critical events as soon as they occur.



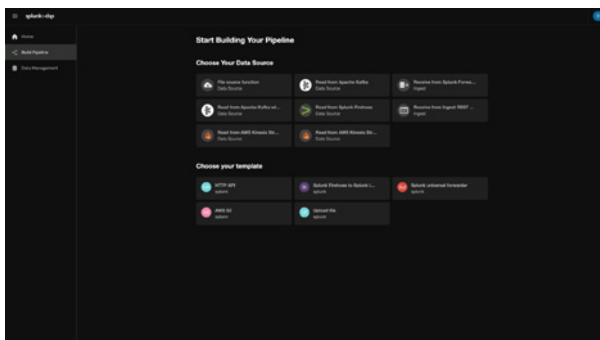
Transform or take action on data in motion

- Aggregate or filter out data that meets specific conditions
- Format, organize, or normalize data based on predefined conditions
- Mask sensitive or private information



Deliver data to multiple destinations across the enterprise

- Give users access to all data streaming into DSP or distribute subsets of data to various destinations
- Robust streaming solution built on top of Apache Kafka® that guarantees data delivery and scales as your organization experiences growth



Stream at enterprise scale with a graphical UI and embedded machine learning

- Use a drag and drop UI to design data pipelines
- Auto detect source types to easily group and evaluate data
- Leverage pre-built pipelines by creating templates for future use
- Setup automated rules to format, filter, and send subsets of data downstream based on predefined conditions
- Track the performance of a pipeline by monitoring data latency and throughput

Technical Requirements

- Minimum Node Requirement
 - CPU: 8 core (16 recommended)
 - Memory: 64GB (128GB recommended)
 - Network: 10Gbps
 - Storage: 1TB
- Minimum 5 Node Cluster

Supported Data Sources: Apache Kafka®, Amazon Kinesis, Amazon S3, CloudTrail, Azure Event Hubs, REST APIs, Splunk (Universal Forwarder, Heavy Weight Forwarder)

Supported Destinations: Apache Kafka®, Amazon Kinesis, Splunk

More sources and destinations to come in future releases

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



Learn more: www.splunk.com/asksales

www.splunk.com