


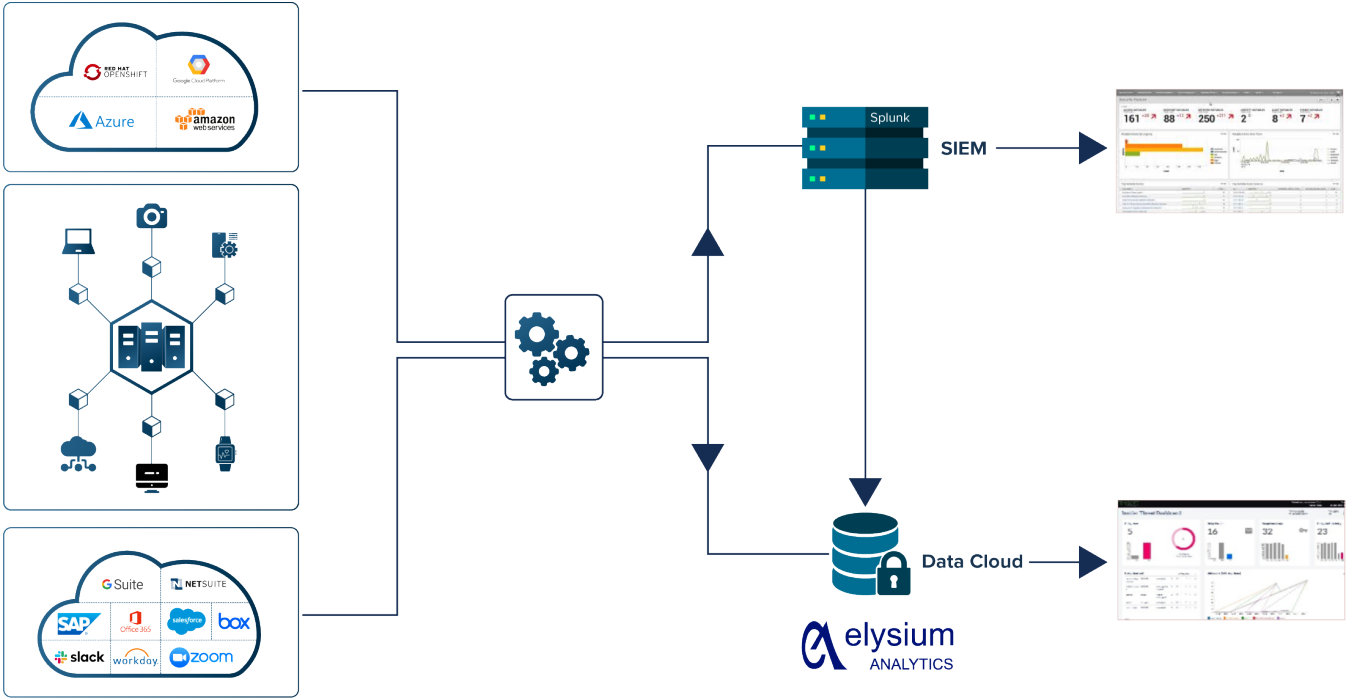
Migrate from Splunk to Snowflake with Elysium Analytics





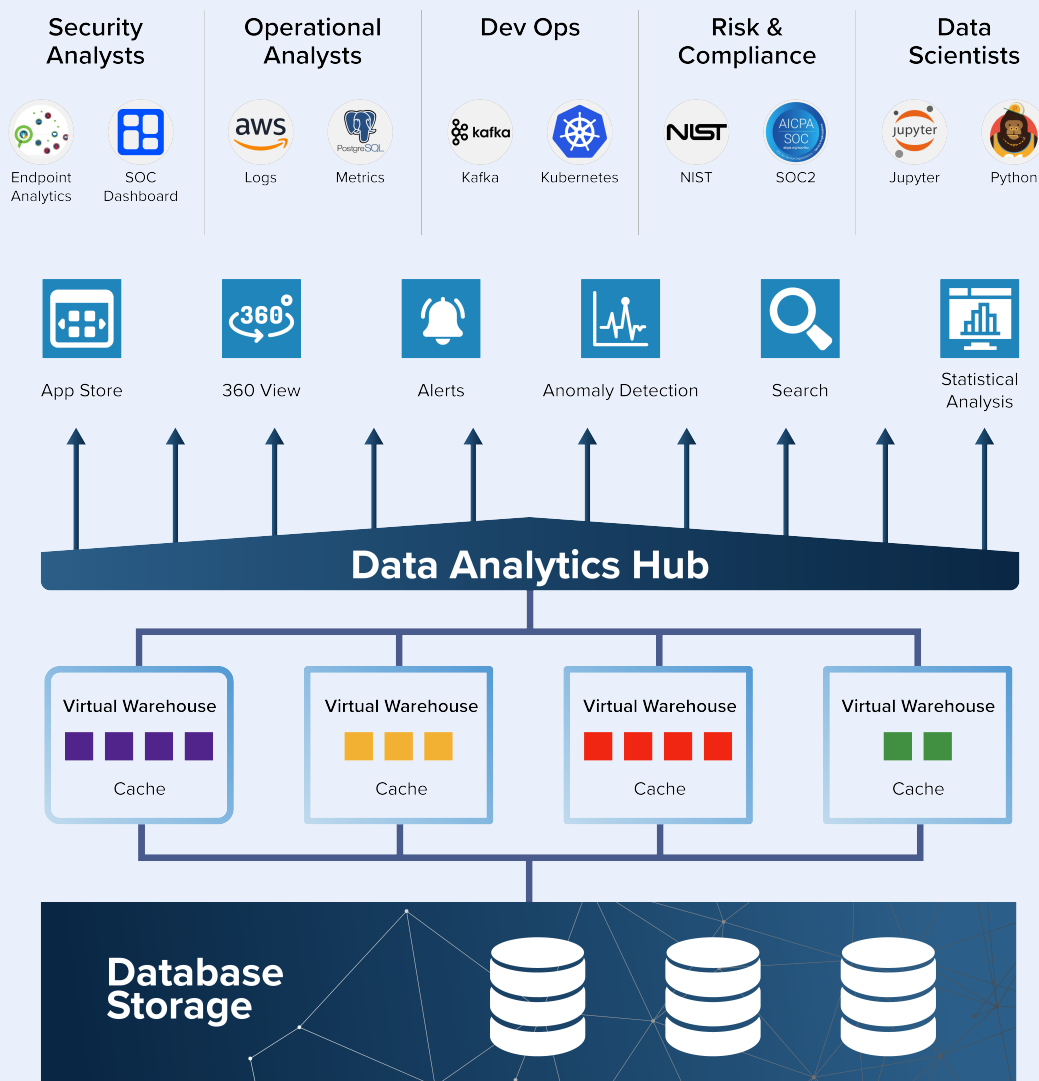
Splunk was launched in 2003 with the promise of giving organizations insights into IT data with its search software that indexes logs and generates alerts. Splunk rapidly became the leading search application for IT data as operators were looking for a quick way to make sense of logs, configuration files, message queues, etc. in order to diagnose problems better. Over the past two decades, the volume, variety, and velocity of data generated by systems and users have grown exponentially, making it challenging for many organizations to ingest all the relevant data due to Splunk's pricing based on the volume of data ingested per day and their tiered node-based architecture. With the demands quickly moving beyond trouble shooting, compliance and basic reporting towards Observability across an enterprise network, a new approach is required.

With the need to be able to collect all log, metrics, and traces from systems on-premises and in the cloud and be able to search and apply machine learning to petabytes of data to gain real-time insights, organizations are increasingly coming to the realization that the data ingest-based pricing system has rendered Splunk cost prohibitive. We are constantly hearing from companies who can no longer afford to add more data to their Splunk implementation and, as a result, are faced with the difficult decision to drop valuable data sources for security and Observability use cases or search for a new and better solution. Augmenting your Splunk implementation with Elysium Analytics, built on Snowflake natively, may be the answer.



Elysium Analytics was built for Observability running on Snowflake natively, used as core data platform for all data pipeline, storage and analytics, a zero-operations data warehouse running on multiple cloud platforms with no concurrency limitations and no degradation of response time regardless of load with cloud-scale compute. A recent study by Forrester shows that Snowflake's customers on average achieve a 3-year ROI of 612% with data warehouse applications and we are seeing similar gains for Elysium Analytics security and Observability use-cases. We have found that organizations moving their data from Splunk can save in excess of 90% by migrating their data to Snowflake with Elysium Analytics.

Elysium Analytics provides an observability solution for IT Ops, Dev Ops, and Security teams that need observability of all their log data, metrics, and traces with longer data retention. We load data from AWS, Azure, or GCP storage containers or load directly from cloud apps and on-premises sources to Snowflake. Leveraging the Snowflake platform with infinite compute scale-up and scale-out capabilities as well as unlimited cloud storage, there is no operational overhead from adding nodes, migrating indexes, and re-adjusting shards. Additionally, you have full access to all the compute and storage you require on a pay-as-you-go basis.



Search and dashboards are provided cost-effectively on Kibana, a search and visualization tool familiar to most operators. Additionally, Looker is integrated with the solution for advanced analytics and out-of-the-box dashboards as well as your own custom dashboards. Machine learning-based analytics gives you user and entity-based anomaly detection across all your data.

Migrating from Splunk in four easy steps

While there is generally not an easy way to migrate data from one platform to another, we have made the process simple with the following process:

1 STEP

Determine data sources you don't ingest into Splunk

Most organizations have outgrown Splunk when it comes to data sources. Important new data sources are often not shipped to Splunk because the license would have to be upgraded at significant cost, a source had to be dropped due to increased data volume, or adding to a Splunk implementation would be complicated and add significant operational overhead. Following an "80/20 rule" for observability may keep you within budget but never without serious compromise to your ability understanding what is happening on your infrastructure and why performance is not what it should be. Ingesting these previously ignored data sources into Snowflake allows for better observability and new use cases. With Elysium Analytics and Snowflake billing on usage only and high compression rates, getting started bringing data into Snowflake can be done at a very low cost.

2 STEP

Identify data sources to migrate from Splunk to Elysium Analytics

An important part of the migration planning is to understand what data sources are currently shipped to Splunk. The easiest way to do this is via a SPL query:

```
| eventcount summarize=false index=* index=_*
| dedup index
| fields index
| map maxsearches=100 search="|metadata type=sourcetypes index=\"\$index\$\""
| eval index=\"\$index\$\"
| fields index sourcetype
```

Run the query in the Splunk search application and choose the list of data in your preferred format.

3 STEP

Migrate existing Splunk data flows to Snowflake with Elysium Analytics

With the Elysium Analytics' data collection, an end-to-end cloud-based service for simple integration of any source, data migration is simple.

If you want to continue to load data to Splunk while you also load to Elysium Analytics, you can bifurcate the data using the Splunk Forwarders. However, Splunk imposes technical and licensing constraints on how you can send data to third-party systems with their Forwarders; review documentation and licensing before you invest effort into setting this up. If this is an option, this is an easy way to augment Splunk with additional data sources and greater data retention on Snowflake.

- If you are using Universal or Light Forwarders, simply configure to send the logs to TCP PORT and stream the logs to Elysium's data collection cloud and we will handle the parsing, loading, and data mapping into Snowflake. Alternatively, you can log to local file and then ship the data to Elysium's data collection cloud.
- If you are using Heavy Forwarders, you can simply configure to ship the logs in Syslog format to Elysium Analytics' data collection service.

To augment your Splunk implementation, load all your data into Snowflake while Splunk receives a sub-set of your data for certain use-cases. This will allow you to reduce the amount of data ingested into Splunk, reducing your subscription or license cost significantly, and reduce the data retention time on expensive Splunk storage while gaining greater Observability and analytics capabilities.

Over time you have the option to replace the Splunk Forwarders and collect the data directly from the source with Elysium Analytics' for loading into Snowflake.



Beats and Minifi compatible for simple integration and leverage of existing enterprise collection frameworks as well as integration with any 3rd party source.

Connect your sources leveraging existing Logstash implementations or use our Kafka and proprietary connectors based on REST APIs and webhooks.

Parse legacy device data sources in Logstash and modern data sources in JSON and Java.

Enrich data in real-time with Identity, Asset, Geolocation, and Threat Intelligence, as well as data from lookup tables built into the storage platform data pipeline.

4 STEP

Migrate existing data in Splunk to Snowflake with Elysium Analytics

Although most organizations will prefer to start fresh with current data, there are often use cases that require migration of old data in Splunk to Snowflake. The easiest way to is to export the data from the Splunk interface. Use the Splunk dump cmd. to export the data as explained here: <https://docs.splunk.com/Documentation/Splunk/8.1.2/Search/Exportdatausingdumpcommand> You can also use the [Splunk API](#) to export data, or you can connect via [ODBC](#). The approach you use will depend on the use case, volume of data, and type of data you are working with.

Whichever way you decide to load your data into Snowflake, Elysium Analytics can assist.

Observability and “data spelunking” with Elysium Analytics


When you have all your log data from all your sources on Snowflake, Elysium Analytics will provide you with Observability with speed, scale, and relevance. With its logical architecture running natively on Snowflake, the solution will scale with your organization and offer the flexibility required for your future needs without operational overhead.





About Elysium Analytics


Elysium Analytics is a machine learning based log analysis solution for Observability-minded, mid-sized to large enterprises who are challenged by the volume of security and operational log data today, both from an infrastructure as well as an analytics and detection perspective. We have simplified onboarding of data, provided a scalable data lake analytics platform, and search on a pay-as-you-go basis. Since we are built on top of Snowflake, our SaaS solution is truly a cloud scale security analytics platform that removes the barriers from ingesting, contextualizing, searching, analyzing, and storing log data with a cost-effective and low-risk service. Unlike other log analysis vendors in the market, our SaaS offering is licensed on a usage basis, lowering cost and removing financial risk. You pay a low price for storage, and compute is billed by the minute of usage. Additionally, we have an open platform with no vendor lock-in, customizable analytics models, as well as APIs for end user development of analytics models.



 Elysium Analytics, Inc. 2550 Great America Way, Santa Clara, CA 95054

 elysiumanalytics.ai

 Phone: +1 (669) 209-0801

 info@elysiumanalytics.ai