Migrate from Elasticsearch to Snowflake with Elysium Analytics

elysium
ANALYTICS

Migrate from **Elasticsearch** to **Snowflake** with **Elysium Analytics**

Elasticsearch, previously known as the ELK Stack, a search engine, Integrating Beats, Logstash, Elasticsearch, and Kibana. Elasticsearch was designed to provide near real-time search of documents with a distributed architecture, which means that indices can be divided into shards which are replicated for resiliency. Scale is achieved by adding nodes which each host one or more shards. Related data is often stored in the same index, which consists of one or more primary shards, and zero or more replica shards. Once an index has been created, the number of primary shards cannot be changed.

Organizations generate ever-increasing amount of data. Every application, server, network device, security device, and IoT device produce an enormous amount of data that has the potential to provide deep business insights with log data analytics solutions such as Elasticsearch. Customers often begin their journey of operational monitoring for their organization with the open source platform ELK Stack since the barrier to entry is low. However, as the monitoring or analytics teams start to ingest more data and develop more use cases, they often run into problems such as:

- Large-scale deployments lead to high cost and operational overhead
- Mid- to long-term data retention requirements lead to high cost and operational overhead
- Kibana not supporting all use case requirements
- Integration of data sources can be complicated
- Open source Elastic Stack is unsupported

As a result, as the analytics initiatives of organizations mature, customers realize there is a need to migrate to a platform that provides broader capabilities, lower cost, and lower operations overhead.

Migrating your Elasticsearch implementation with Elysium Analytics and Snowflake may be the answer.

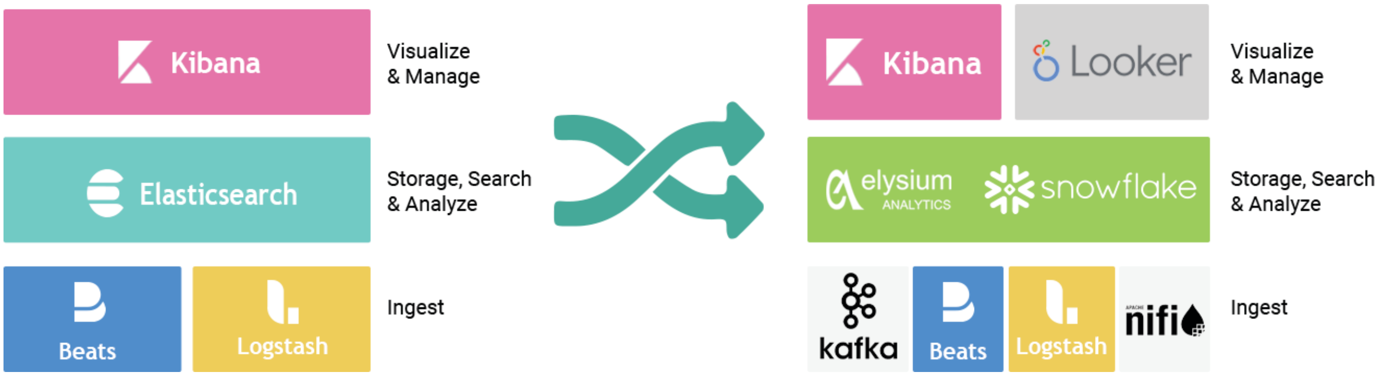

Elysium Analytics was built for cloud scale observability leveraging Snowflake, a zero-operations data warehouse running on multiple cloud platforms with no concurrency limitations and no degradation of response time regardless of load with cloud scale compute. A recent study by Forrester shows that 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 Elasticsearch can save in excess of 90% by migrating their data to Snowflake.

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. You have full access to all the compute and storage you require on a pay-as-you-go basis. 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 with Kibana, retaining the search and visualization experience identical that of Elasticsearch. 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.

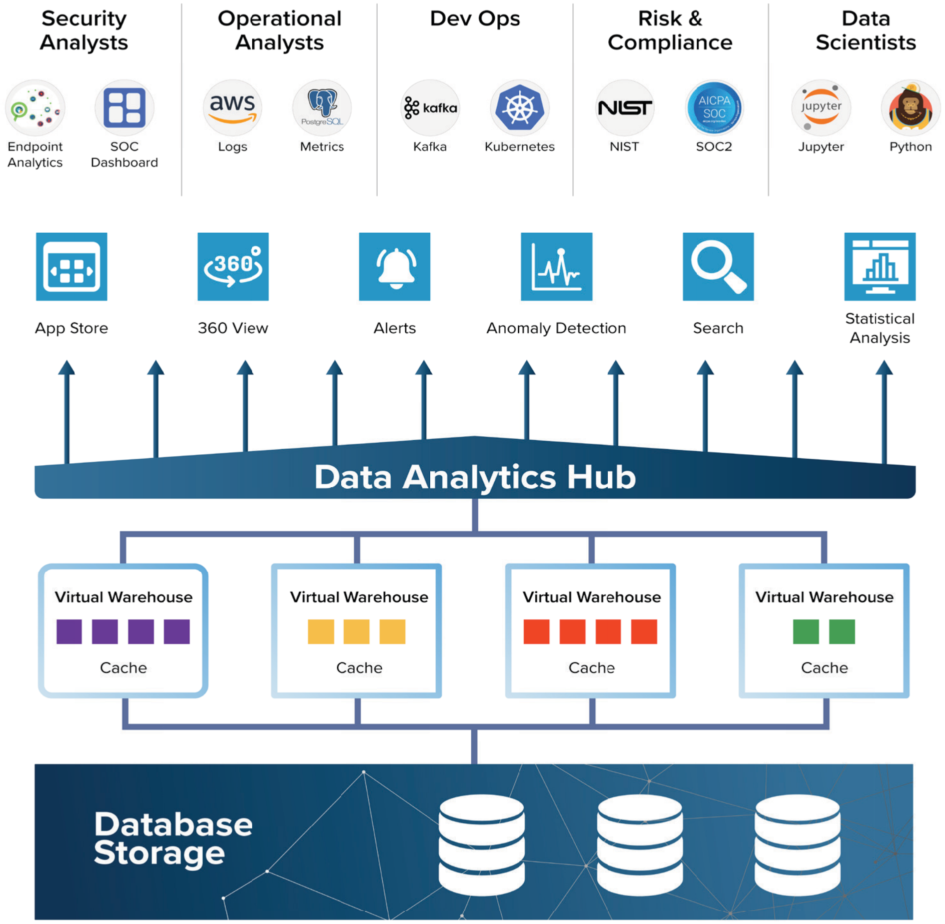

Migrating from **Elasticsearch in Five easy steps**

Elysium Analytics has architected a unique solution by replacing the middle layer of the Elastic Stack, Elasticsearch, with our own proprietary software running on Snowflake. This allows for an organization to retain all the features and benefits of data collection and parsing with Beats and Logstash as well as search and visualization capabilities of Kibana while providing a true cloud-scale platform removing operational overhead and lowering overall TCO dramatically.

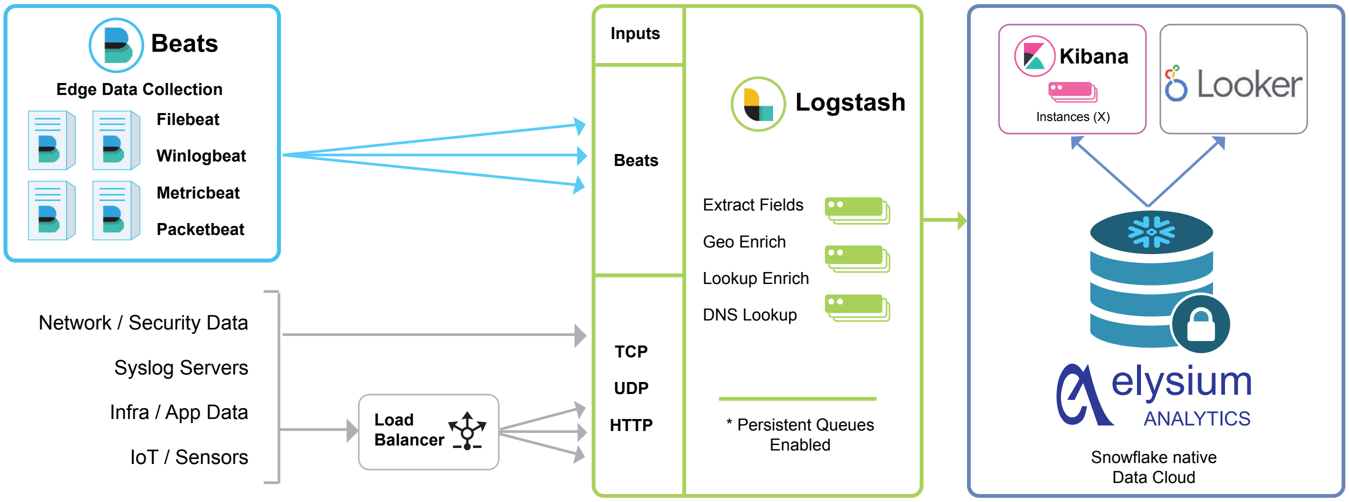

Migration from Elasticsearch can be done with a simple configuration on Logstash: Add the Elysium Analytics Logstash plugin which directs traffic to Elysium Analytics in parallel to traffic to Elasticsearch. Once you have validated that all the data is flowing to Elysium Analytics, you can terminate the feed to Elasticsearch without any disruption to the existing data collection framework.

If you are not quite ready to fully migrate from Elasticsearch right away, you can augment your Elasticsearch implementation by loading your data to both solutions in parallel for a period of time. This is also a great way for an organization that appreciates the fast Elasticsearch response time on smaller data sets, from the past day to 7 days for example, while gaining greater observability, analytics capabilities, and cost-effective data retention with the Elysium Analytics solution.

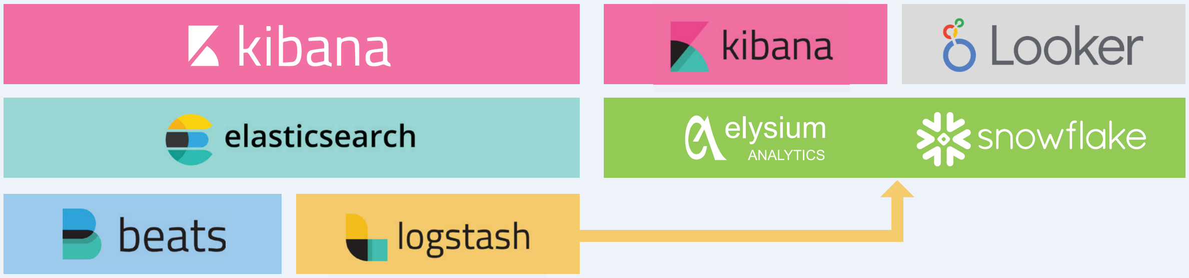

Step 1: Determine data sources you don’t ingest into Elasticsearch

Most organizations have outgrown Elasticsearch when it comes to data sources. Important new data sources are often not shipped to Elasticsearch 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 an Elasticsearch 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.

Step 2: Identify data sources to migrate from Elasticsearch to Elysium Analytics

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

1. Run any of below command to get list of indices in Elasticsearch

For non-ssl enabled Elasticsearch:

```
curl -XGET “http://<ELASTIC_SRC_HOST>:9200/_cat/indices?h=index” -u < ELASTIC_SRC_USERNAME >:< ELASTIC_SRC_PASSWORD >
```

Or

For ssl enabled Elasticsearch:

```
curl -XGET "https://< ELASTIC_SRC_HOST >:9200/_cat/indices?h=index" -u < ELASTIC_SRC_USERNAME >:< ELASTIC_SRC_PASSWORD > --insecure
```

2. For each index that we got from above command, run below elasticdump commands to migrate indices to Elysium Analytics

a.
```
Elasticdump --input=http://<ELASTIC_SRC_USERNAME>:<ELASTIC_SRC_PASSWORD>@<ELASTIC_SRC_HOST>:9200/<INDEXNAME_HERE> --output=http://<ELYSIUM_ELASTIC_USERNAME>:<ELYSIUM_ELASTIC_PASSWORD>@<ELYSIUM_ELASTICSEARCH_HOST>:9200/<INDEXNAME_HERE> --type=mapping
```

b.
```
elasticdump --input=http://<ELASTIC_SRC_USERNAME>:<ELASTIC_SRC_PASSWORD>@<ELASTIC_SRC_HOST>:9200/<INDEXNAME_HERE> --output=http://<ELYSIUM_ELASTIC_USERNAME>:<ELYSIUM_ELASTIC_PASSWORD>@<ELYSIUM_ELASTICSEARCH_HOST>:9200/<INDEXNAME_HERE> --type=data
```

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

Step 3: Migrate existing Elastic 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 ElasticSearch while you also load to Elysium Analytics, you can bifurcate the data using the Elysium Analytics Logstash plugin. This is an easy way to validate data loading to Elysium Analytics and Snowflake while it also allows you to add any additional data sources.

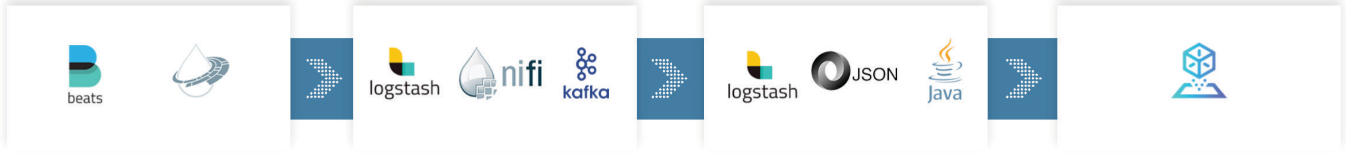

Beats and Minifi compatible for simple integration and leverage of existing enterprise collection frameworks as well as integration with any 3rd party source.

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.

Step 4: Migrate existing Kibana dashboards and visualizations to Elysium Analytics

1. Export all the saved objects (saved search / visualizations / dashboards / index-patterns etc.) from ELK/Kibana

Cmd to run:

a) go to the Kibana from which we need to import all the saved objects
b) go to Kibana -> management -> stack management (as shown highlighted in red below)

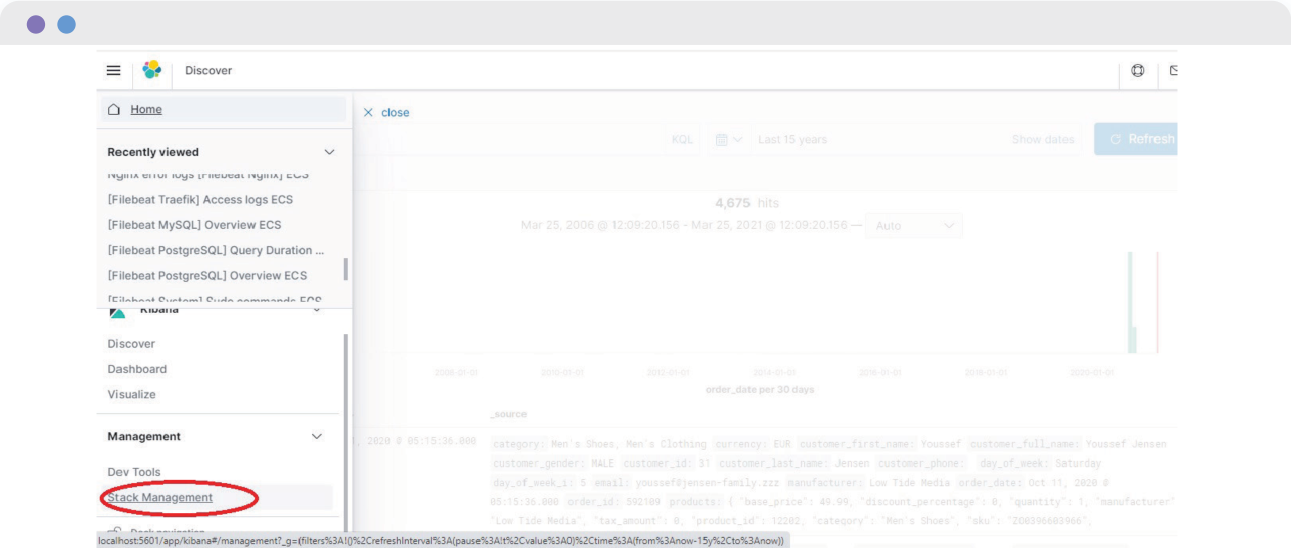

```
, 2020 @ 05:15:36.000   category: Men's Shoes, Men's Clothing currency: EUR customer_first_name: Youssef customer_full_name: Youssef Jensen
                        customer_gender: MALE customer_id: 31 customer_last_name: Jensen customer_phone:  day_of_week: Saturday
                        day_of_week_i: 5 email: youssef@jensen-family.zzz manufacturer: Low Tide Media order_date: Oct 11, 2020 @
                        05:15:36.000 order_id: 592109 products: { "base_price": 49.99, "discount_percentage": 0, "quantity": 1, "manufacturer"
                        "Low Tide Media", "tax_amount": 0, "product_id": 12202, "category": "Men's Shoes", "sku": "ZO0396603966",
```

```
localhost:5601/app/kibana#/management?_g=(filters%3A!()%2CrefreshInterval%3A(pause%3A!t%2Cvalue%3A0)%2Ctime%3A(from%3Anow-15y%2Cto%3Anow))
```

c) Click on Saved Objects as shown highlighted in red below

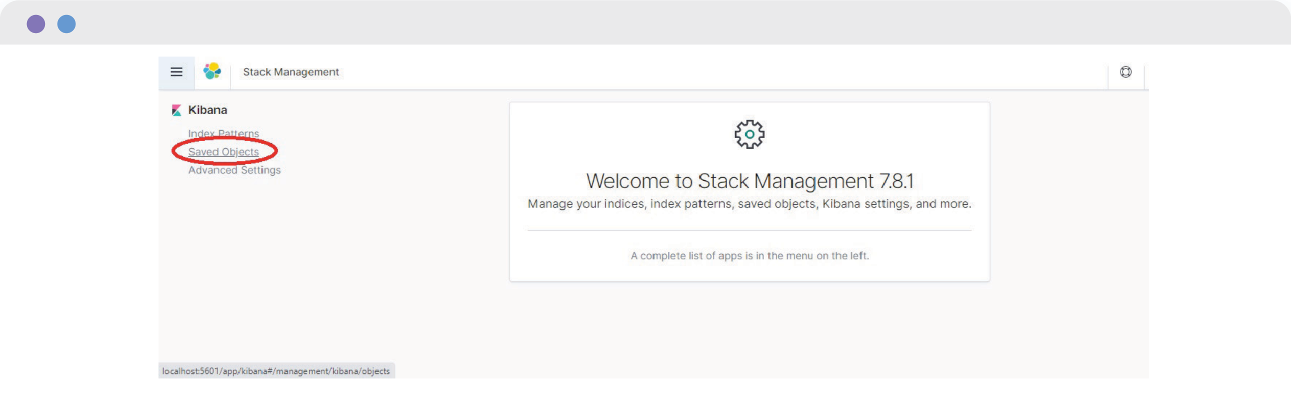

Welcome to Stack Management 7.8.1

Manage your indices, index patterns, saved objects, Kibana settings, and more.

A complete list of apps is in the menu on the left.

```
localhost:5601/app/kibana#/management/kibana/objects
```

d) Export Saved objects as shown highlighted in red below

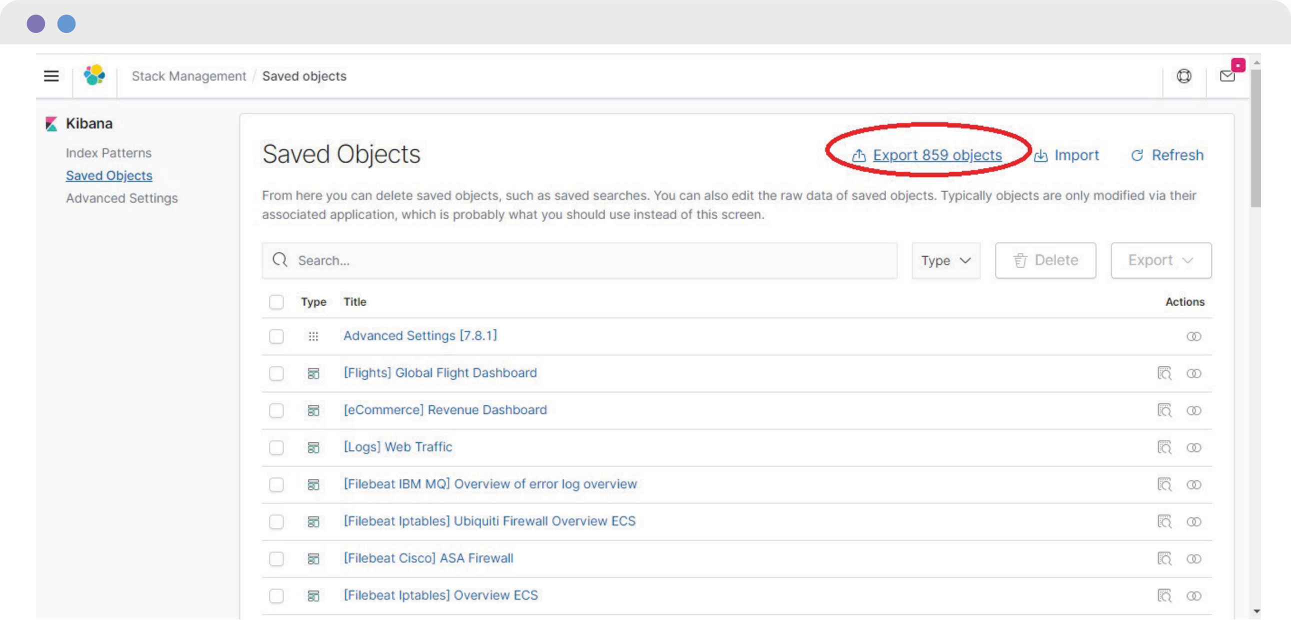

Saved Objects

From here you can delete saved objects, such as saved searches. You can also edit the raw data of saved objects. Typically objects are only modified via their associated application, which is probably what you should use instead of this screen.

	Type	Title	Actions
□		Advanced Settings [7.8.1]	
□		[Flights] Global Flight Dashboard	
□		[eCommerce] Revenue Dashboard	
□		[Logs] Web Traffic	
□		[Filebeat IBM MQ] Overview of error log overview	
□		[Filebeat Iptables] Ubiquiti Firewall Overview ECS	
□		[Filebeat Cisco] ASA Firewall	
□		[Filebeat Iptables] Overview ECS	

f) Select to include all the objects that are required to be exported to new Kibana instance as highlighted in red below

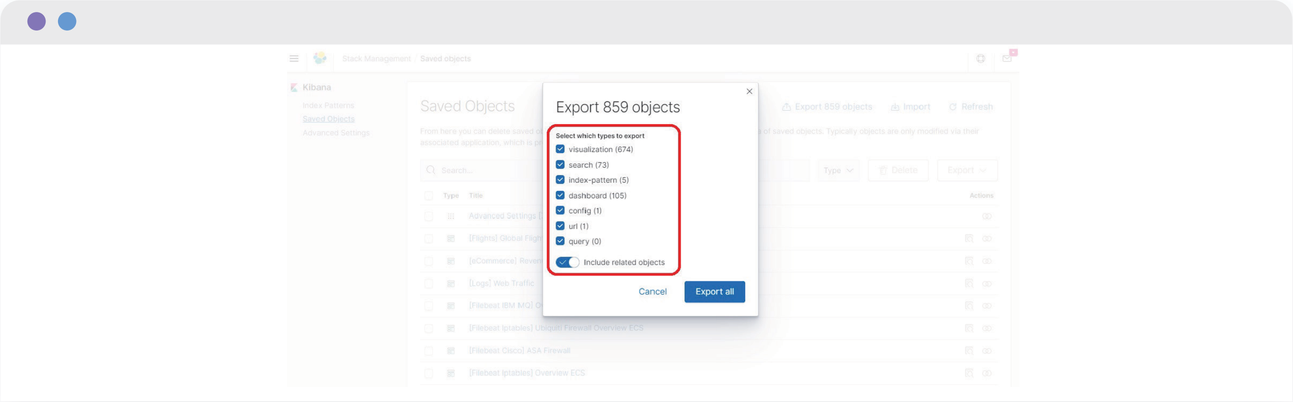

Stack Management / Saved objects

Kibana

Index Patterns

Saved Objects

Advanced Settings

Saved Objects

Export 859 objects

Select which types to export

☑ visualization (674)

☑ search (73)

☑ index-pattern (5)

☑ dashboard (105)

☑ config (1)

☑ url (1)

☑ query (0)

Include related objects

Cancel | Export all

g) Click on "Export All" button and save the ndjson file with a unique name such as “ElysiumKibanaExport.ndjson”

2. Import all the saved objects (saved search / visualizations / dashboards / index-patterns etc) to Elysium Analytics

a. Login to Elysium Analytics

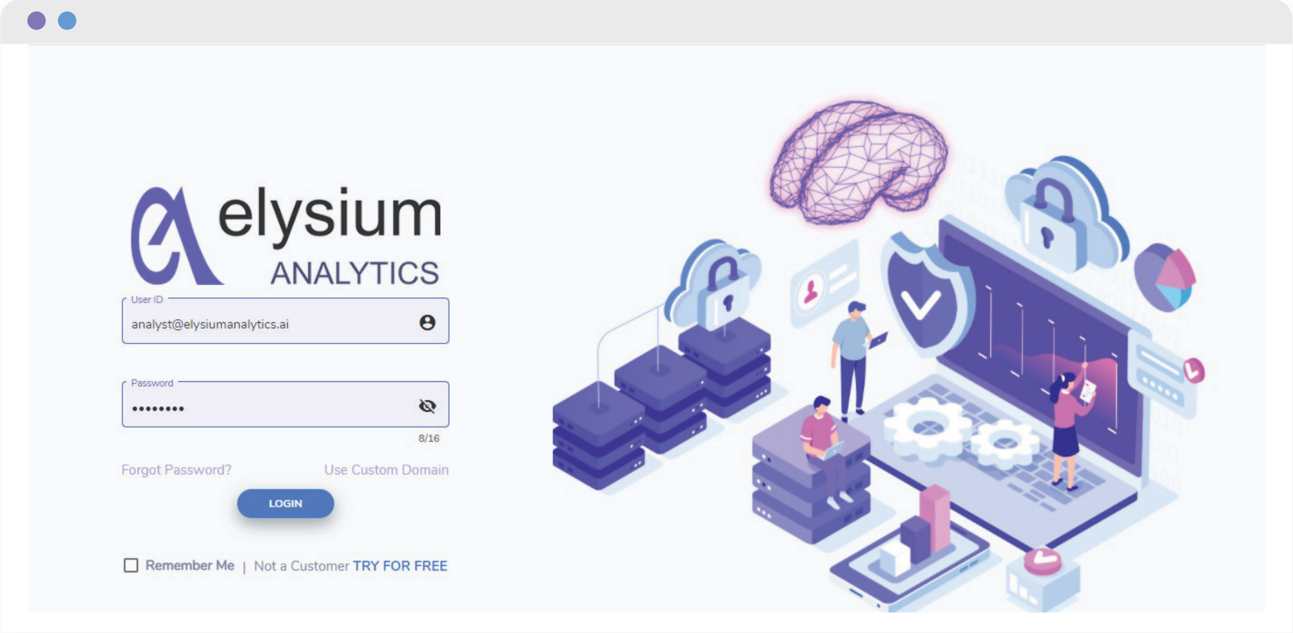

elysium ANALYTICS

User ID: analyst@elysiumanalytics.ai

Password: ••••••••

8/16

Forgot Password? | Use Custom Domain

LOGIN

□ Remember Me | Not a Customer TRY FOR FREE

b. Go to Kibana by clicking on the search icon as highlighted in red below.

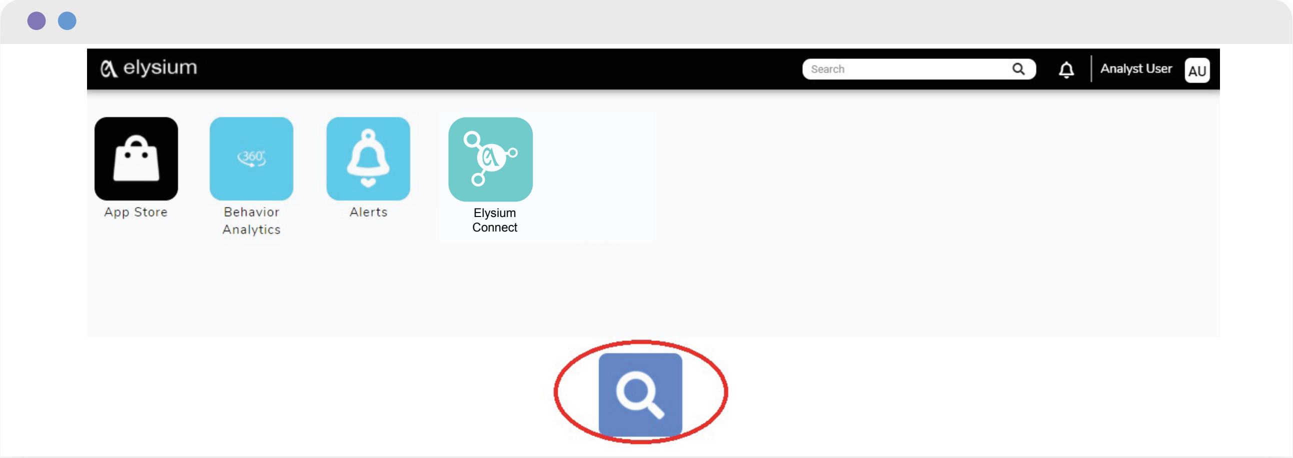

c. Go to Kibana -> management -> stack management (as shown highlighted in red below)

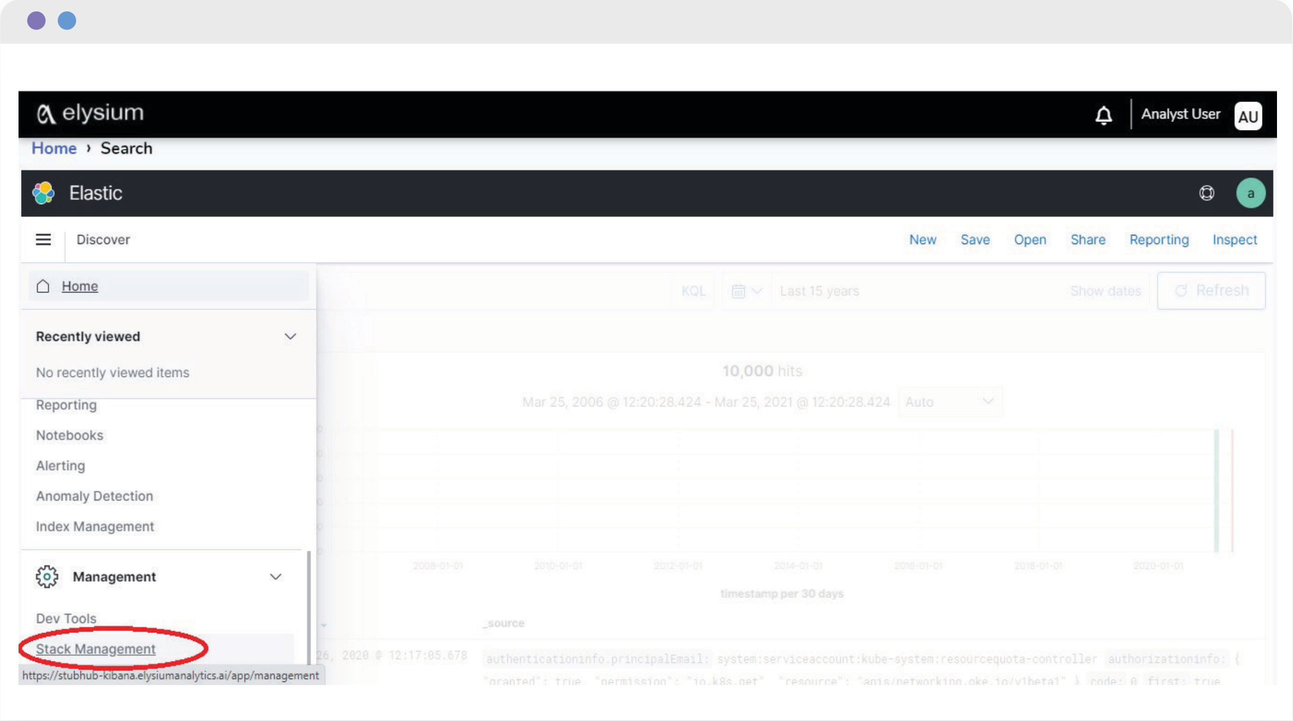

https://stubhub-kibana.elysiumanalytics.ai/app/management

d. Click on Saved Objects as shown highlighted in red below

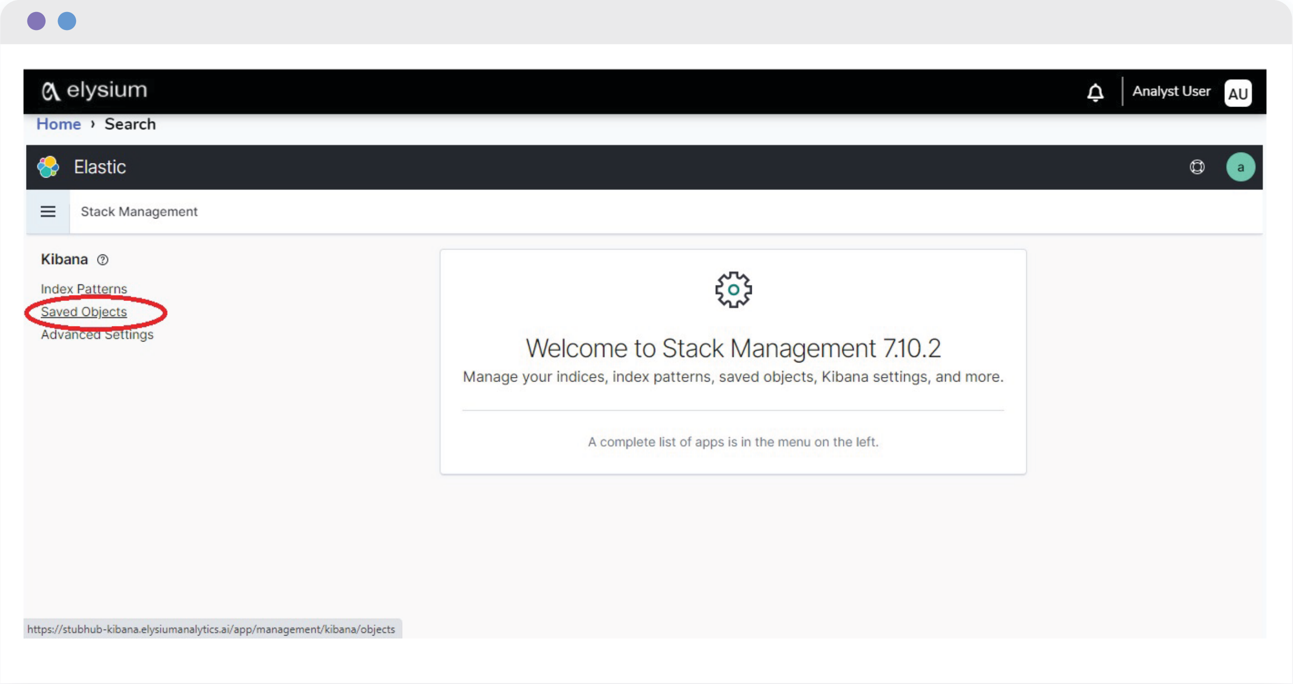

Welcome to Stack Management 7.10.2

Manage your indices, index patterns, saved objects, Kibana settings, and more.

A complete list of apps is in the menu on the left.

https://stubhub-kibana.elysiumanalytics.ai/app/management/kibana/objects

e. Click on “Import” button as shown highlighted in red below

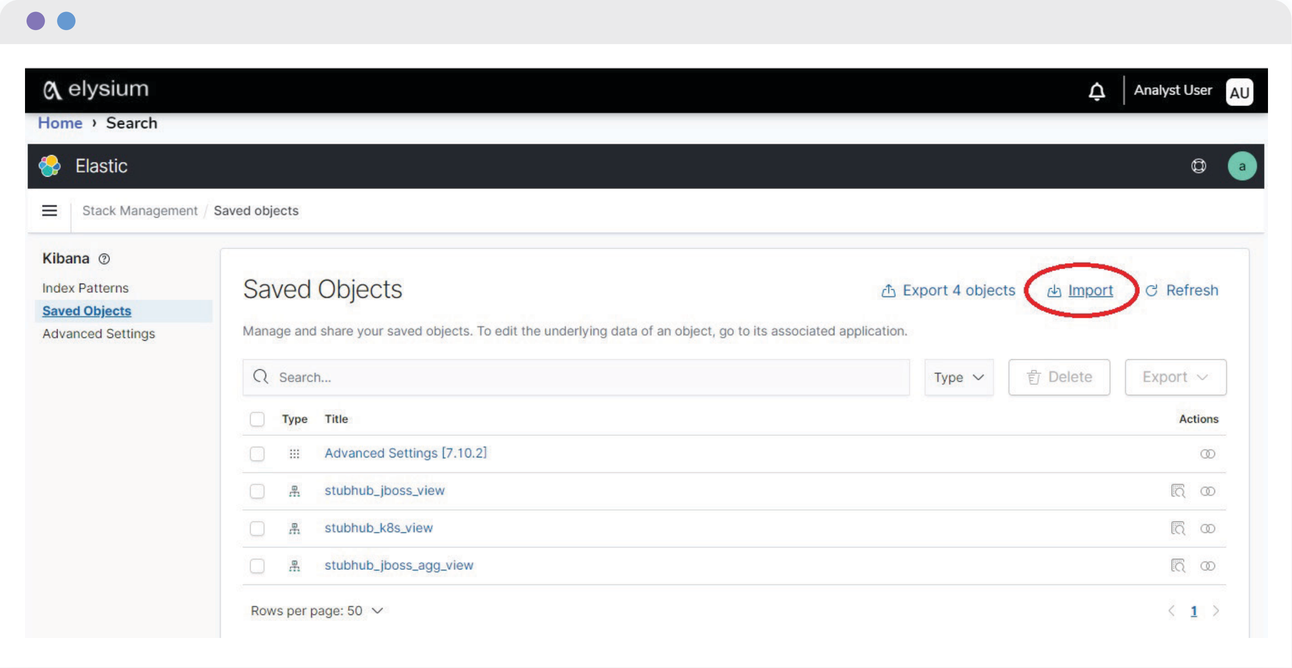

Manage and share your saved objects. To edit the underlying data of an object, go to its associated application.

Type	Title	Actions
	Advanced Settings [7.10.2]	
	stubhub_jboss_view	
	stubhub_k8s_view	
	stubhub_jboss_agg_view	

f. In the popup window, click on “select a file to import” as shown highlighted in red below and choose the above “ElysiumKibanaExport.ndjon” file

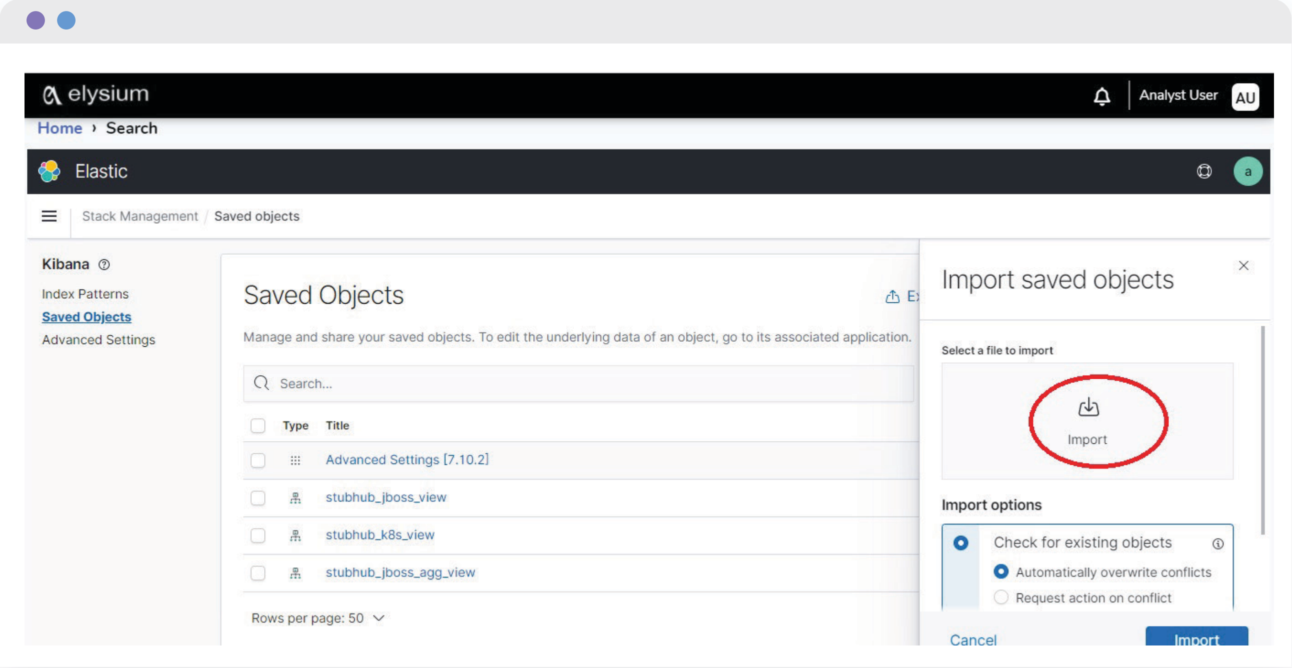

Manage and share your saved objects. To edit the underlying data of an object, go to its associated application.

Type	Title
	Advanced Settings [7.10.2]
	stubhub_jboss_view
	stubhub_k8s_view
	stubhub_jboss_agg_view

Import options

● Check for existing objects
- ● Automatically overwrite conflicts
- ○ Request action on conflict

g. After choosing the ndjson file, choose “check” for existing objects as highlighted in red and click on the “Import“ button as shown highlighted in red below

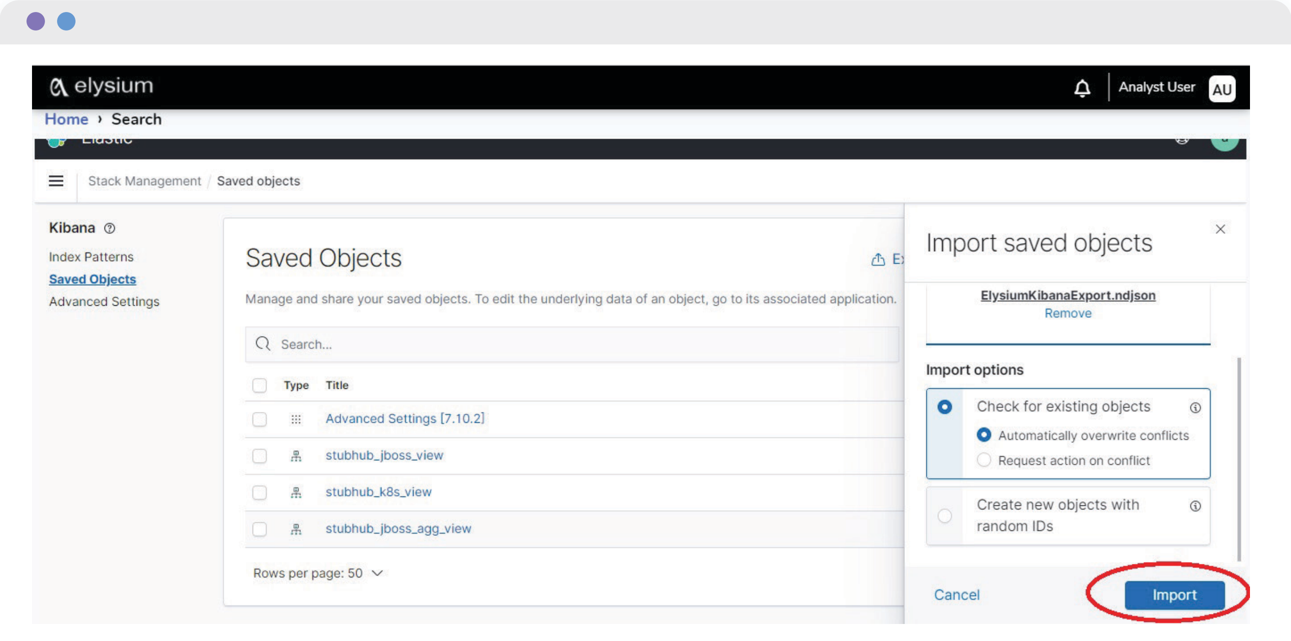

Manage and share your saved objects. To edit the underlying data of an object, go to its associated application.

Type	Title
	Advanced Settings [7.10.2]
	stubhub_jboss_view
	stubhub_k8s_view
	stubhub_jboss_agg_view

Rows per page: 50

Import saved objects

ElysiumKibanaExport.ndjson
Remove

Import options

- ● Check for existing objects
  - ● Automatically overwrite conflicts
  - ○ Request action on conflict
- ○ Create new objects with random IDs

h. All the objects in “ElysiumKibanaExport.ndjson” file are imported to the Elysium Analytics instance as illustrated below. Now click “Done”.

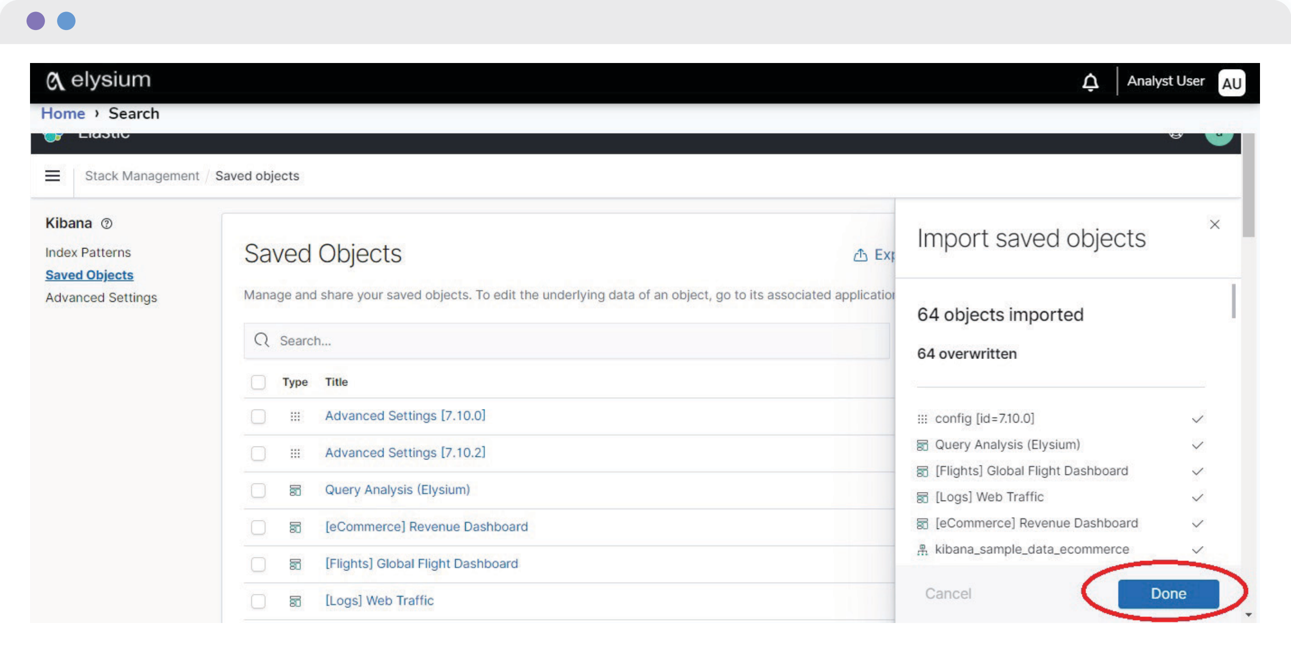

Manage and share your saved objects. To edit the underlying data of an object, go to its associated applicatio

Type	Title
	Advanced Settings [7.10.0]
	Advanced Settings [7.10.2]
	Query Analysis (Elysium)
	[eCommerce] Revenue Dashboard
	[Flights] Global Flight Dashboard
	[Logs] Web Traffic

Import saved objects

64 objects imported

64 overwritten

- config [id=7.10.0] ✓
- Query Analysis (Elysium) ✓
- [Flights] Global Flight Dashboard ✓
- [Logs] Web Traffic ✓
- [eCommerce] Revenue Dashboard ✓
- kibana_sample_data_ecommerce ✓

3. You can now use your imported dashboards.

Step 5: **Migrate existing data in Elasticsearch to Snowflake**

Although most organizations will prefer to start fresh with current data, there are often use cases that require migration of old data in Elasticsearch to Snowflake. The approach you use will depend on the use case, volume of data, and type of data you are working with and whichever way you decide to load your data into Snowflake, Elysium Analytics can assist.

Cloud Scale Observability and Search 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

Elysium Analytics, Inc.
2550 Great America Way, Ste 101,
Santa Clara, CA 95054

www.elysiumanalytics.ai

info@elysiumanalytics.ai