elastic stack

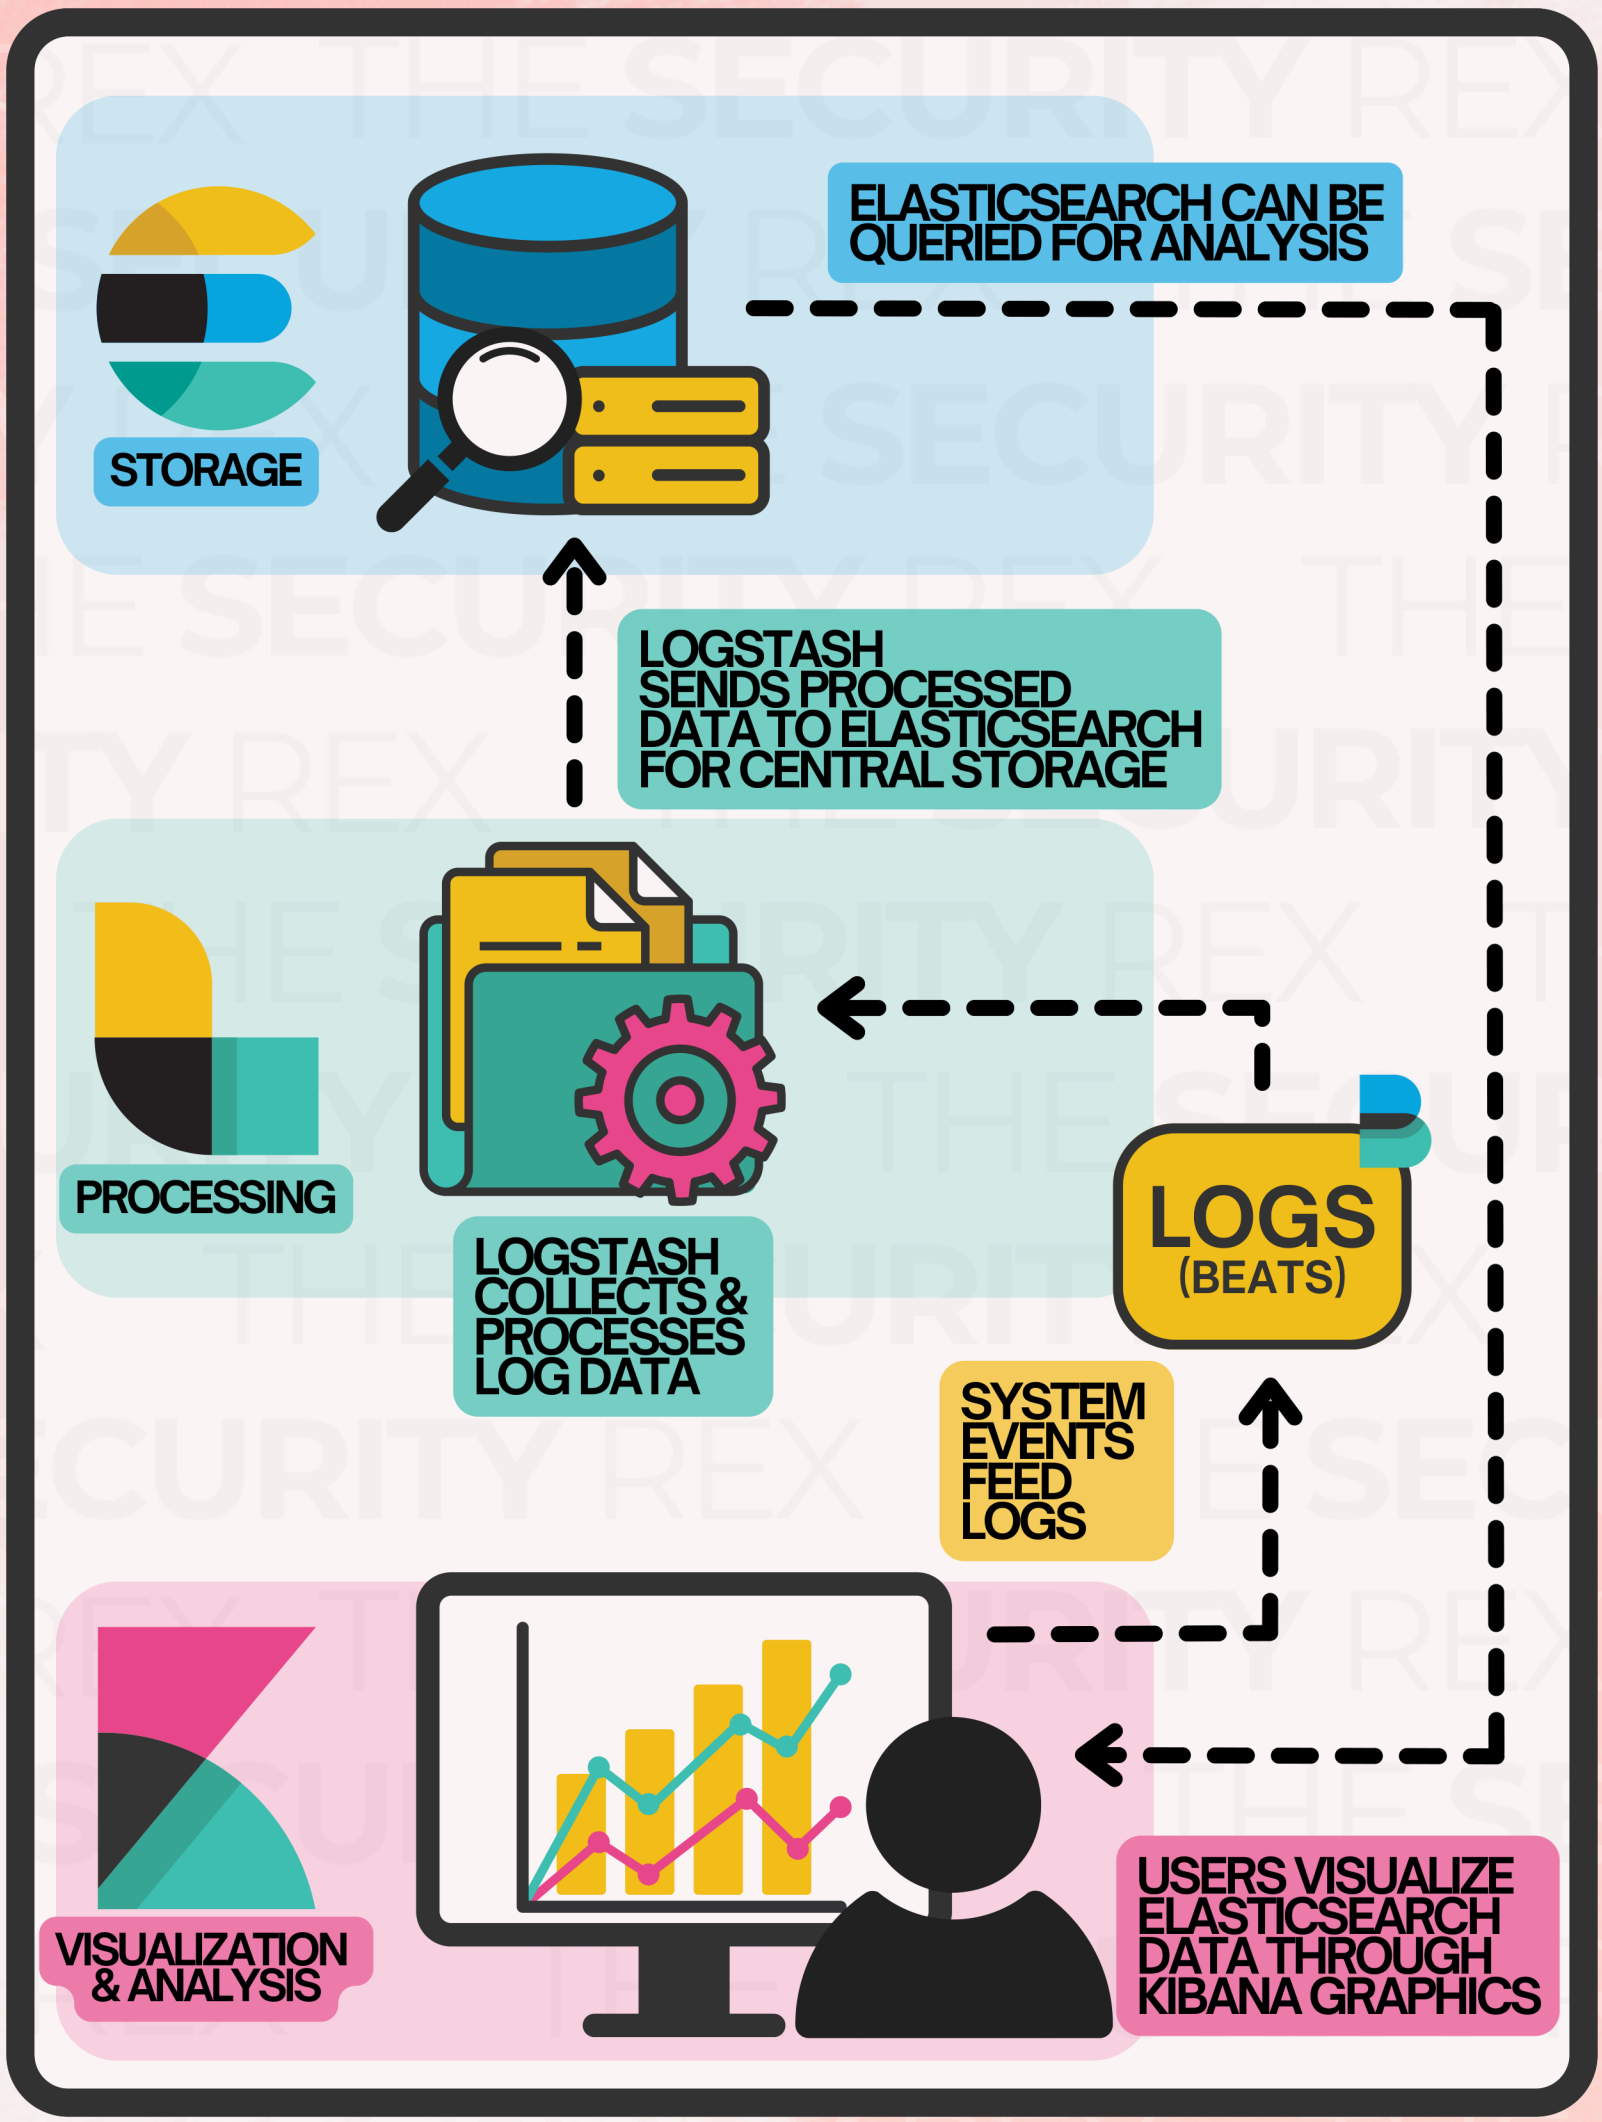

Log management involves the systematic collection, aggregation, and analysis of logs and audit trails generated by systems, applications, and devices.

The **ELK Stack** comprises **Elasticsearch**, **Logstash**, and **Kibana**.

- **Elasticsearch** is a distributed search and analytics engine.
- **Logstash** is for data processing and enrichment
- **Kibana** provides the user interface for data visualization.

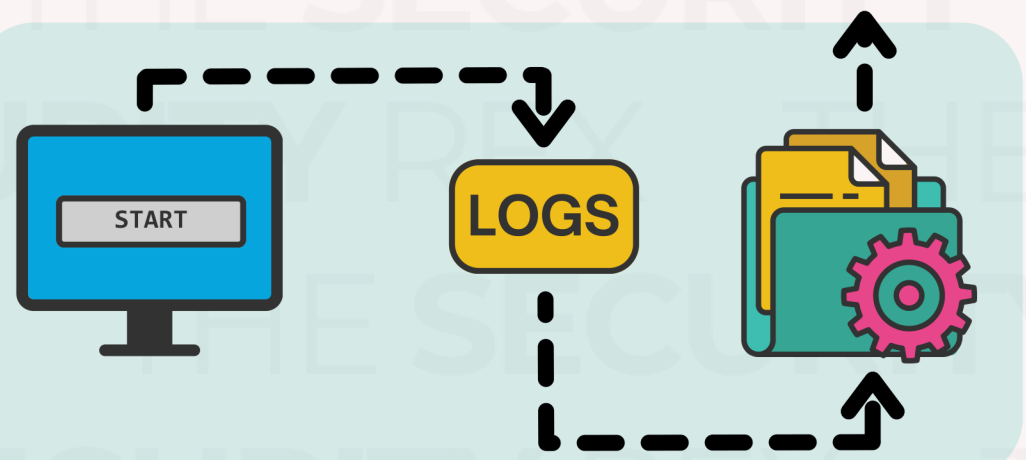

Logstash ingests logs and events (referred to as **Beats**) using ***pipelines*** to define the flow. Filters transform the data and then forward it to an output, often **Elasticsearch**.

Suppose you have Apache web logs and you want to parse this data before storing it in **Elasticsearch**.

By creating a **Logstash** pipeline, you can enrich log data, making it more useful and accessible in **Kibana**.

Create a Logstash configuration file

```
$ nano /etc/httpd/conf/<apache_logstash.conf>
```

Ensure Logstash is running

```
$ sudo systemctl status logstash
```

Point to the configuration file

```
$ bin/logstash -f
/etc/httpd/conf/<apache_logstash.conf>
```

```
# an example Logstash configuration file

input { file { path =>
"/path/to/apache/logs/access.log" start_position =>
"beginning" type => "apache_access" } }

filter { if [type] == "apache_access" { grok { match
=> { "message" => "%{COMBINEDAPACHELOG}" } } date {
match => [ "timestamp", "dd/MMM/yyyy:HH:mm:ss Z" ] }
} }

output { elasticsearch { hosts =>
["http://localhost:9200"] index => "apache_logstash"
} }
```

Input: Reads the Apache access log from the beginning.

Filter: Here, the ***grok*** filter parses a predefined pattern and the date filter parses the timestamp.

Output: Sets where processed logs will be stored.

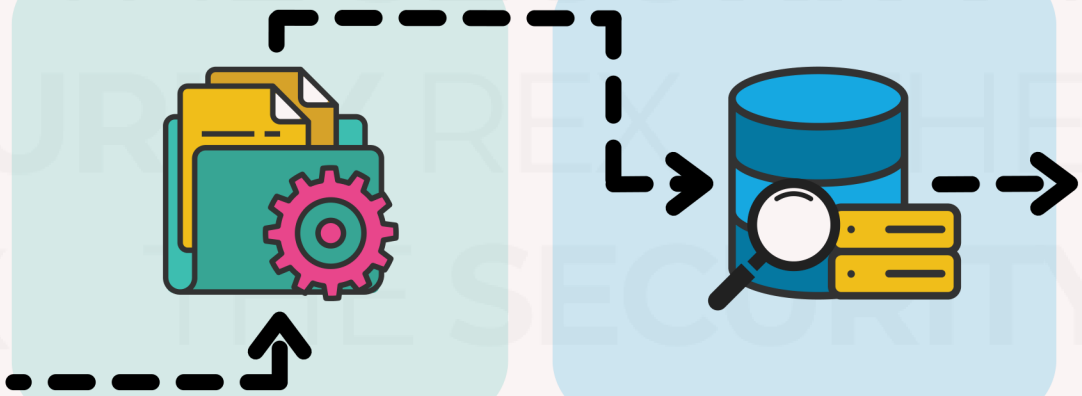

Elasticsearch stores data as JSON files, employing a schema-free nature to enable rapid indexing. Its RESTful APIs make data instantly accessible to **Kibana**.

Setting up an Elasticsearch instance is easy!

Import the Elastic GPG key

```
$ sudo wget -qO - https://artifacts.elastic.co/GPG-KEY-elasticsearch | sudo apt-key add -
```

Add the Elastic APT repository

```
$ sudo echo "deb https://artifacts.elastic.co/packages/oss-7.x/apt stable main" | sudo tee /etc/apt/sources.list.d/elastic-7.x.list
```

Update the package list and install Elasticsearch

```
$ sudo apt update
$ sudo apt install elasticsearch
```

Elasticsearch runs on port 9200 by default.

Restrict accesses in the ***elasticsearch.yml*** file located in ***/etc/elasticsearch/***.

Start Elasticsearch and enable boot-start

```
$ sudo systemctl start elasticsearch
$ sudo systemctl enable elasticsearch
```

Test if Elasticsearch is running

```
$ curl -X GET "localhost:9200/"
```

Elasticsearch security includes user authentication methods, role-based access controls, and secure communication via SSL/TLS.

Kibana is a web-based interface that adds visualization to **Elasticsearch** data. It does not directly feed data to **Logstash**. However, all system activity generates event logs which are ingested by **Logstash** for parsing.

Kibana Query Language (KQL) is used to query and analyze data.

```
field_name: “<field>”
```

```
timestamp: “<timeframe>”
```

```
numeric_field: “<num>”
```

You can integrate with **OSQuery** to centralize and analyze data at scale.

OSQuery monitors the real-time state of operating systems. **Logstash** facilitates the communication between **OSQuery** and **Elasticsearch**.