elastic

The webinar will start soon...

Elasticsearch Best Practice Architecture

whoami

Eric Westberg

Solutions Architect

Stockholm, Sweden

Joined Elastic last year

Webinar Housekeeping & Logistics

- **Slides and recording** will be available following the webinar
- Please ask questions via “Q&A”

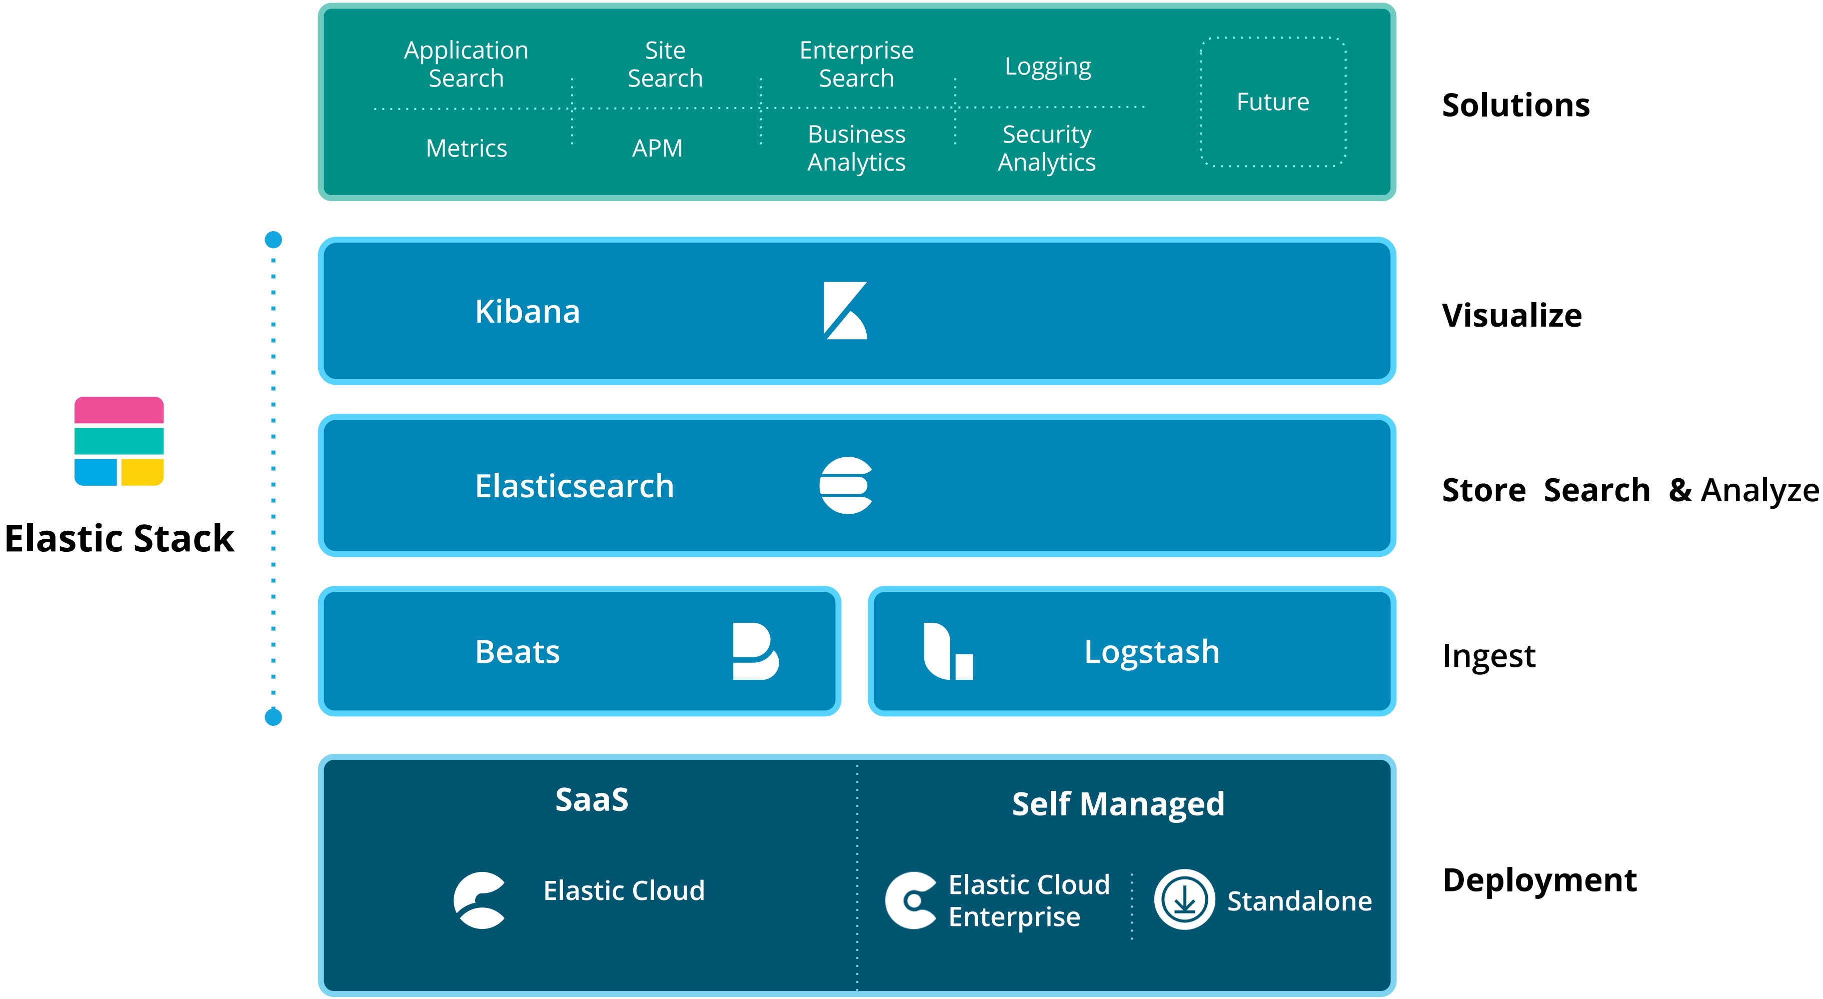

Elastic Stack

Solutions

Application Search | Site Search | Enterprise Search | Logging

Metrics | APM | Business Analytics | Security Analytics

Future

Visualize

Kibana

Store Search & Analyze

Elasticsearch

Ingest

Beats

Logstash

Deployment

SaaS

Elastic Cloud

Self Managed

Elastic Cloud Enterprise

Standalone

Inside an Elasticsearch Cluster

Elasticsearch Node Types

Nodes can play one or more roles, for workload isolation and scaling

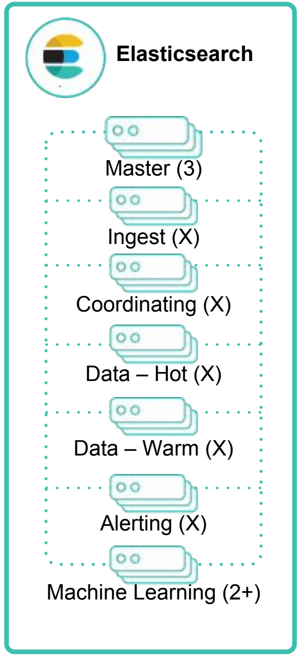

- Master Nodes
  - Control the cluster, requires a minimum of 3, one is active at any given time
- Data Nodes
  - Hold indexed data and perform data related operations
  - Differentiated Hot and Warm Data nodes can be used
- Ingest Nodes
  - Use ingest pipelines to transform and enrich before indexing
- Coordinating Nodes
  - Route requests, handle search reduce phase, distribute bulk indexing
  - All nodes function as coordinating nodes
- Alerting Nodes
  - Run alerting jobs
- Machine Learning Nodes
  - Run machine learning jobs

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Split Brain

- Cluster with 3 master eligible nodes
- Concern if network becomes partitioned
- The cluster would inadvertently elect two masters, which is referred to "split brain"

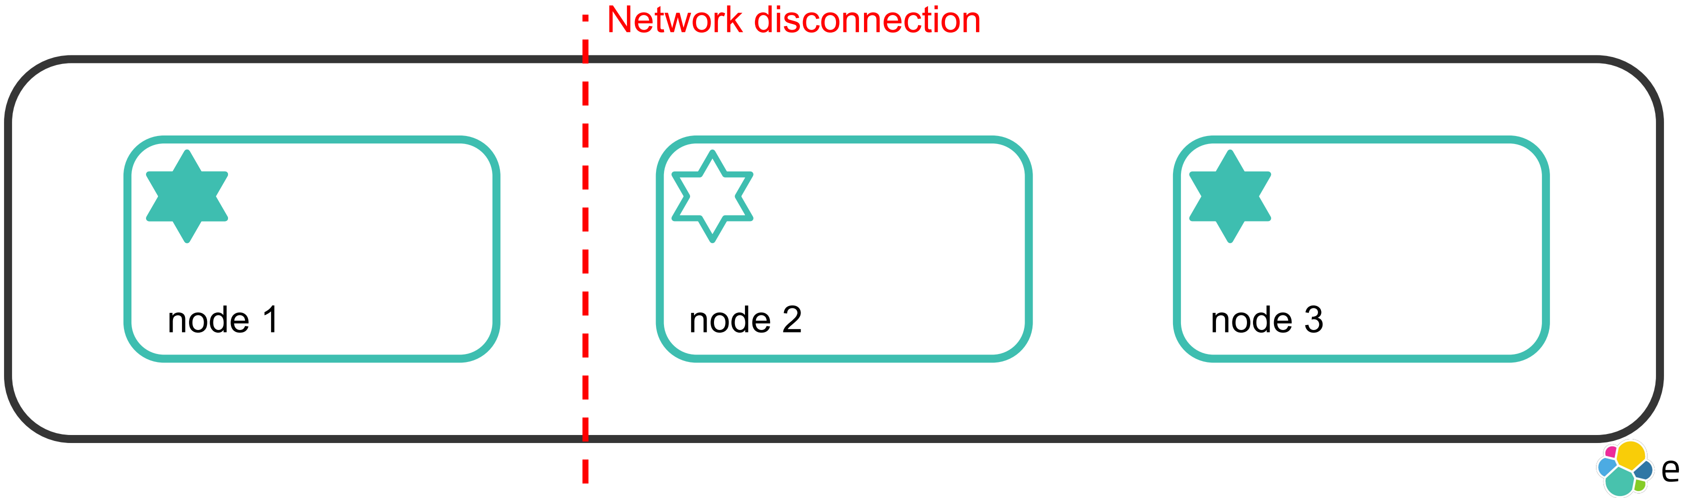

Avoiding Split Brain

- A master eligible node needs at least minimum_master_nodes votes to win an election
  - Setting it to a quorum prevents the split brain scenario
- Recommendation for production clusters is to have 3 dedicated master eligible nodes
  - with the setting `minimum_master_nodes = 2`

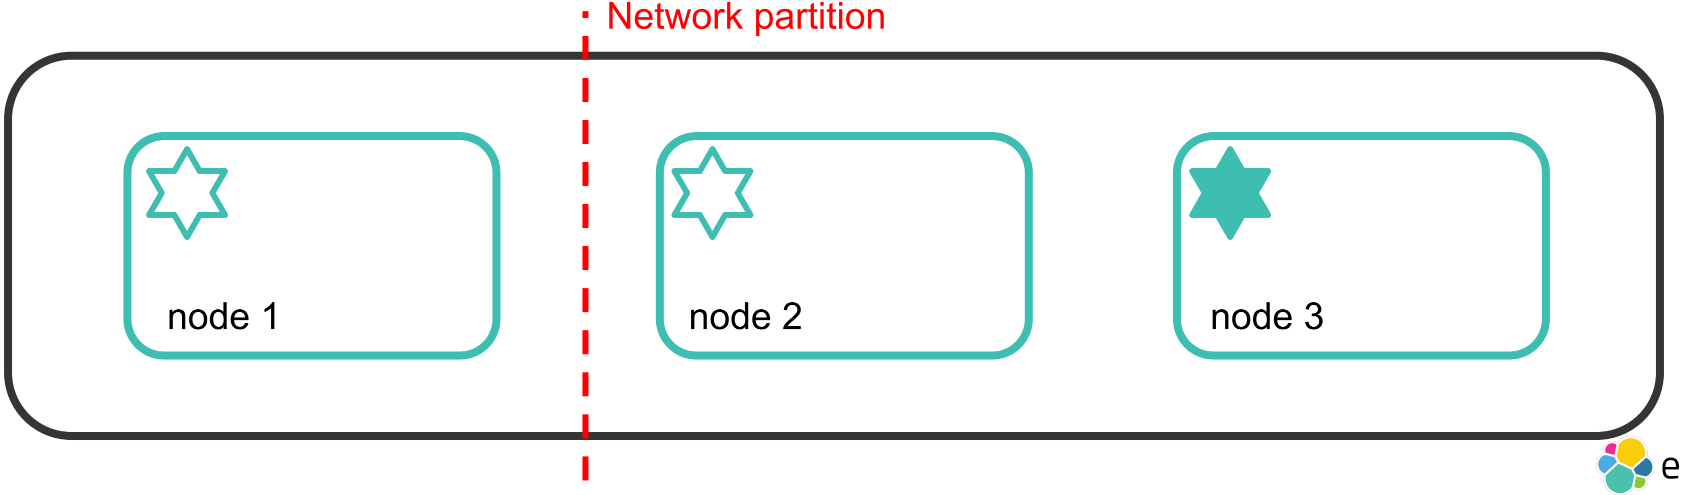

Inside a Large Elasticsearch Logging Cluster

Reduce infrastructure costs, isolate workloads, and manage data lifecycle

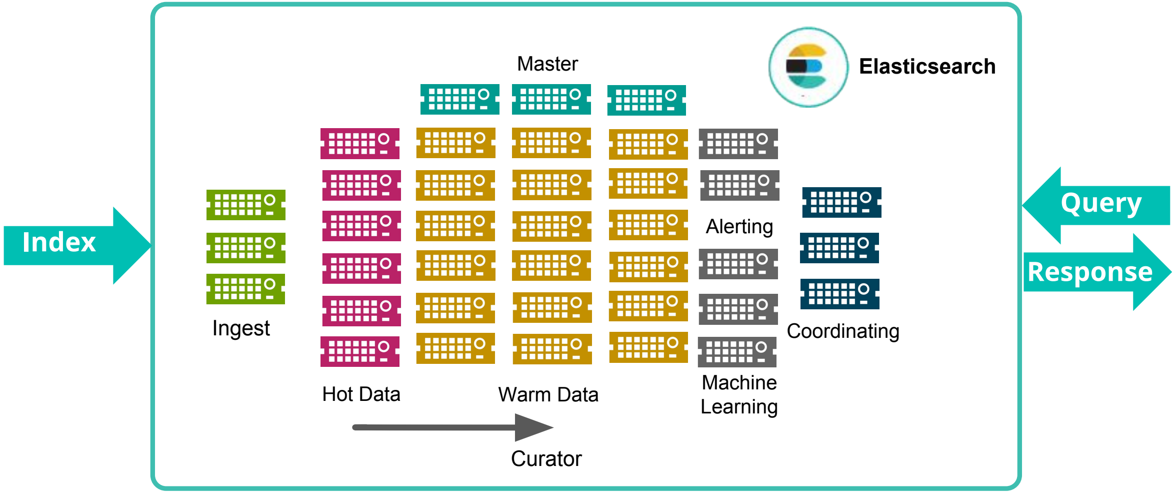

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Logging Architectures

Quick Start

Beats, Elasticsearch and Kibana

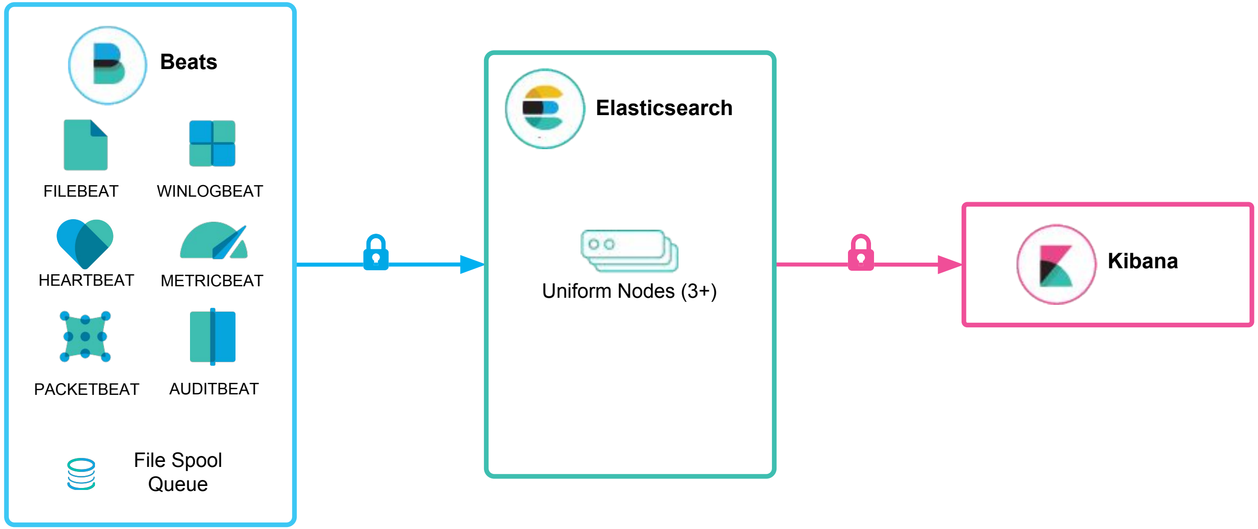

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Advanced Processing and Resiliency

Adding Logstash processing, differentiated Elasticsearch node types

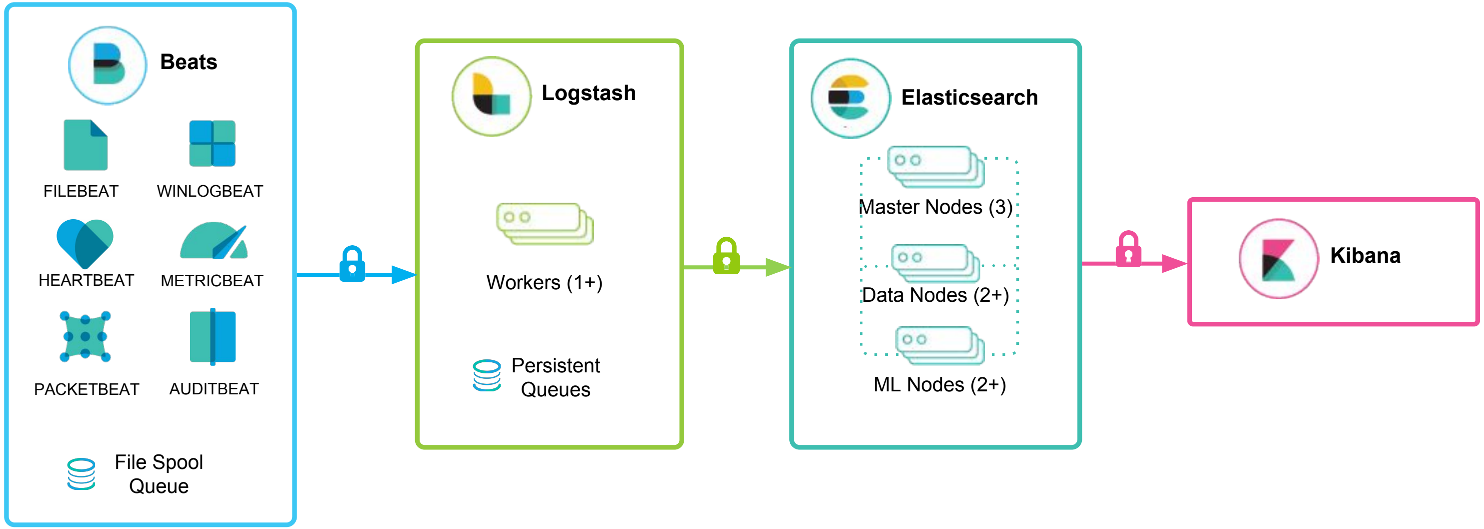

Flexible ingestion and input sources

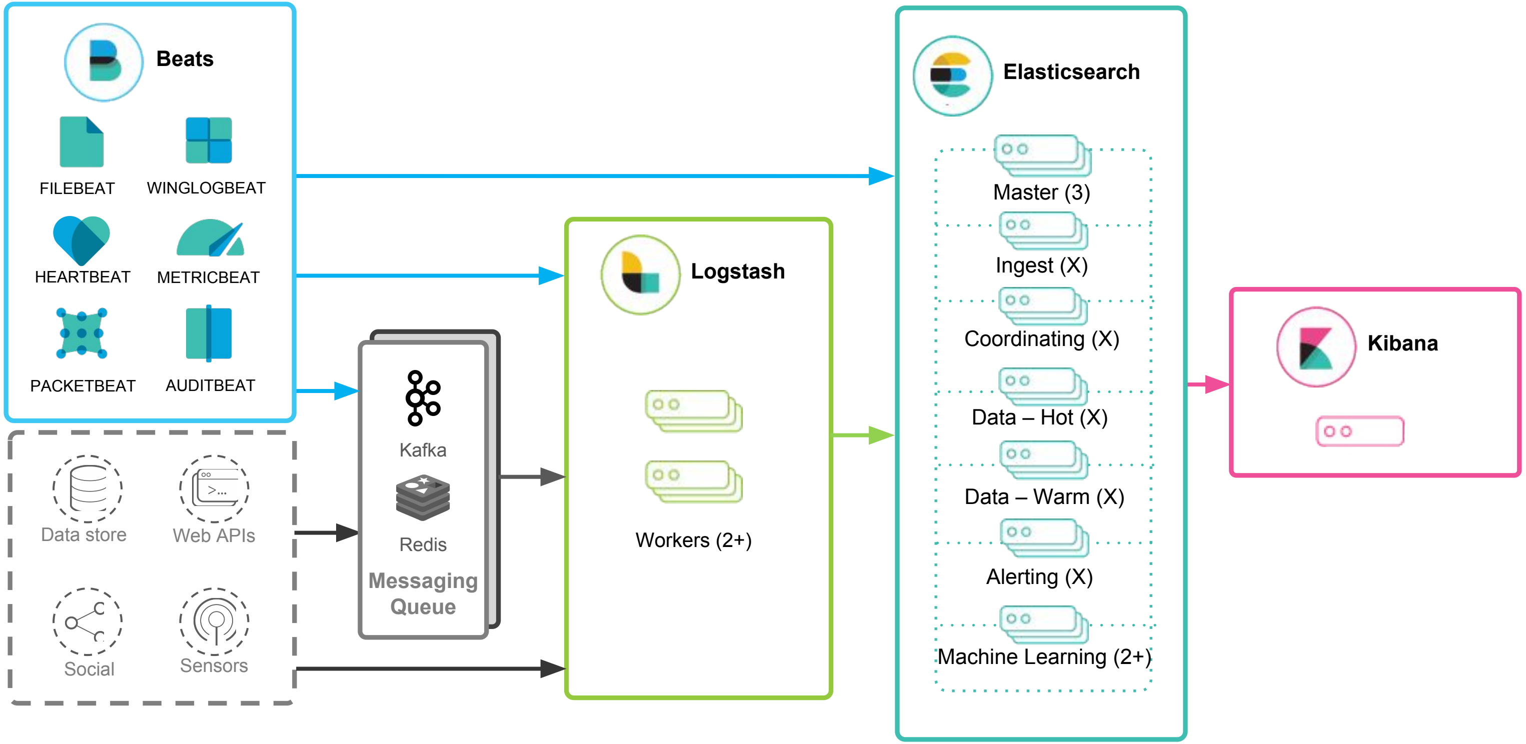

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Application Metric Collection with Elastic APM

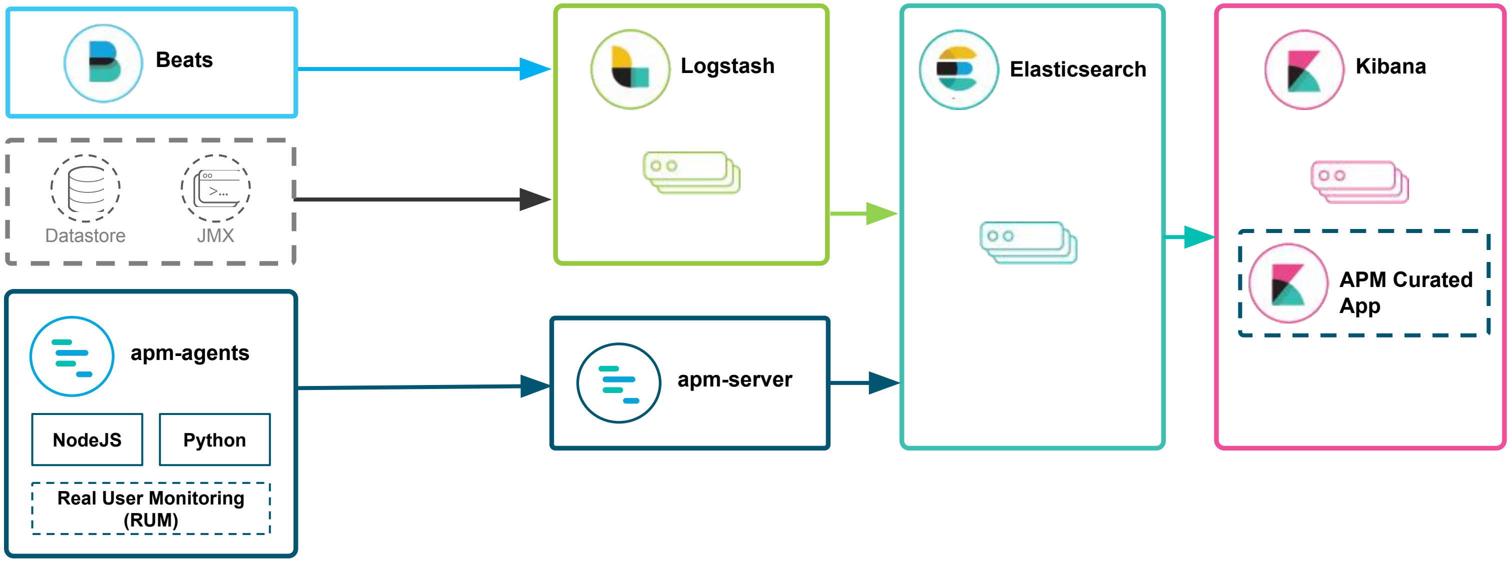

Deployment Best Practices

elastic

Centralized Monitoring Cluster

Maintain isolated monitoring cluster for monitoring workload isolation

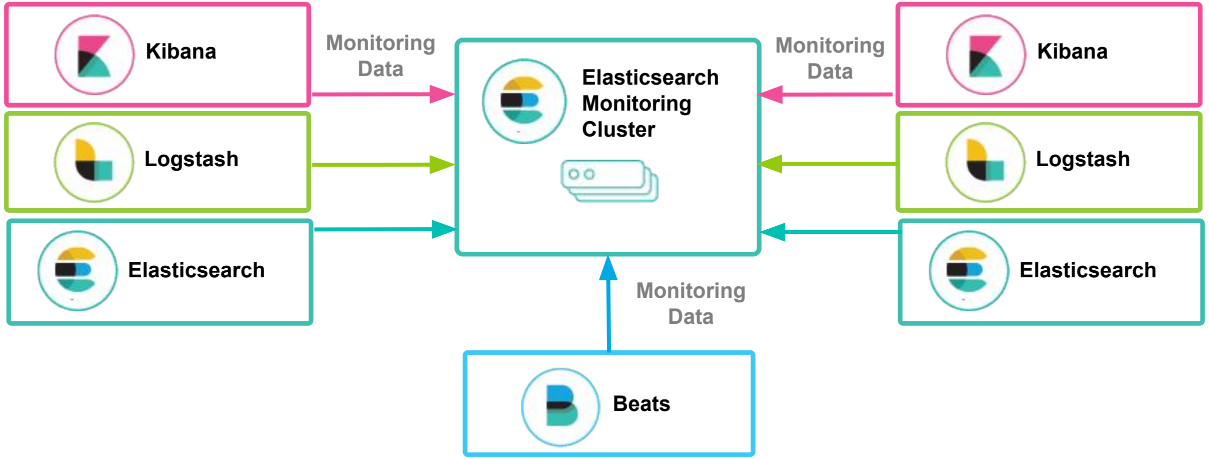

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Cloud Monitoring Cluster

Opt-in Elastic Cloud cluster for monitoring on-premise stack

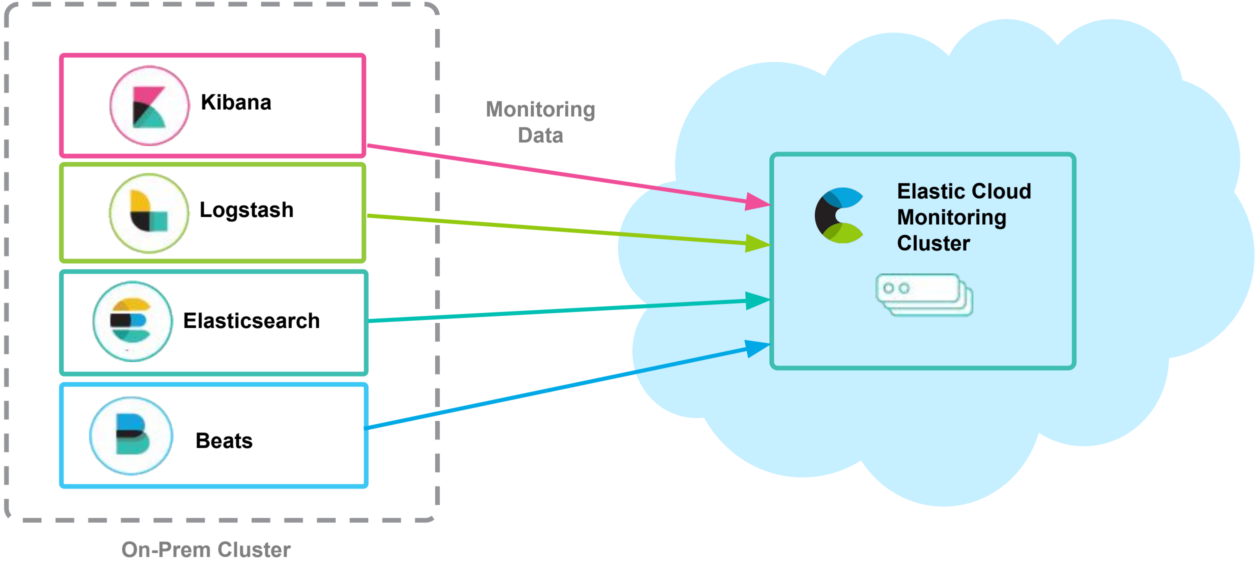

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Isolated Audit Logging Cluster

Maintain isolated audit logging cluster for increased security and compliance

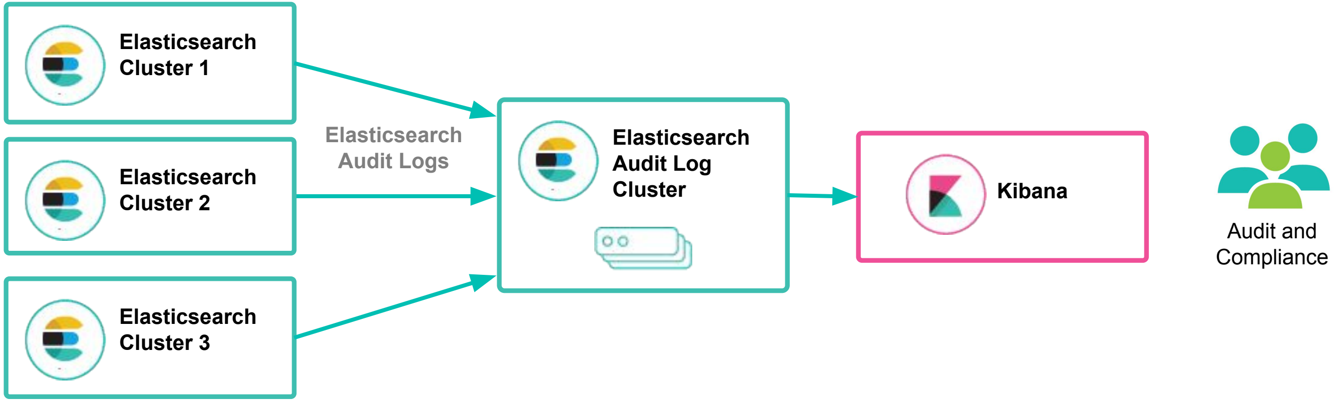

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Multi Data Center

Multiple Data Centers, Duplicate Data

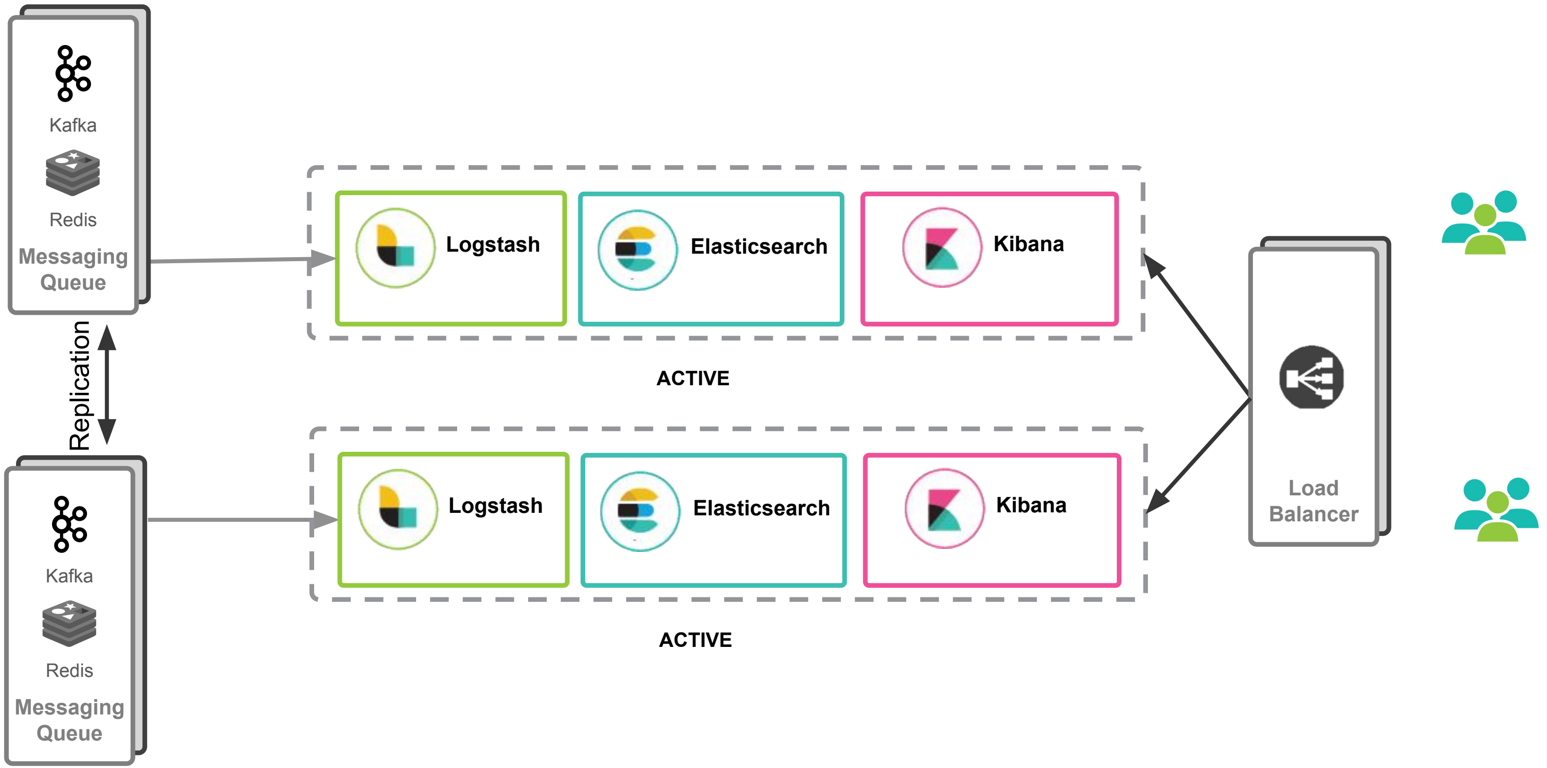

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Multi Data Centers with a Queue at Each DC

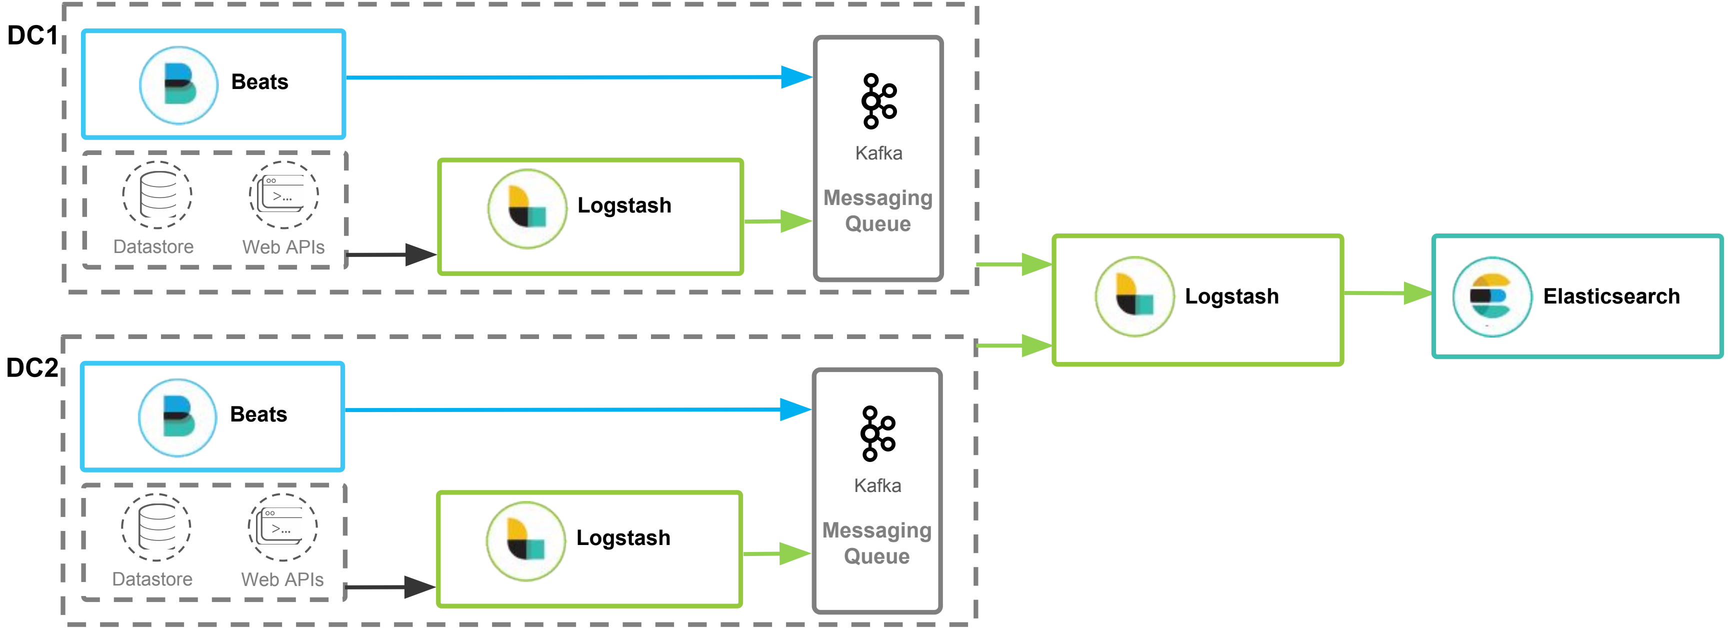

Multi Data Center, Distinct Data and Cross-Cluster Search

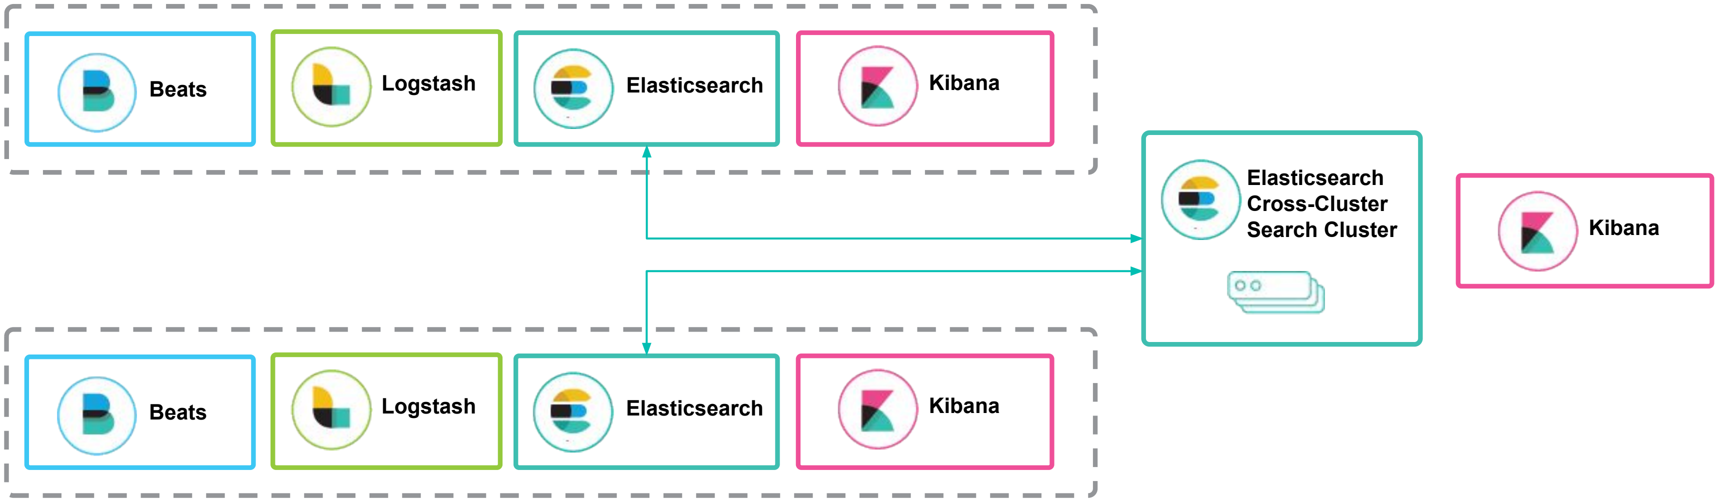

All product names, logos, and brands are property of their respective owners and are used only for identification purposes. This is not an endorsement.

Scaling Kibana

High Availability

Pair two coordinating nodes with two independent Kibana nodes

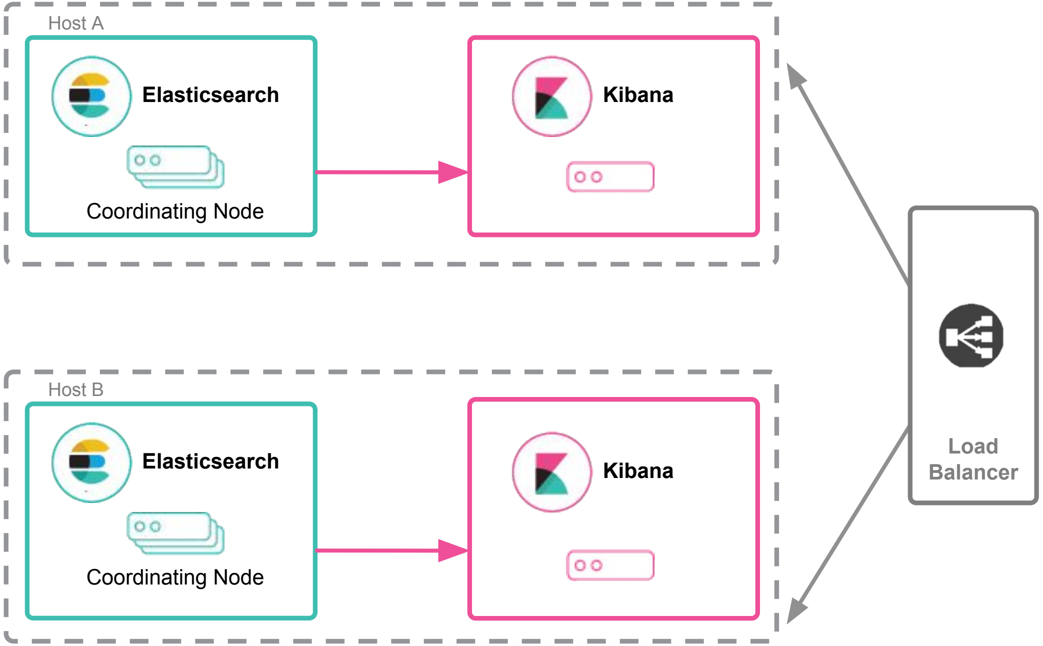

Separating Dashboards by Groups

Isolate user content by group in different Kibana instances

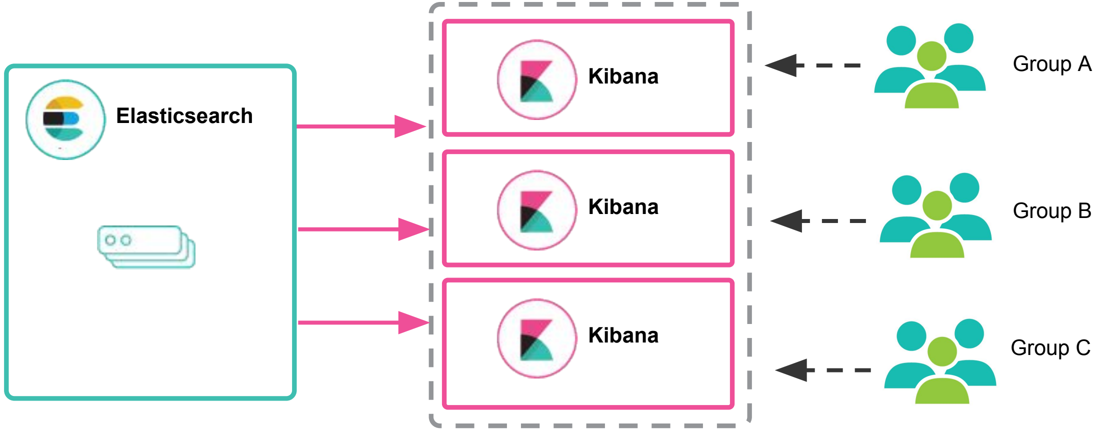

Questions?

elastic

Thank You

- **Web** : www.elastic.co
- **Products** : https://www.elastic.co/products
- **Forums** : https://discuss.elastic.co/
- **Community** : https://www.elastic.co/community/meetups
- **Twitter** : @elastic
- **Contact us** : www.elastic.co/contact