





Elasticsearch多集群管理最佳实践

罗厚付，技术总监
极限科技，2023/04/08


分享嘉宾




集群管理面临的挑战



集群管理常用的工具



客户实际运维案例和经验分享



最佳实践建议



集群管理面临的挑战

Elasticsearch版本发布历程




2015.10

基于Lucene 5.x

删除_shutdown API

Mapping changes



2016.10

基于Lucene 6.x

Profile API

Shrink API

Rollover API

Reindex

Ingest Node




2017.8

基于Lucene 7.x

Cross-cluster search

Removal of types

Load aware shard routing



2019.4


基于Lucene 8.x

内置OpenJDK支持快速使用

支持JSON格式日志跨集群复制 (CCR)

快照生命周期管理 (SLM)

索引生命周期管理 (ILM)



2022.2

基于Lucene 9.x

需要Java 17 才能运行

彻底删除_type

没有了High-level REST API

没有了 Low-level REST API



集群管理面临的挑战

安全性与权限

7.1开始

Security功能免费使用

A

8.0开始

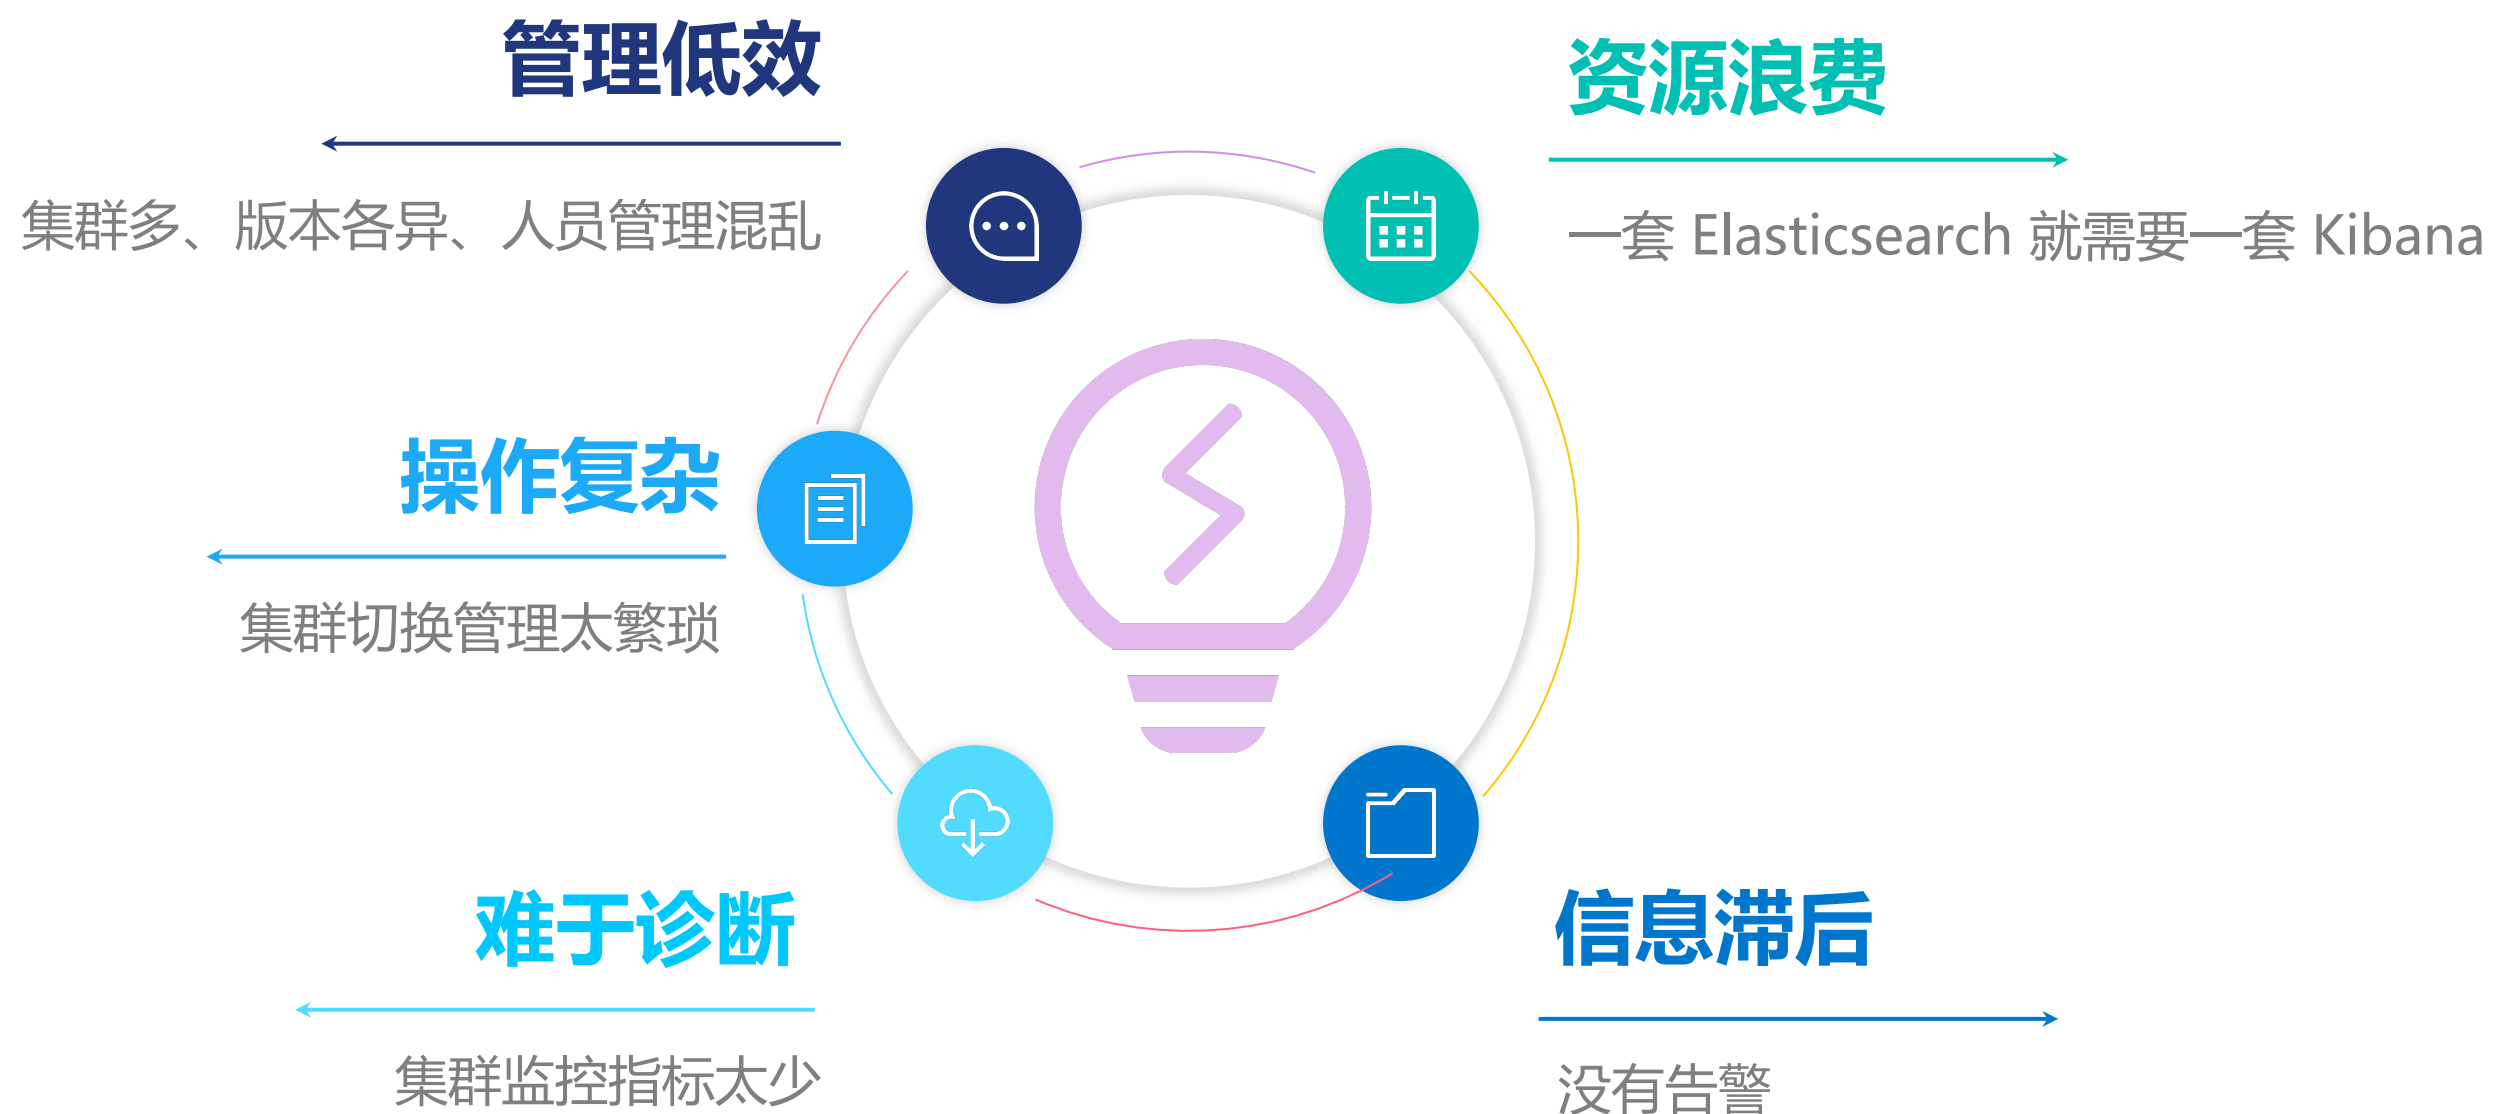
默认开启安全配置（三层安全）
对ES内置索引的保护加强

B



集群管理面临的挑战

运维挑战

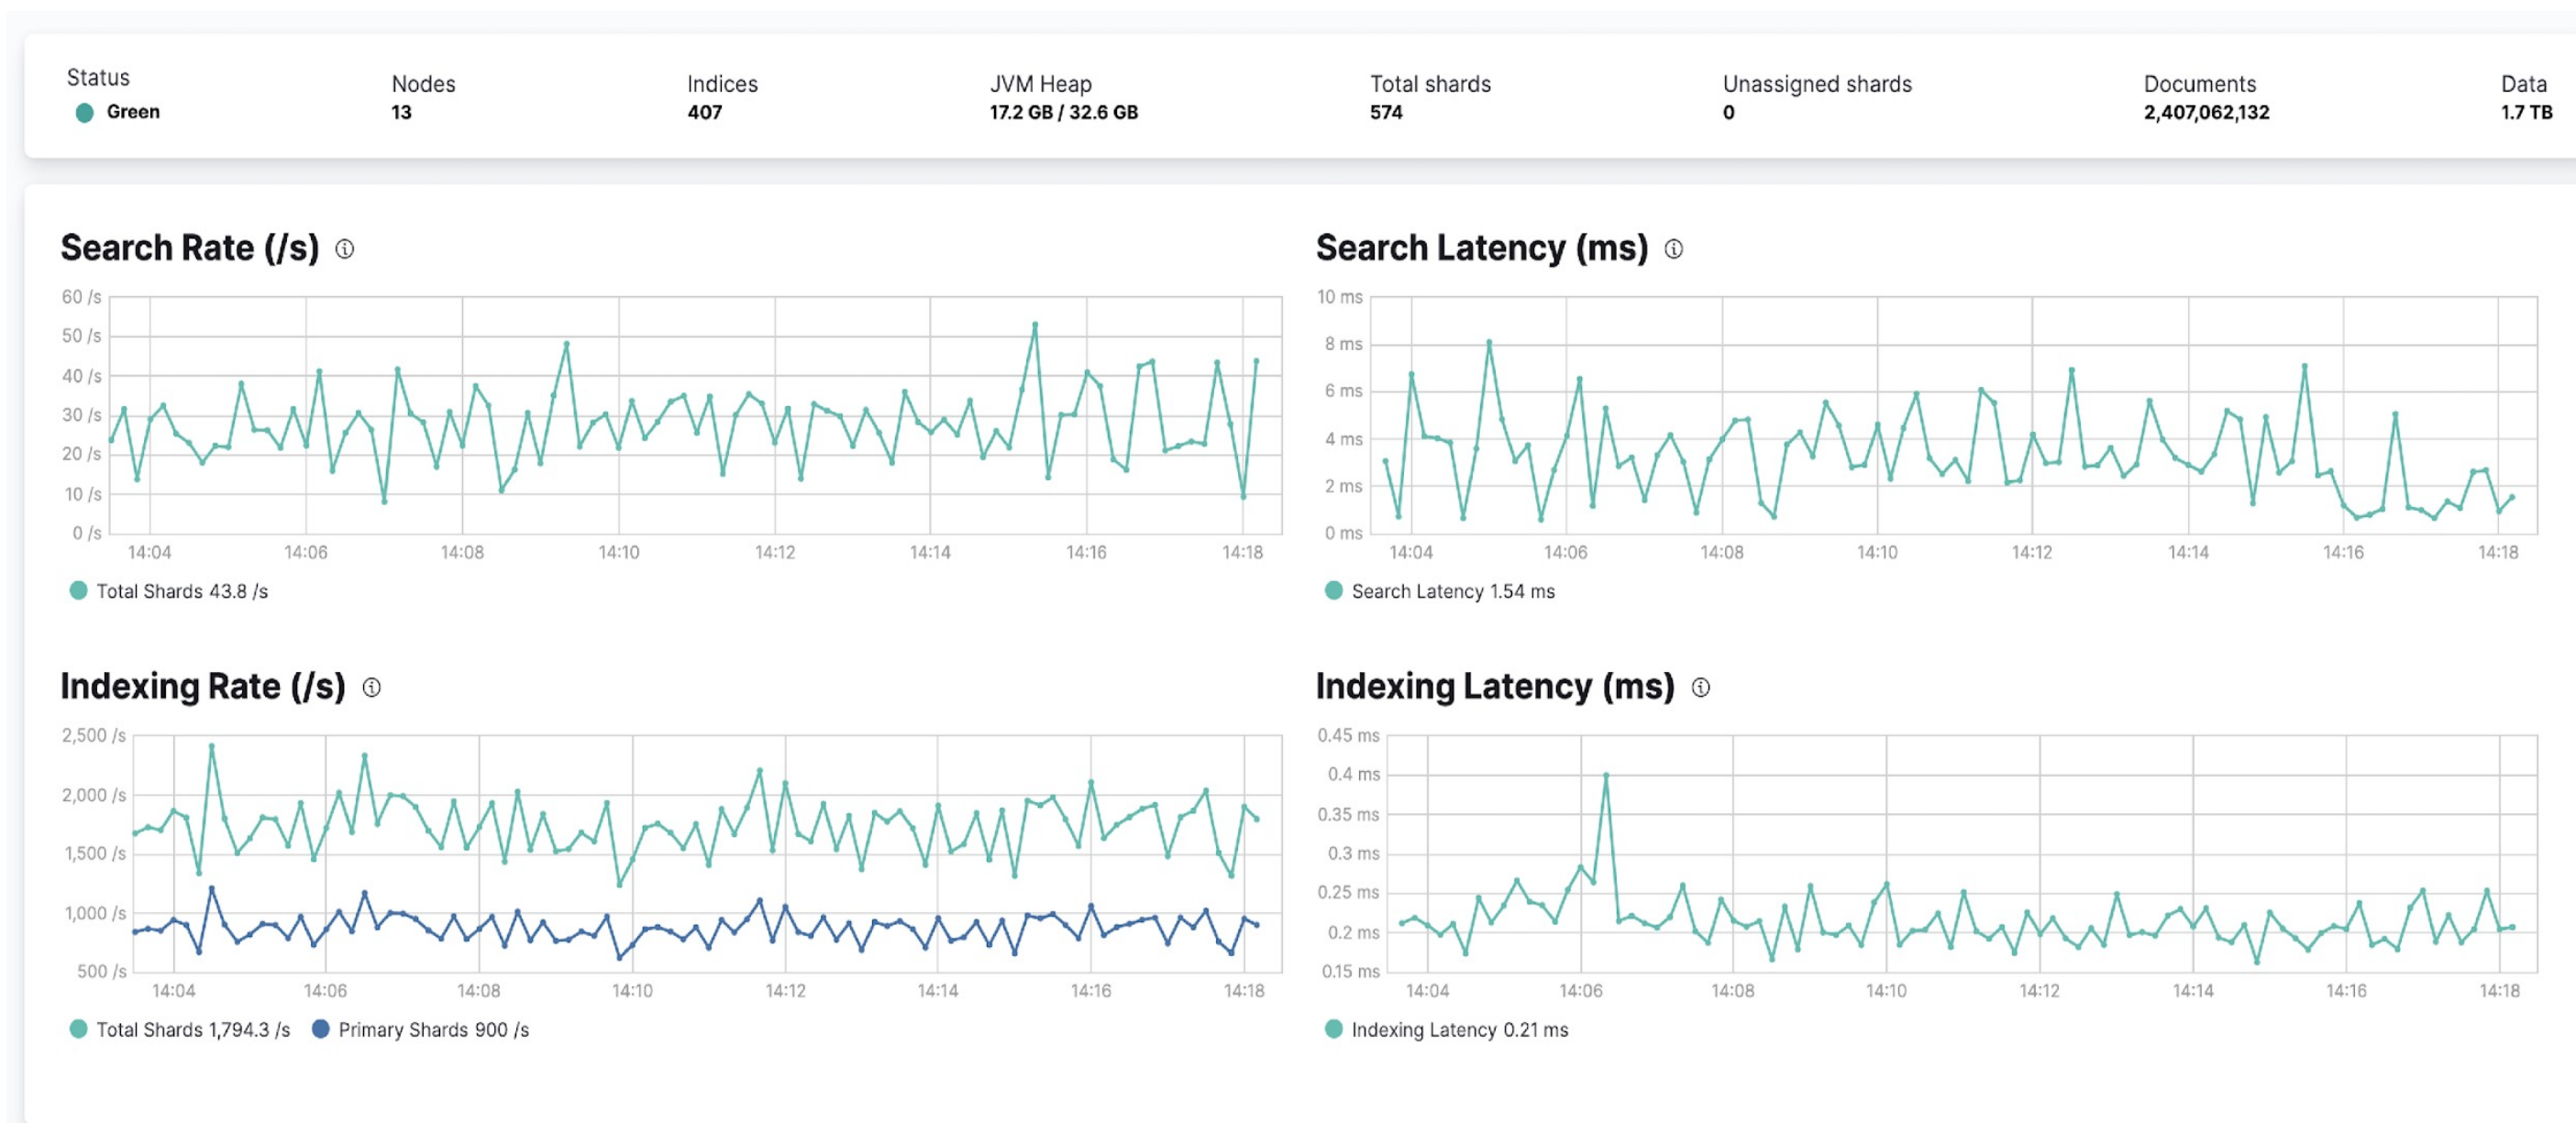




集群管理常用的工具



Kibana

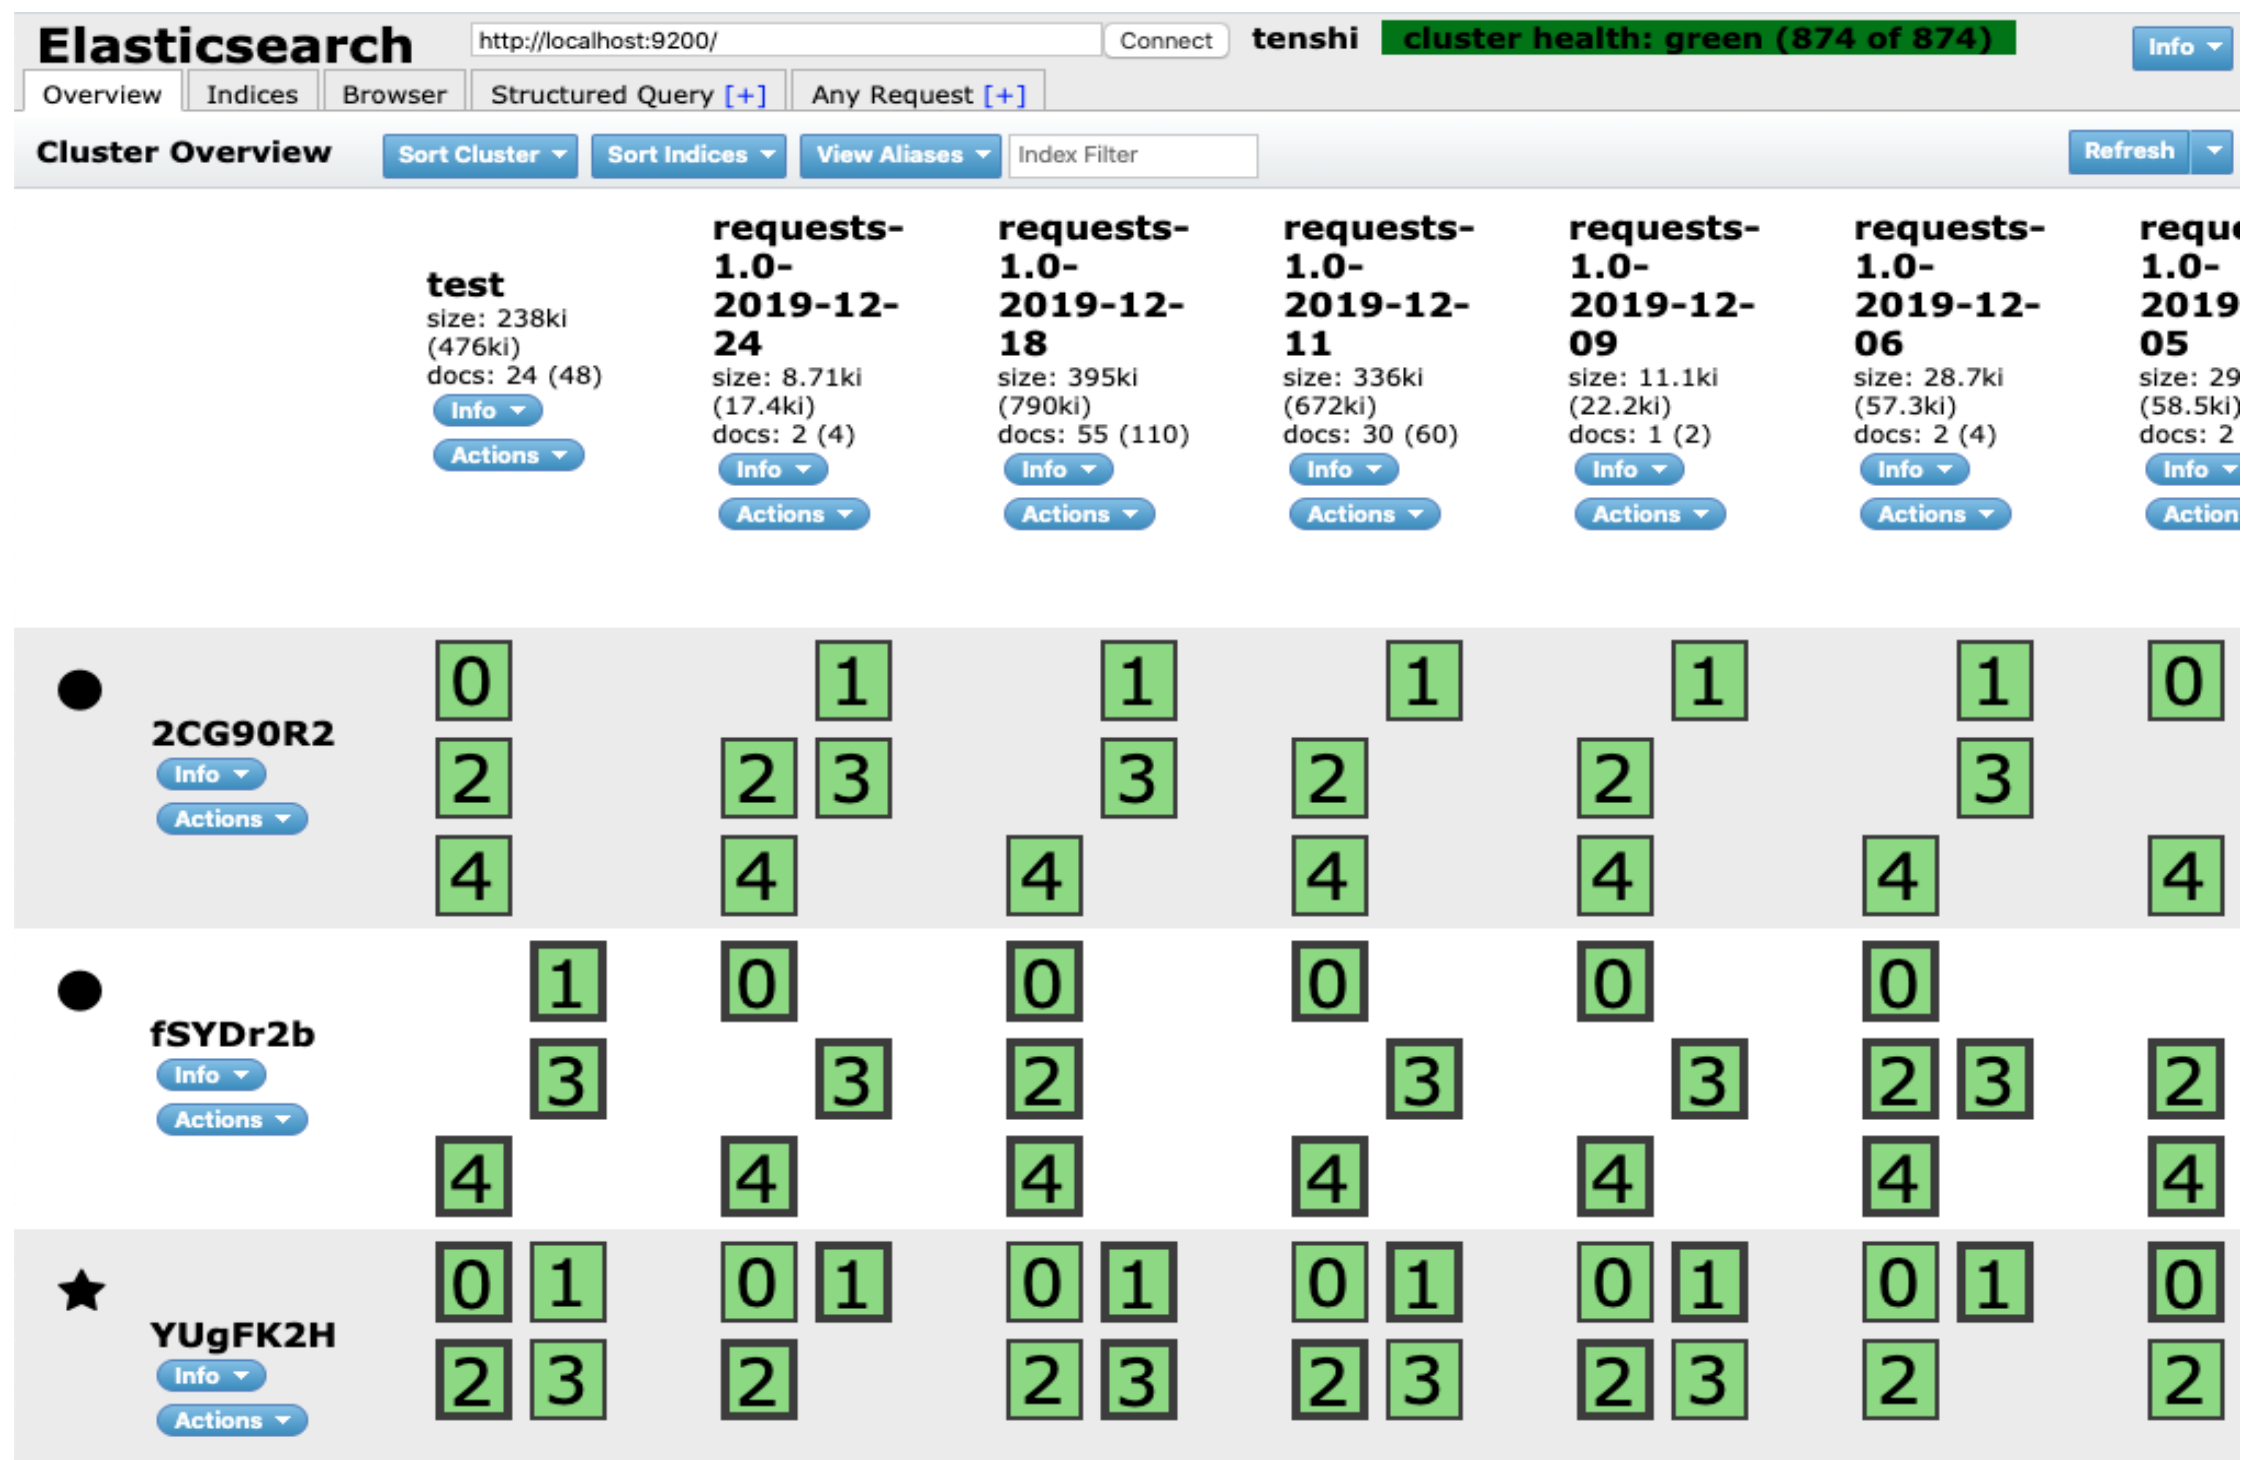




集群管理常用的工具



Elasticsearch-head





集群管理常用的工具



dejavu

URL for cluster goes here. e.g. `https://username:password@scalr.api.appbase.io` Headers AppName (aka inde... Connect

Connecting to Elasticsearch
To make sure you have enabled CORS settings for your Elasticsearch instance, add the following lines in the ES configuration file:

```
http.port: 9200
http.cors.allow-origin: http://localhost:1358,http://127.0.0.1:1358
http.cors.enabled: true
http.cors.allow-headers : X-Requested-With,X-Auth-Token,Content-Type,Content-Length,Authorization
http.cors.allow-credentials: true
```

If you are running Elasticsearch via Docker, use the following command:

```
docker run --name es -d -p 9200:9200 -e http.port=9200 -e http.cors.enabled=true -e http.cors.allow-origin=http://localhost:1358,http://127.0.0.1:1358 -e http.cors.allow-headers=X-Requested-With,X-Auth-Token,Content-Type,Content-Length,Authorization -e http.cors.allow-credentials=true docker.elastic.co/elasticsearch/elasticsearch-oss:6.5.2
```

gitxplore-live
asry3
2a-recon
MovieAppFinal
marketplacev1
shopect-upload-lo...
shopify-plugin-appb...
dejavu-test

Star 5,017 Watch Video Create your **Elasticsearch** in cloud with appbase.io



集群管理常用的工具

The screenshot displays the ElasticHQ interface for an Elasticsearch 6.4.3 cluster. The top navigation bar includes 'Clusters', 'Settings', 'Docs', and 'Star us on GitHub'. The main content area is divided into several sections:

- Summary:** Shows 12 Nodes and 81 Active Shards. A large '7 GB' size indicator is visible on the right.
- Nodes List:** A table listing nodes such as 'elasticsearch-data-2', 'elasticsearch-master-0', and 'elasticsearch-client-8694cb5b4d-stfx8'. The 'elasticsearch-master-0' node is checked as the master.
- Node Details:** Three detailed views for 'elasticsearch-data-0', 'elasticsearch-client-8694cb5b4d-d9nsk', and 'elasticsearch-data-2'. Each view includes a 'Summary' (IP Address, Node ID), 'File System' (Documents Deleted, Merge Rate, Disk space used), 'Actions' (Indexing, Search, Get, Refresh, Flush), and 'Cache' (Field Evictions) sections.
- Alerts:** A tooltip titled 'How this rule is calculated' is shown, explaining that high values indicate complex or inefficient queries, insufficient use of filters, insufficient RAM for caching, slow I/O or CPU. It lists keys: 'indices.search.query_time_in_millis / indices.search.query_total' and 'Keys: 406700 / 2423'.
- Cache Size Table:** A table at the bottom right showing cache sizes for various nodes, including 'logstash-2018.12.13' (4.1 GB) and 'logstash-2018.12.06' (3.8 GB).



集群管理常用的工具

The screenshot shows the ELASTICSEARCHCOMRADE interface. At the top, the 'Cluster Stats' section is highlighted with a purple box, displaying the following metrics:

- Indices: 8
- Primary Shards: 75
- Replica Shards: 70
- Docs: 148,773,916
- Nodes: 6

Below the stats, the 'Index Cell' view is shown, displaying a grid of index shards across five nodes. The 'Node Cells' section on the left is highlighted with a red box, showing node details for node-1 through node-5, including metrics like HEAP, DISK, CPU, and LOAD. The 'Shards' section in the center is highlighted with a yellow box, showing the distribution of shards across nodes. A red arrow points to a specific shard in the 'Shards' section, labeled 'Relocation progress'.

Node	dummy_201902	dummy_201903	dummy_201904	dummy_201905	dummy_201906	dummy_201907
node-1	6 3 5	1 4 7	1 0 3	1 6 3 5	3 7 8 9	0 1 2
node-2	3 9 8	7 8 5	4 7 6	3 9 8	1 3 9 5	3 6 9 5
node-3	0 2 8 7	0 6 1 3	5 6 8	2 8 7	0 6 8	8 7 9
node-4	1 7 6 9	3 9 6	1 2 9	7 6 9	2 7 6	7 4 6 8
node-5	5 1 4	5 2 9	0 3 9 5	0 5 1 4	5 1 4	5 1 3



集群管理常用的工具



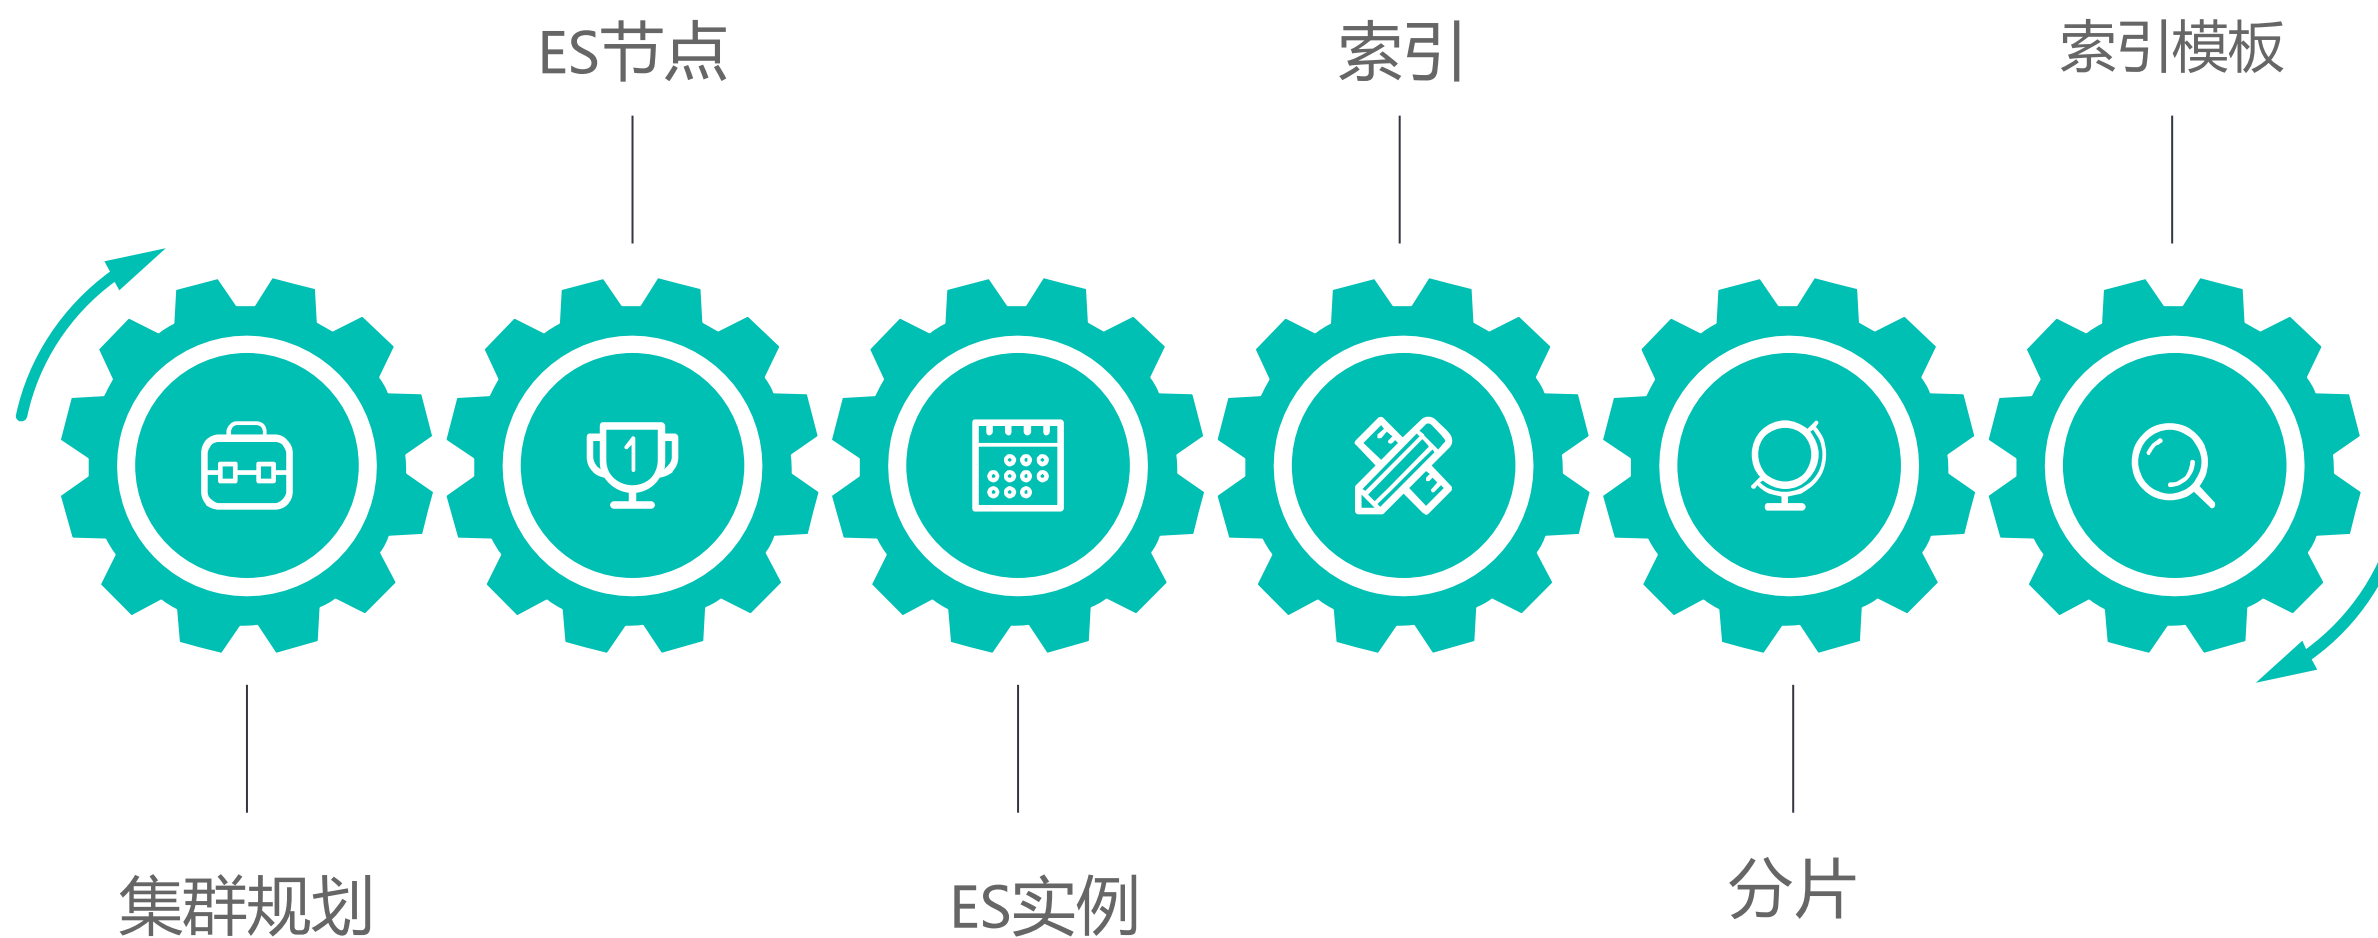
Cerebro

The screenshot displays the Cerebro web interface for an Elasticsearch cluster. At the top, there are navigation tabs for 'overview', 'nodes', 'rest', and 'more'. The current view is 'overview', showing a summary of the cluster: 'elasticsearch' with 12 nodes, 17 indices, 162 shards, 129,538,388 docs, and 106.93GB of data. Below this, there are filters for indices (closed: 0, .special: 2) and nodes. The main part of the interface is a table with columns for indices and rows for nodes. Each cell in the table shows a grid of colored boxes representing shards. The indices listed are logstash-2018.12.04 through logstash-2018.12.08. The nodes listed are elasticsearch-data-0, elasticsearch-data-1, and elasticsearch-data-2. The 'load' column for each node shows a red bar indicating the current load level.

	logstash-2018.12.04 shards: 5 * 2 docs: 7,652,404 size: 3.24GB	logstash-2018.12.05 shards: 5 * 2 docs: 9,849,405 size: 4.26GB	logstash-2018.12.06 shards: 5 * 2 docs: 8,433,022 size: 3.53GB	logstash-2018.12.07 shards: 5 * 2 docs: 11,460,206 size: 4.40GB	logstash-2018.12.08 shards: 5 * 2 docs: 5,951,811 size: 2.27GB
elasticsearch-data-0 10.233.72.29	0 1 4	0 2 3 4	0 1 4	0 1 2	0 2 3
elasticsearch-data-1 10.233.65.121	1 2 3 4	0 1 2	1 2 3 4	2 3 4	1 2 3 4
elasticsearch-data-2 10.233.71.115	0 2 3	1 3 4	0 2 3	0 1 3 4	0 1 4

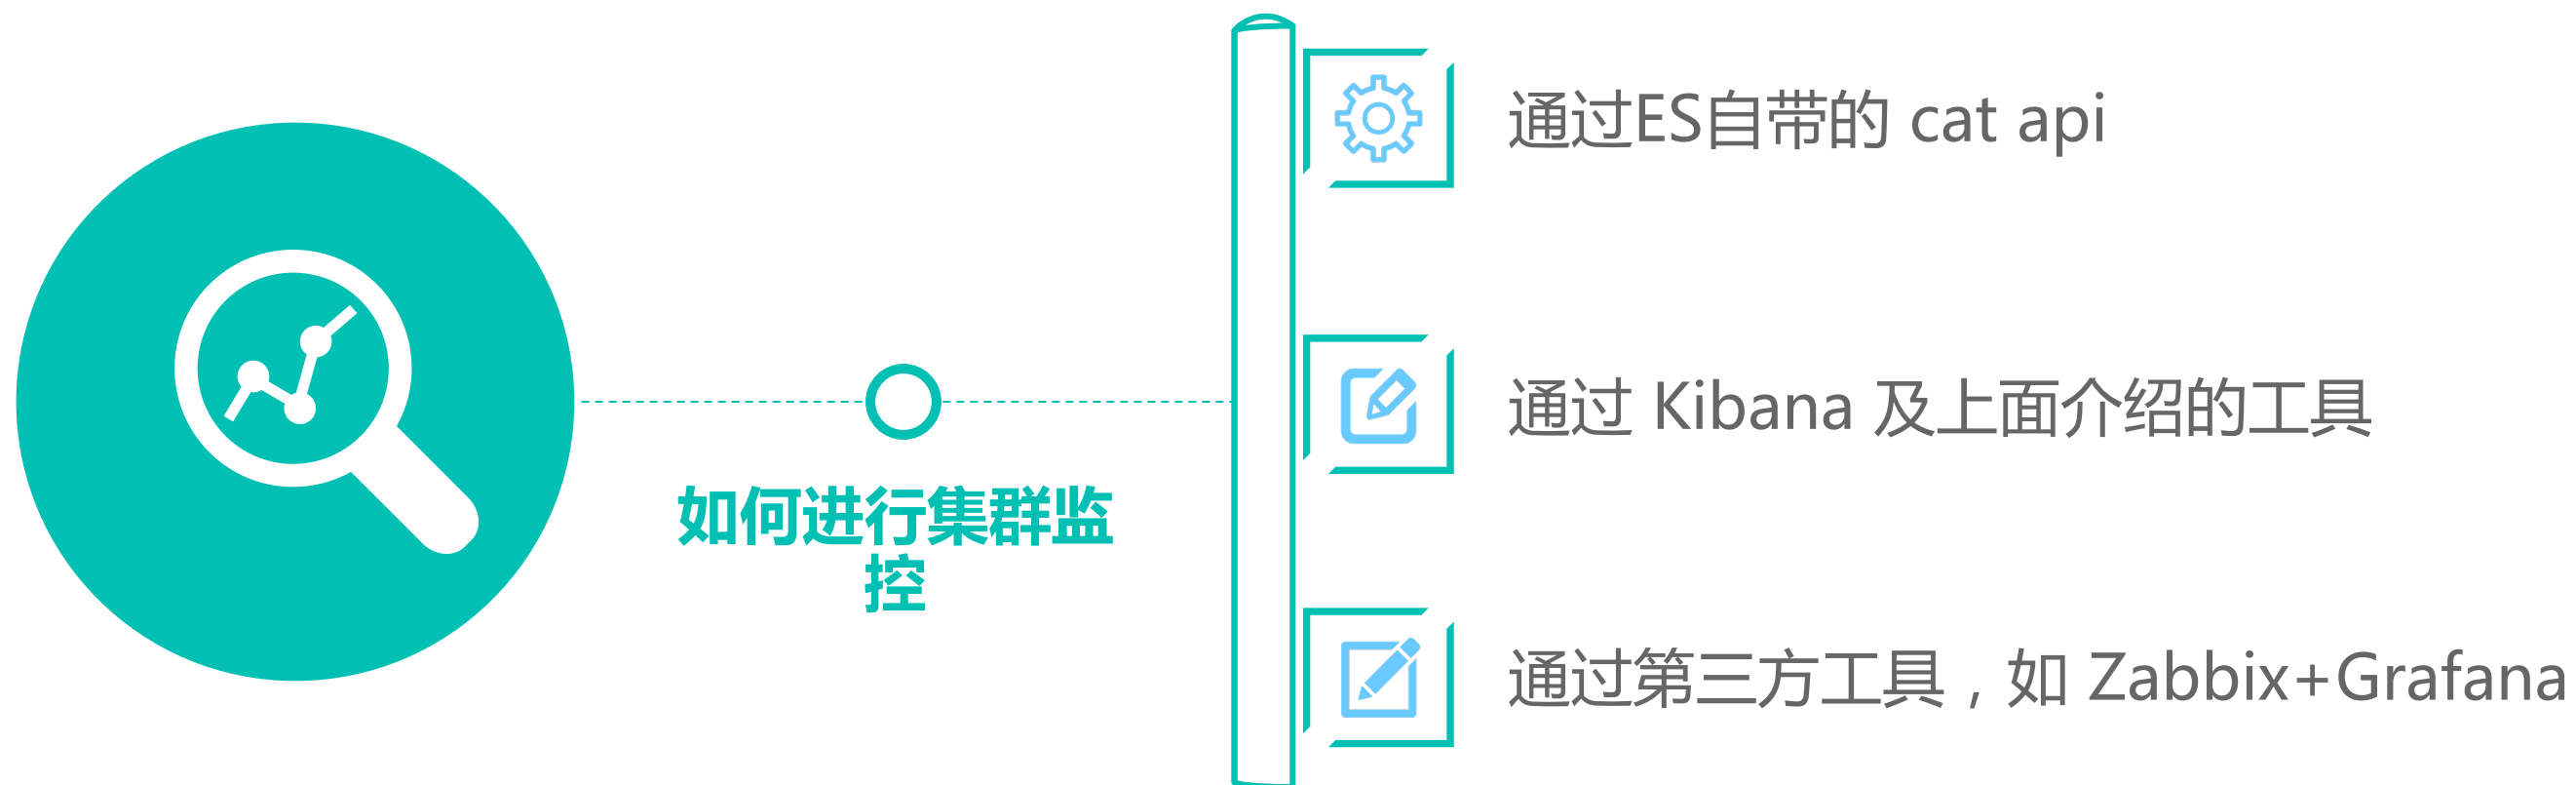
客户实际运维案例和经验分享

如何诊断集群问题





客户实际运维案例和经验分享





最佳实践建议

INFINI Console

工作台
平台管理
平台概览
监控报表
集群动态
数据管理
数据分析
告警管理
开发工具
数据工具
资源管理
系统管理

Enterprise (0.8.2_SNAPS...)

集群总数: 15
主机总数: 3
节点总数: 11
已用存储: 1.5_{TB}

Clusters Nodes Indices Hosts

Filters Type keyword to search Search

green X

Cluster	Version	Host	Provider	Region
es-v7102	7.10.2	192.168.3.185:9200	腾讯云	东京
es-v5616	5.6.16	192.168.3.188:9208	私有部署	昆明

Cluster: es-v7102
Nodes: 1 Indices: 111 Shards: 113 Docs: 21.86B Disk Usage: 1.2TB
Mem Used: 2.6GB
Recent Cluster Status (14 day)
Disk Usage: 22% JVM Heap: 9% 0 Indexing/s 5.5 Search/s

Cluster: es-v5616
Nodes: 1 Indices: 6 Shards: 30 Docs: 12.33M Disk Usage: 762.1GB
Mem Used: 1.8GB
Recent Cluster Status (14 day)
Disk Usage: 80% JVM Heap: 46% 0 Indexing/s 0 Search/s

Sort By: Relevance

HEALTH STATUS
Filter health status
yellow 9
unavailable 4
green 2


DISTRIBUTION
Filter distribution
elasticsearch 11
opensearch 2
easysearch 1

VERSION
Filter version
7.10.2 7
1.0.0 2
2.4.6 1
2.5.0 1

1-2 of 2 items < 1 > 10 条/页



最佳实践建议



[首页](#) / [数据工具](#) / [数据迁移](#) / [cgjt6n53q95pjfu9q17g](#) / [任务详情](#)

- [工作台](#)
- [平台管理](#)
- [数据管理](#)
- [数据分析](#)
- [告警管理](#)
- [开发工具](#)
- [数据工具](#)
- [资源管理](#)
- [系统管理](#)

Task Detail Back

Source Cluster	easysearch-v1.0	Workers	181-Gateway
Target Cluster	es-v7102		185-Gateway
Creator	medcl	Time Window	00:00~23:59
Total Indices	2	Task Logging	View
Completed Indices	2		

Indices 刷新

	easysearch-v1.0			>>	es-v7102			Progress	Error Partitions
	Index	Type	Documents		Index	Type	Documents		
+ test-es-log			12,326,146	>>	newlog-1	_doc	12,326,146	100%	0
- test-10-22.10.21			983,440	>>	new-log2	_doc	983,440	100%	0

Start Time: 2023-04-01 14:47:12 Duration: 00:01:42 Partition: 32 [View Log](#)

End Time: 2023-04-01 14:48:54 Step: 1m Completed Partition: 32



Progress

Enterprise (0.8.2_SNAPS... >

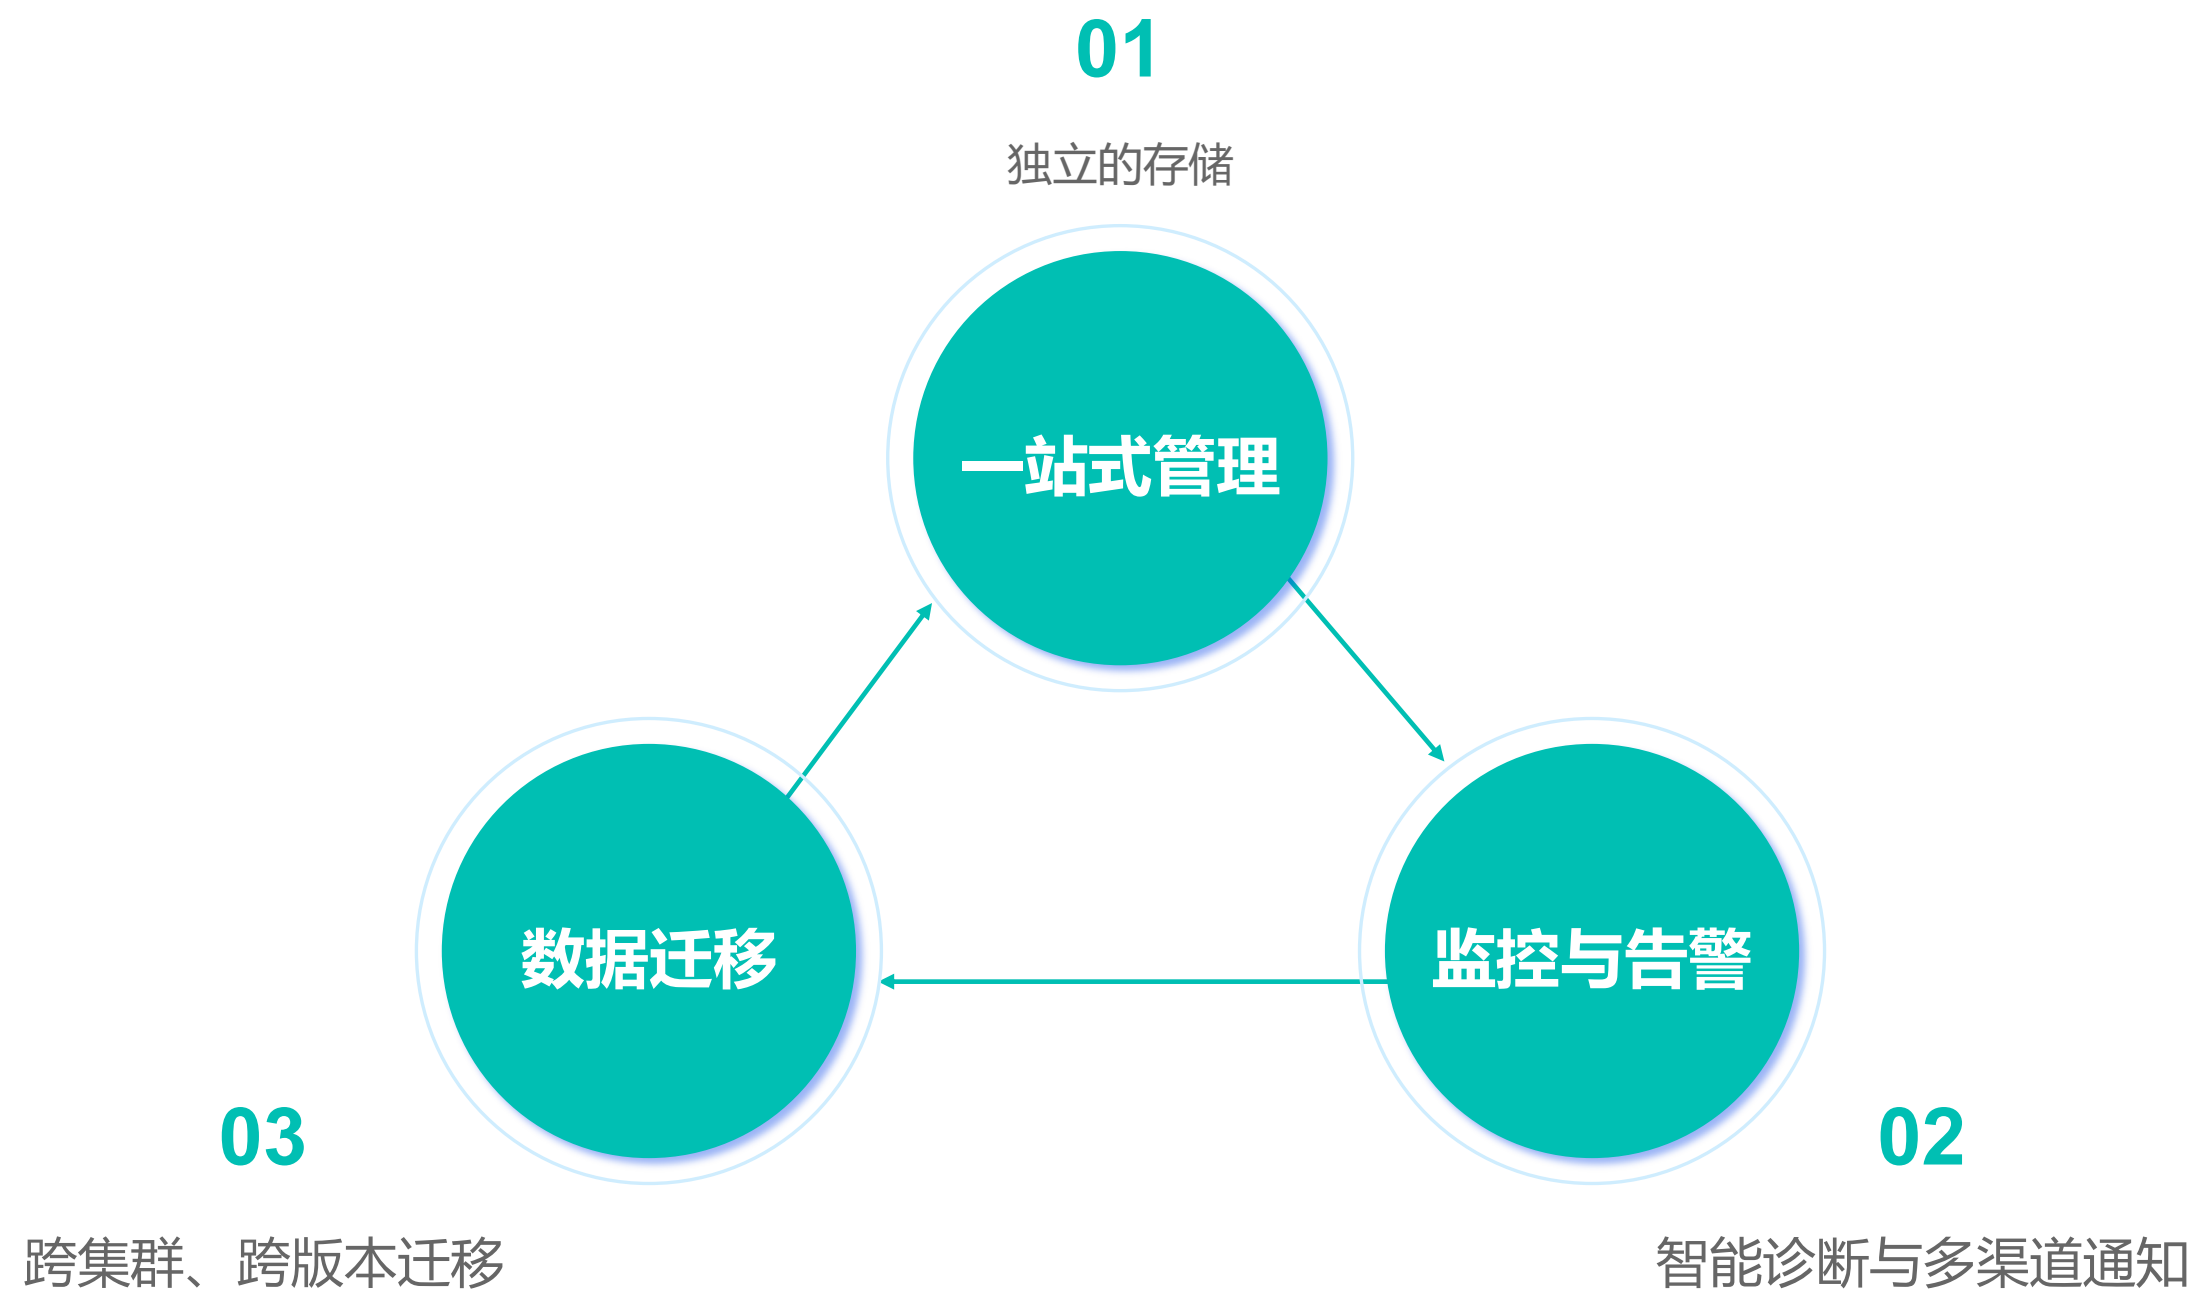




体验 INIFNI Console



总结





感谢观看



专业、垂直、纯粹的 Elastic 开源技术交流社区

<https://elasticsearch.cn/>