elastic

Community Conference 2021

Troubleshooting your Elasticsearch cluster like a Support Engineer

Janko Strassburg, Imma Valls
Sr. Support Engineers, Elastic
@jankopueh, @eyeveebee

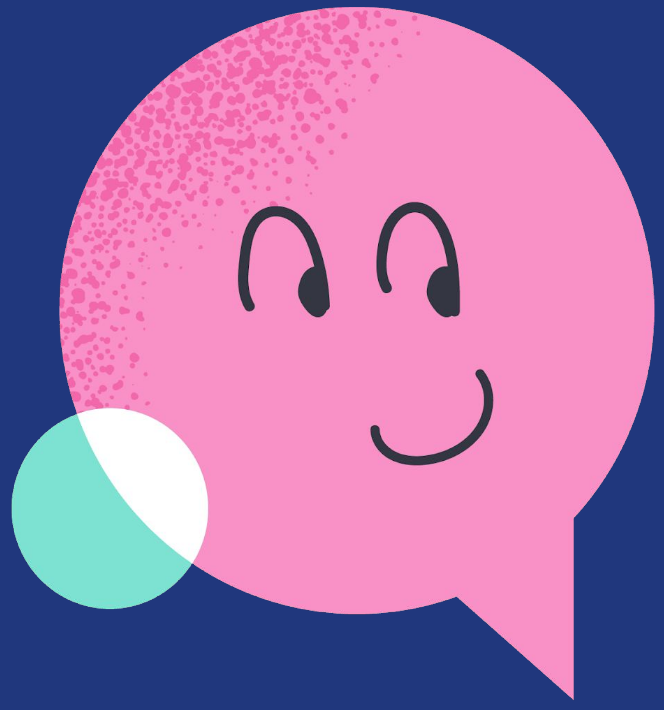

Cluster down!

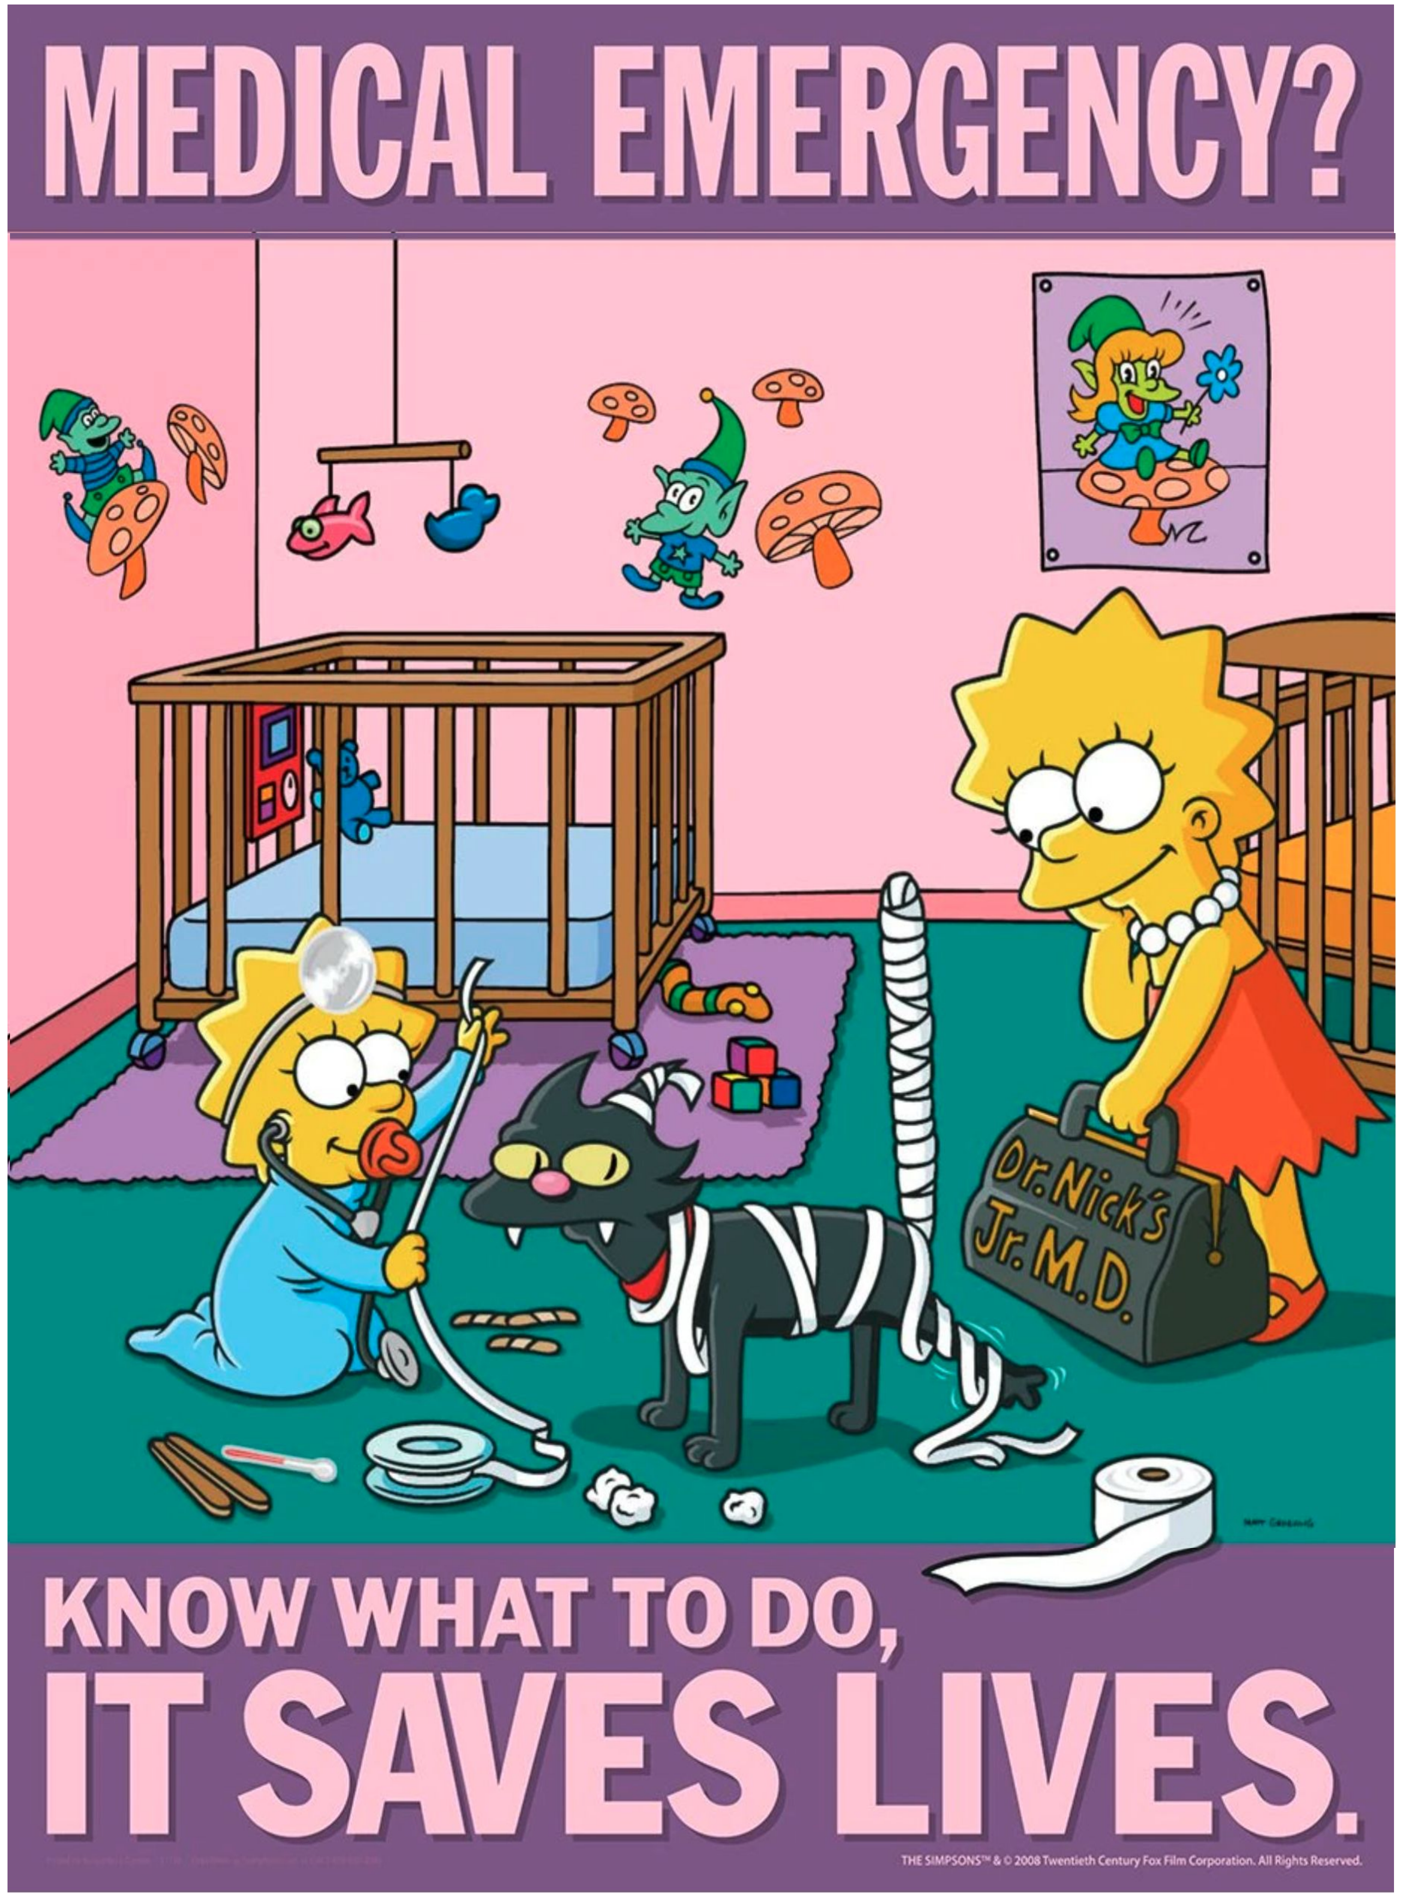
MEDICAL EMERGENCY?
Dr. Nick's
Jr. M.D.
KNOW WHAT TO DO,
IT SAVES LIVES.
THE SIMPSONS™ & © 2008 Twentieth Century Fox Film Corporation. All Rights Reserved.

https://safetyposter.com/products/simpsons-safety-poster-medical-emergency-know-what-to-do

How can we approach troubleshooting?

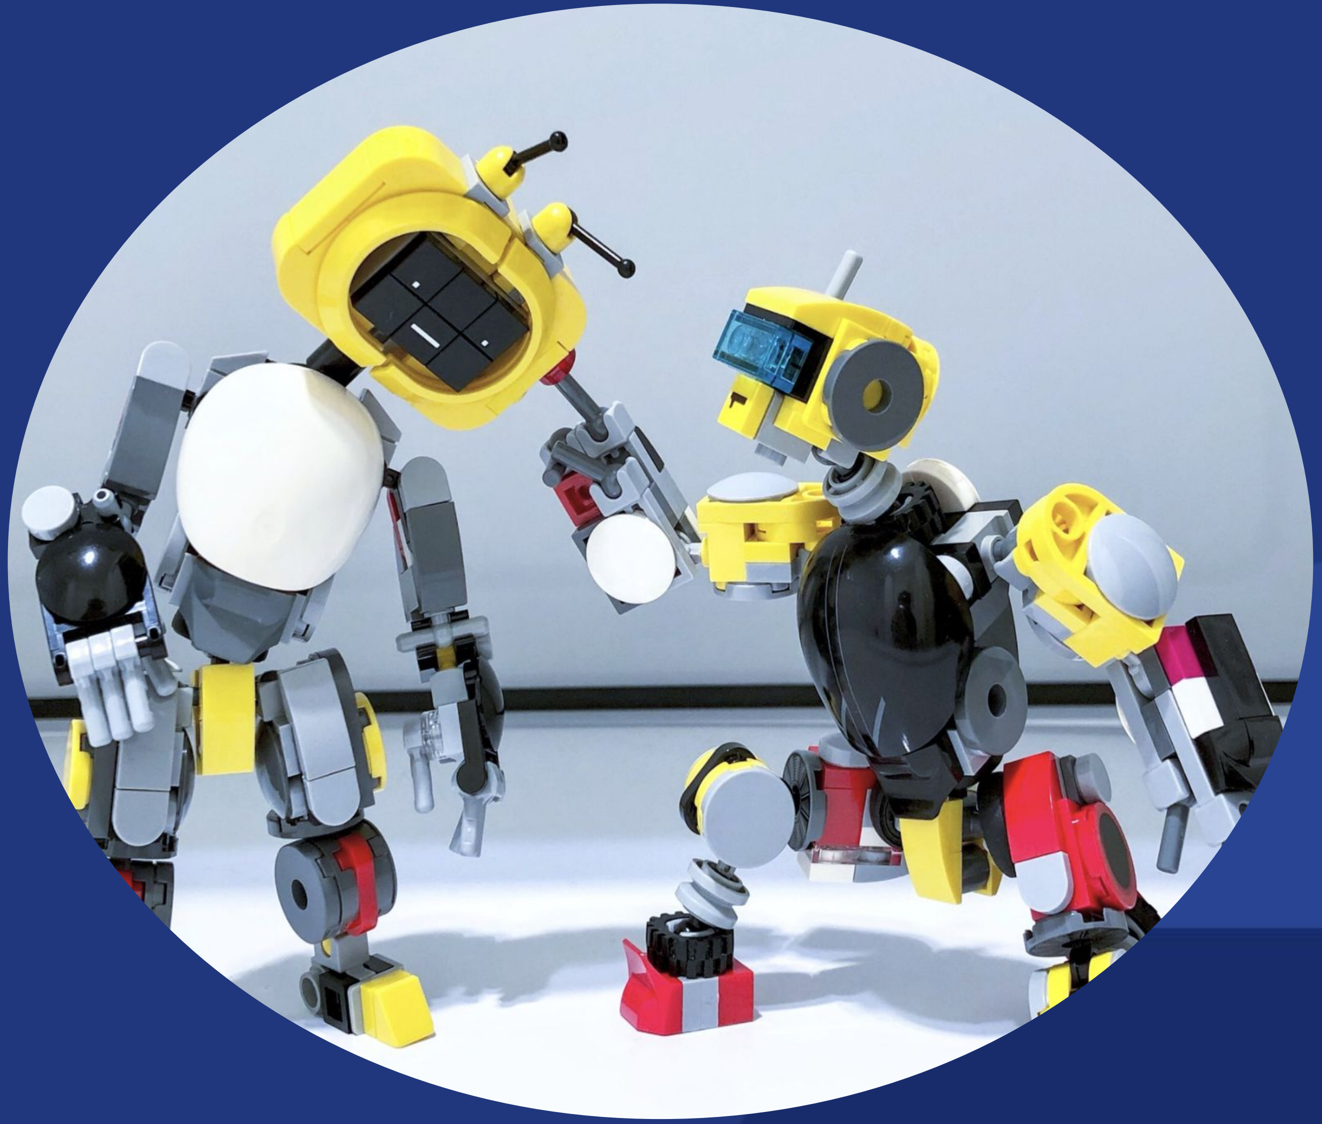

The hospital Emergency Room model

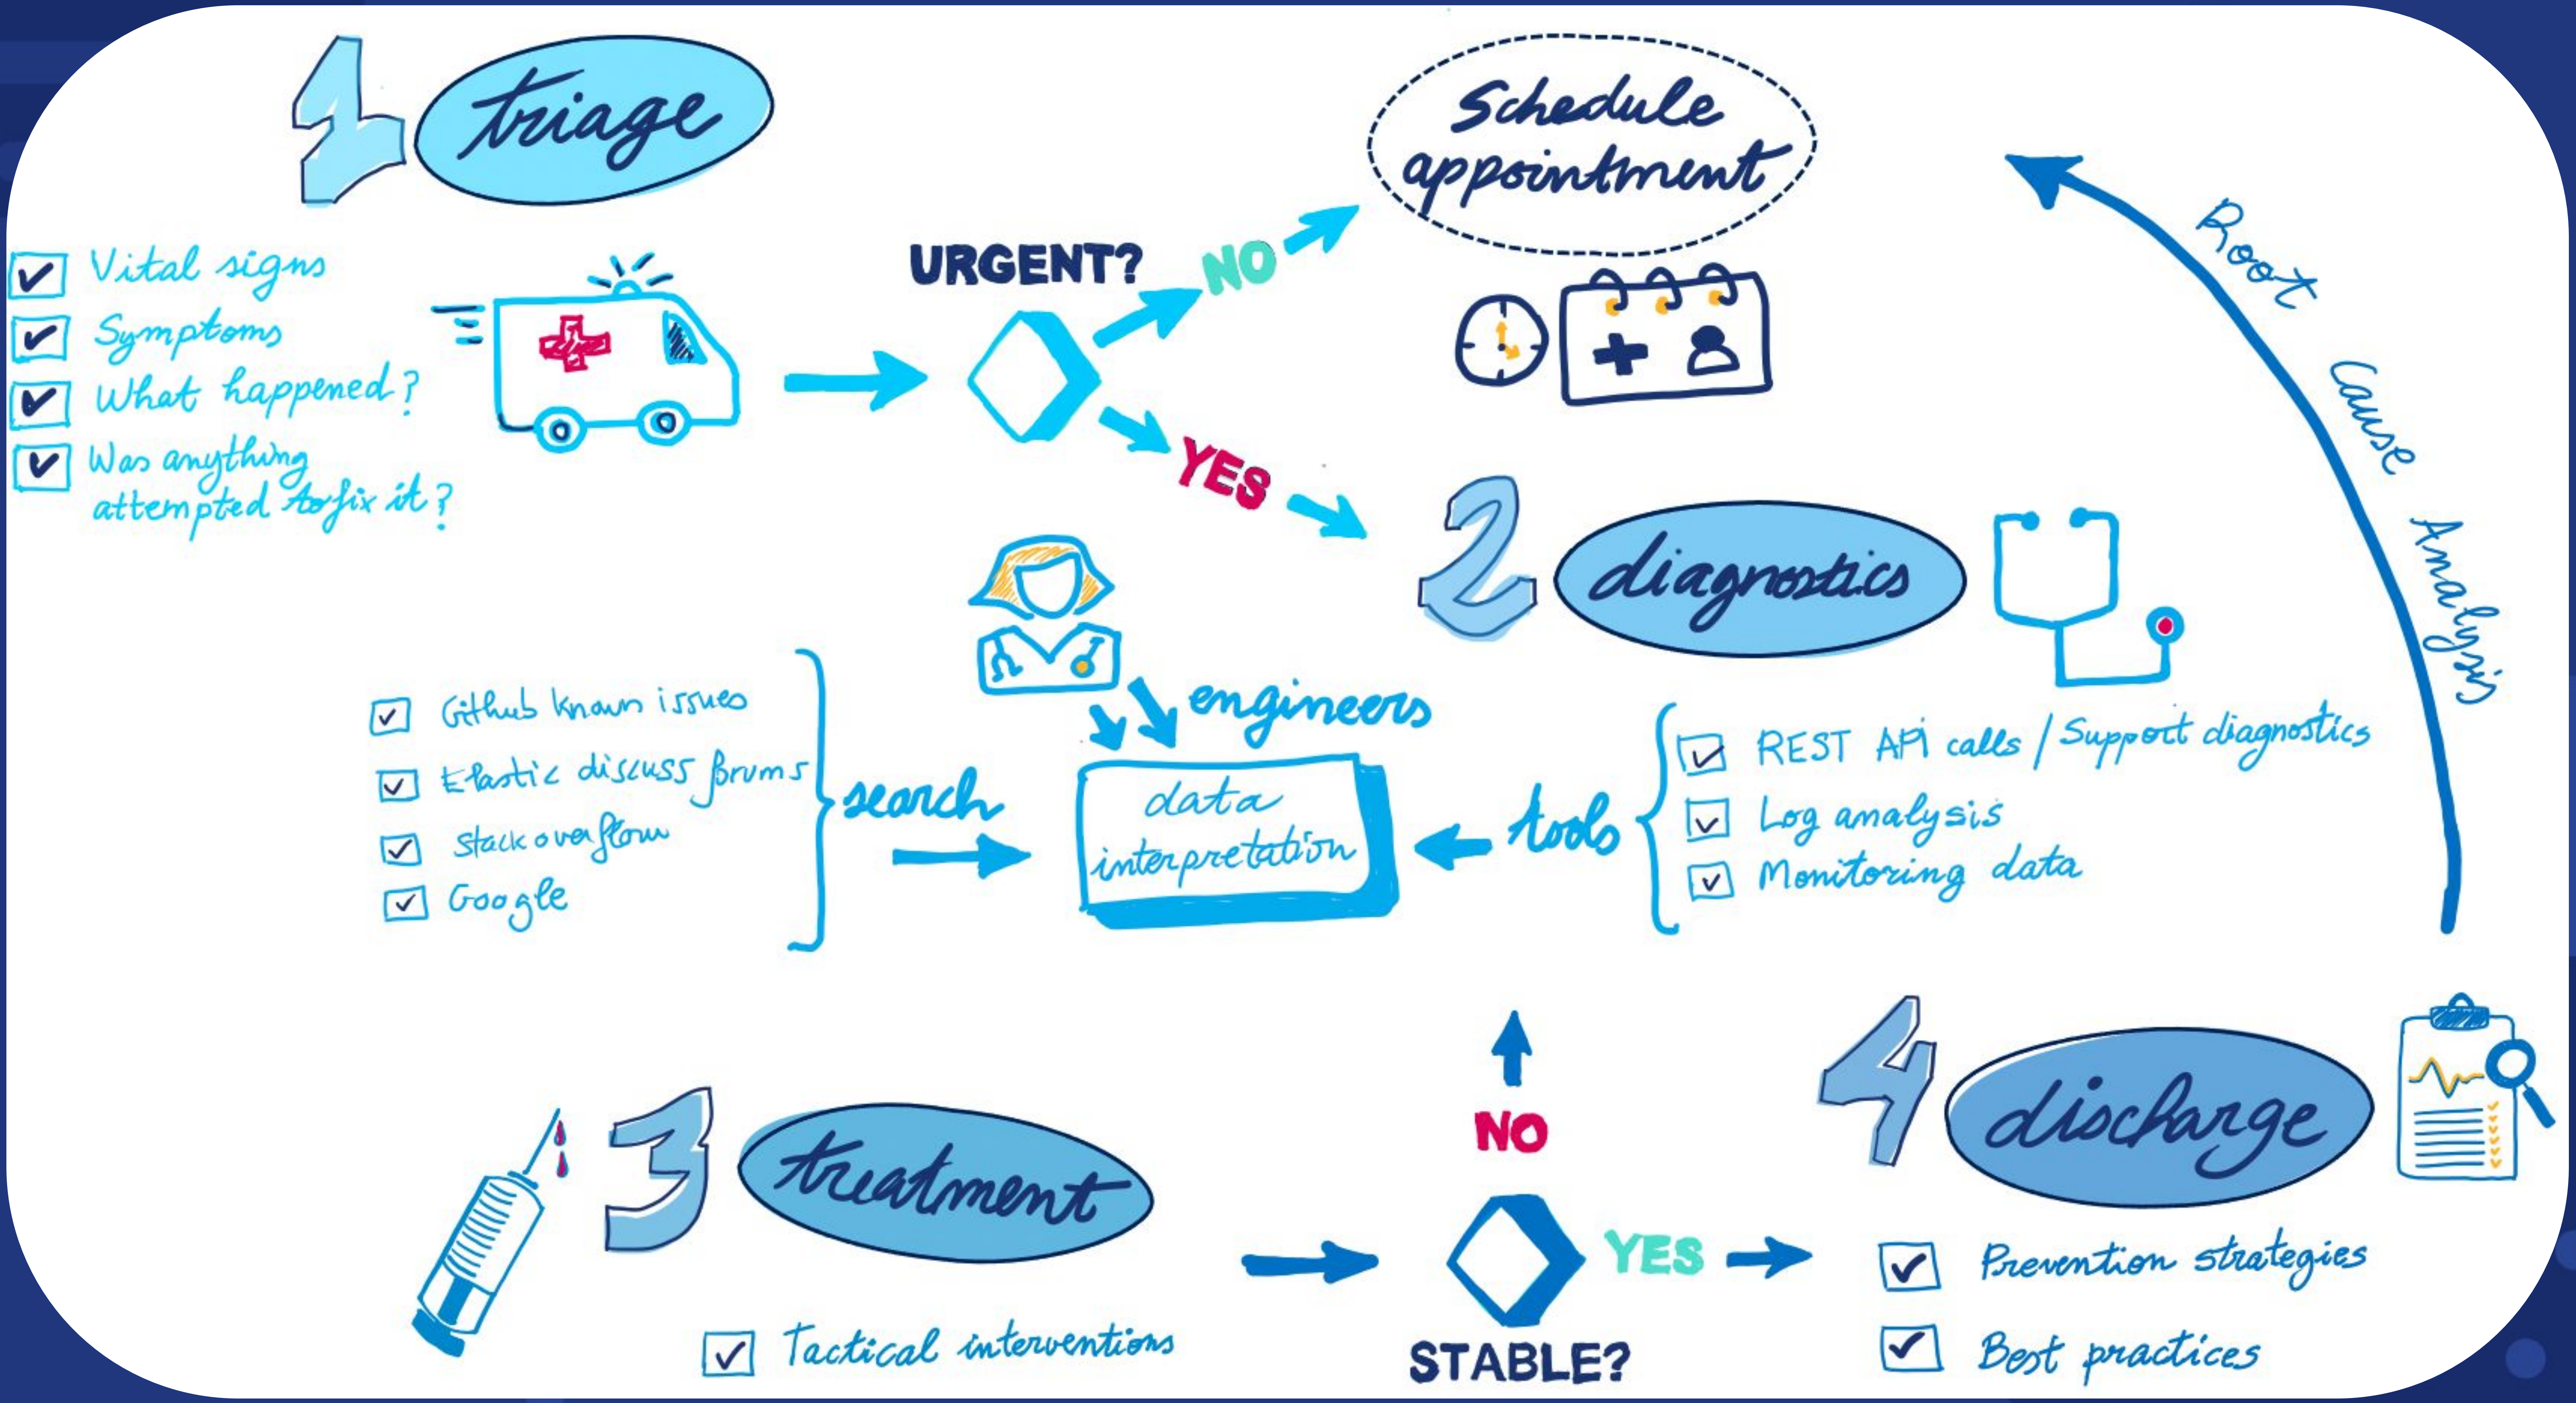

☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix it?

☑ Github known issues
☑ Elastic discuss forums
☑ stack overflow
☑ Google

☑ REST API calls / Support diagnostics
☑ Log analysis
☑ Monitoring data

☑ Tactical interventions

☑ Prevention strategies
☑ Best practices

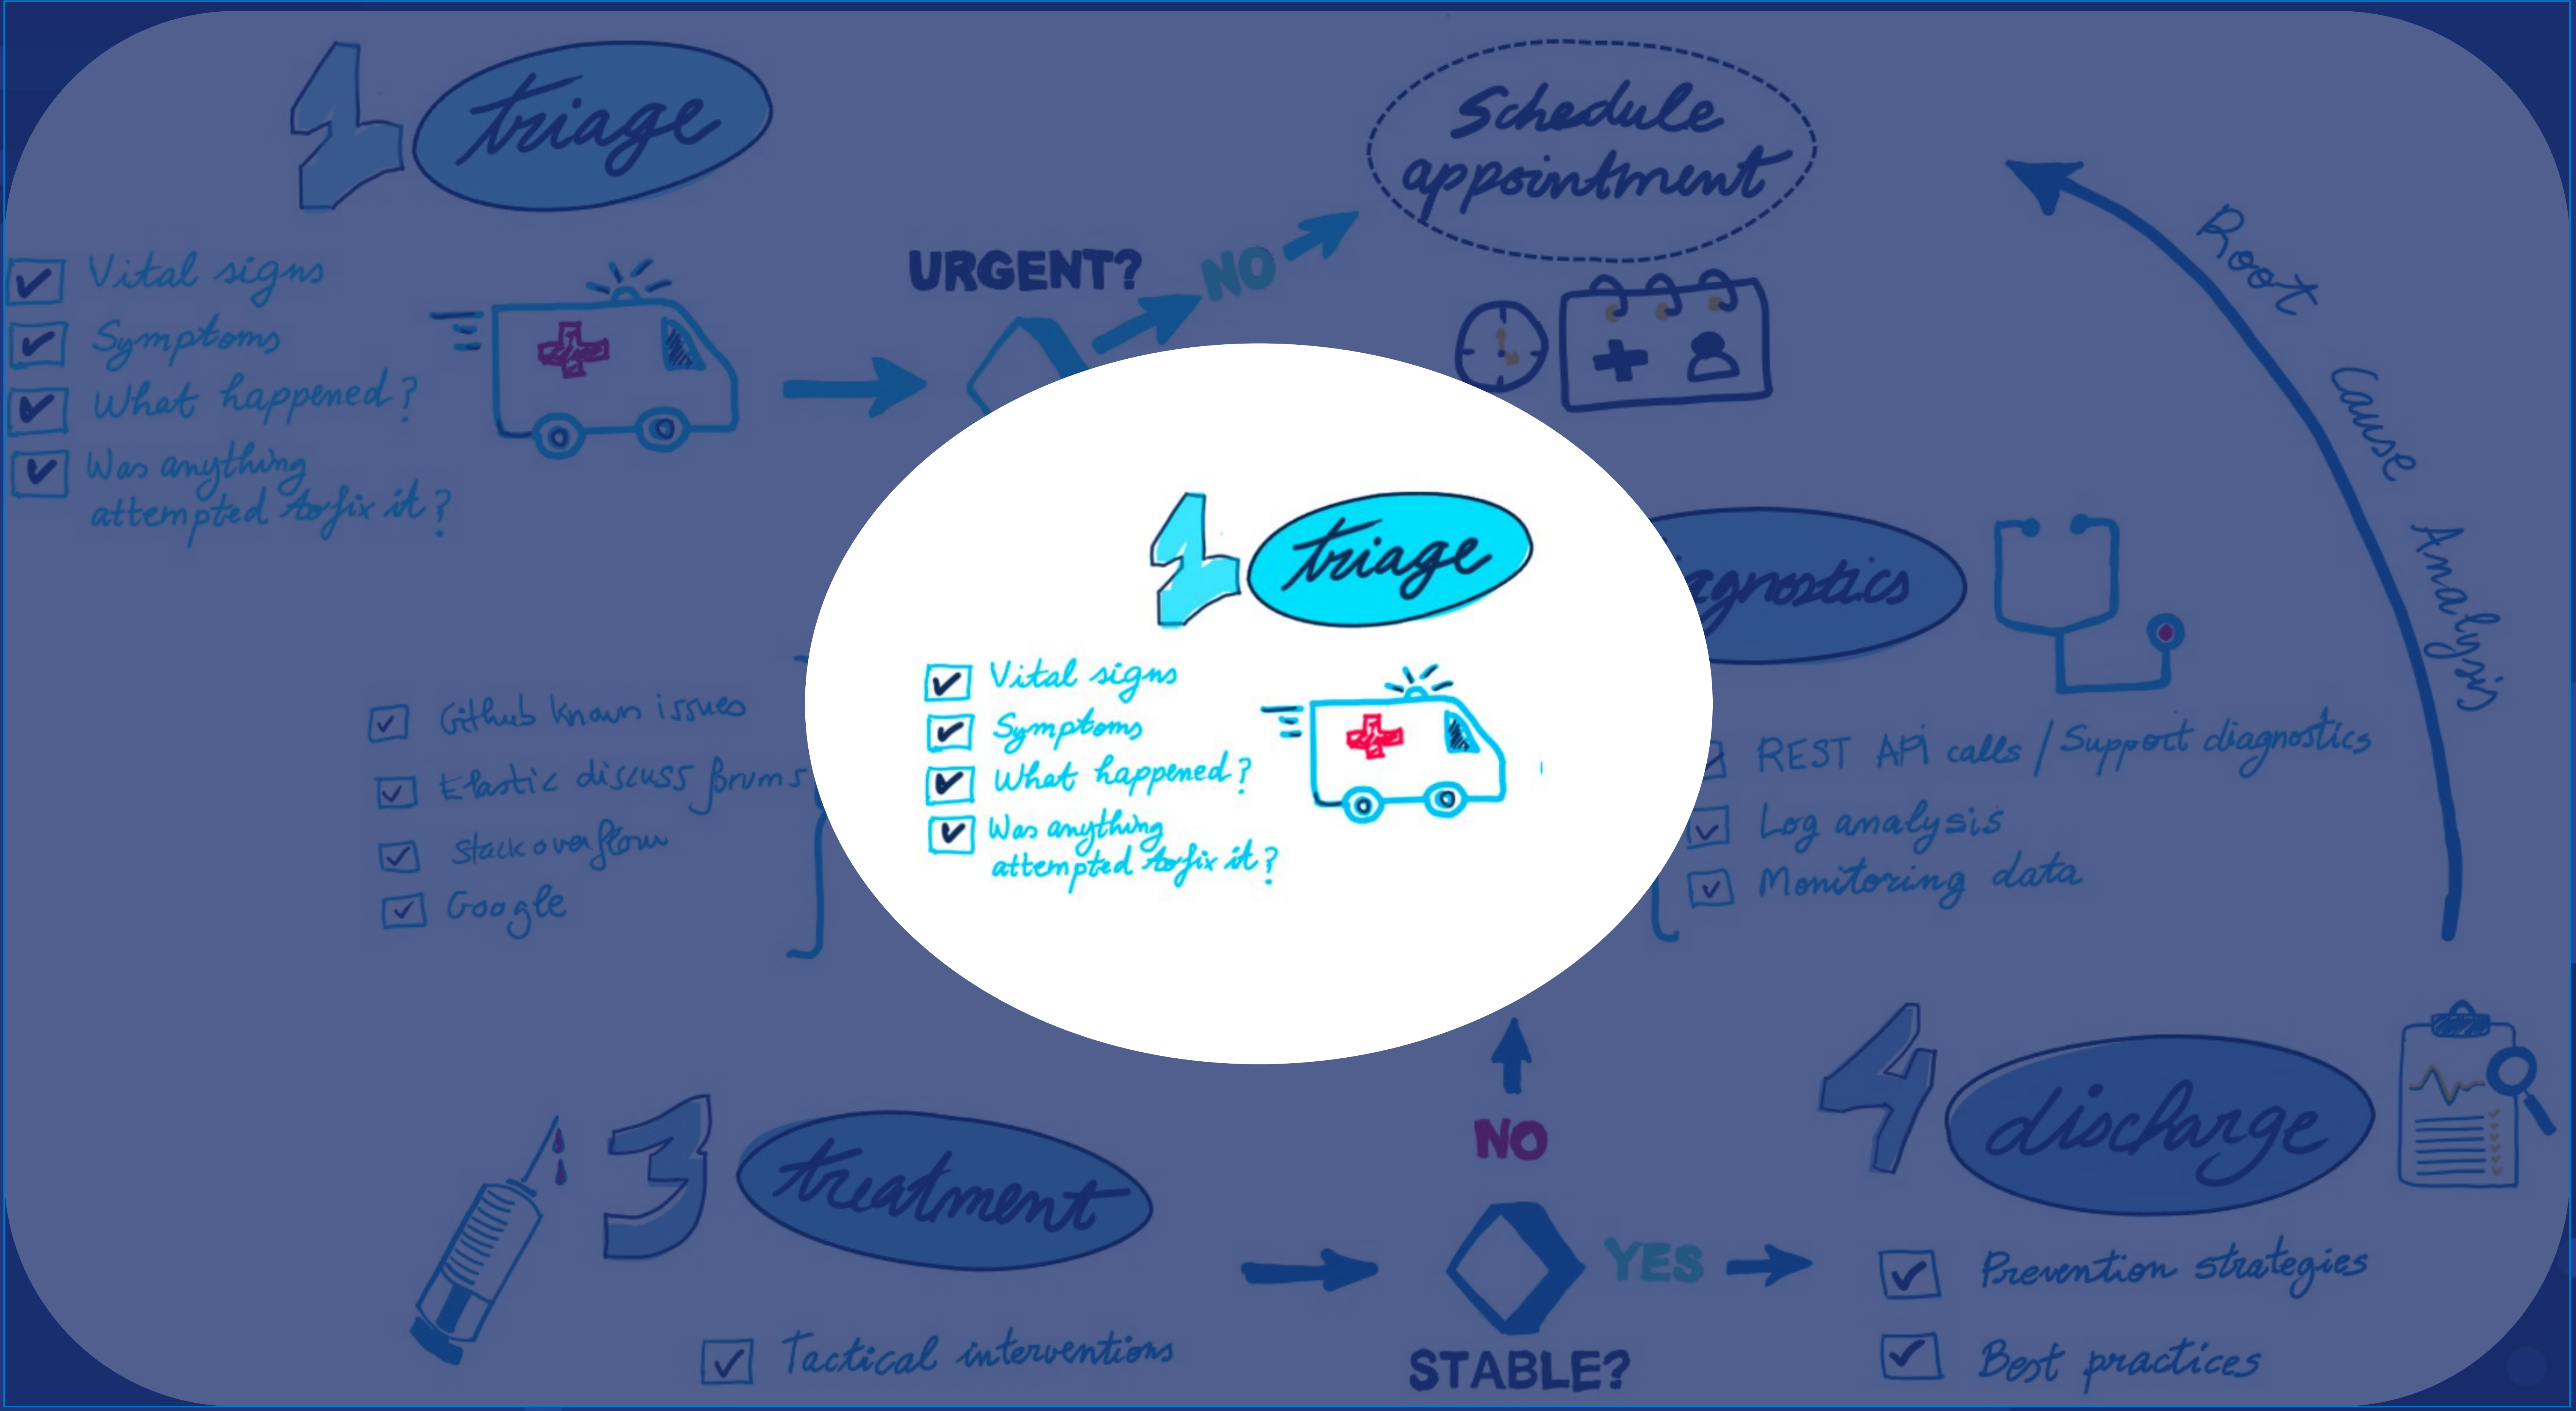
1 triage
☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix it?
URGENT?
NO
Schedule appointment
Root Cause Analysis
gnostics
☑ Github known issues
☑ Elastic discuss forums
☑ stackoverflow
☑ Google
☑ REST API calls / Support diagnostics
☑ Log analysis
☑ Monitoring data
3 treatment
☑ Tactical interventions
NO
STABLE?
YES
4 discharge
☑ Prevention strategies
☑ Best practices

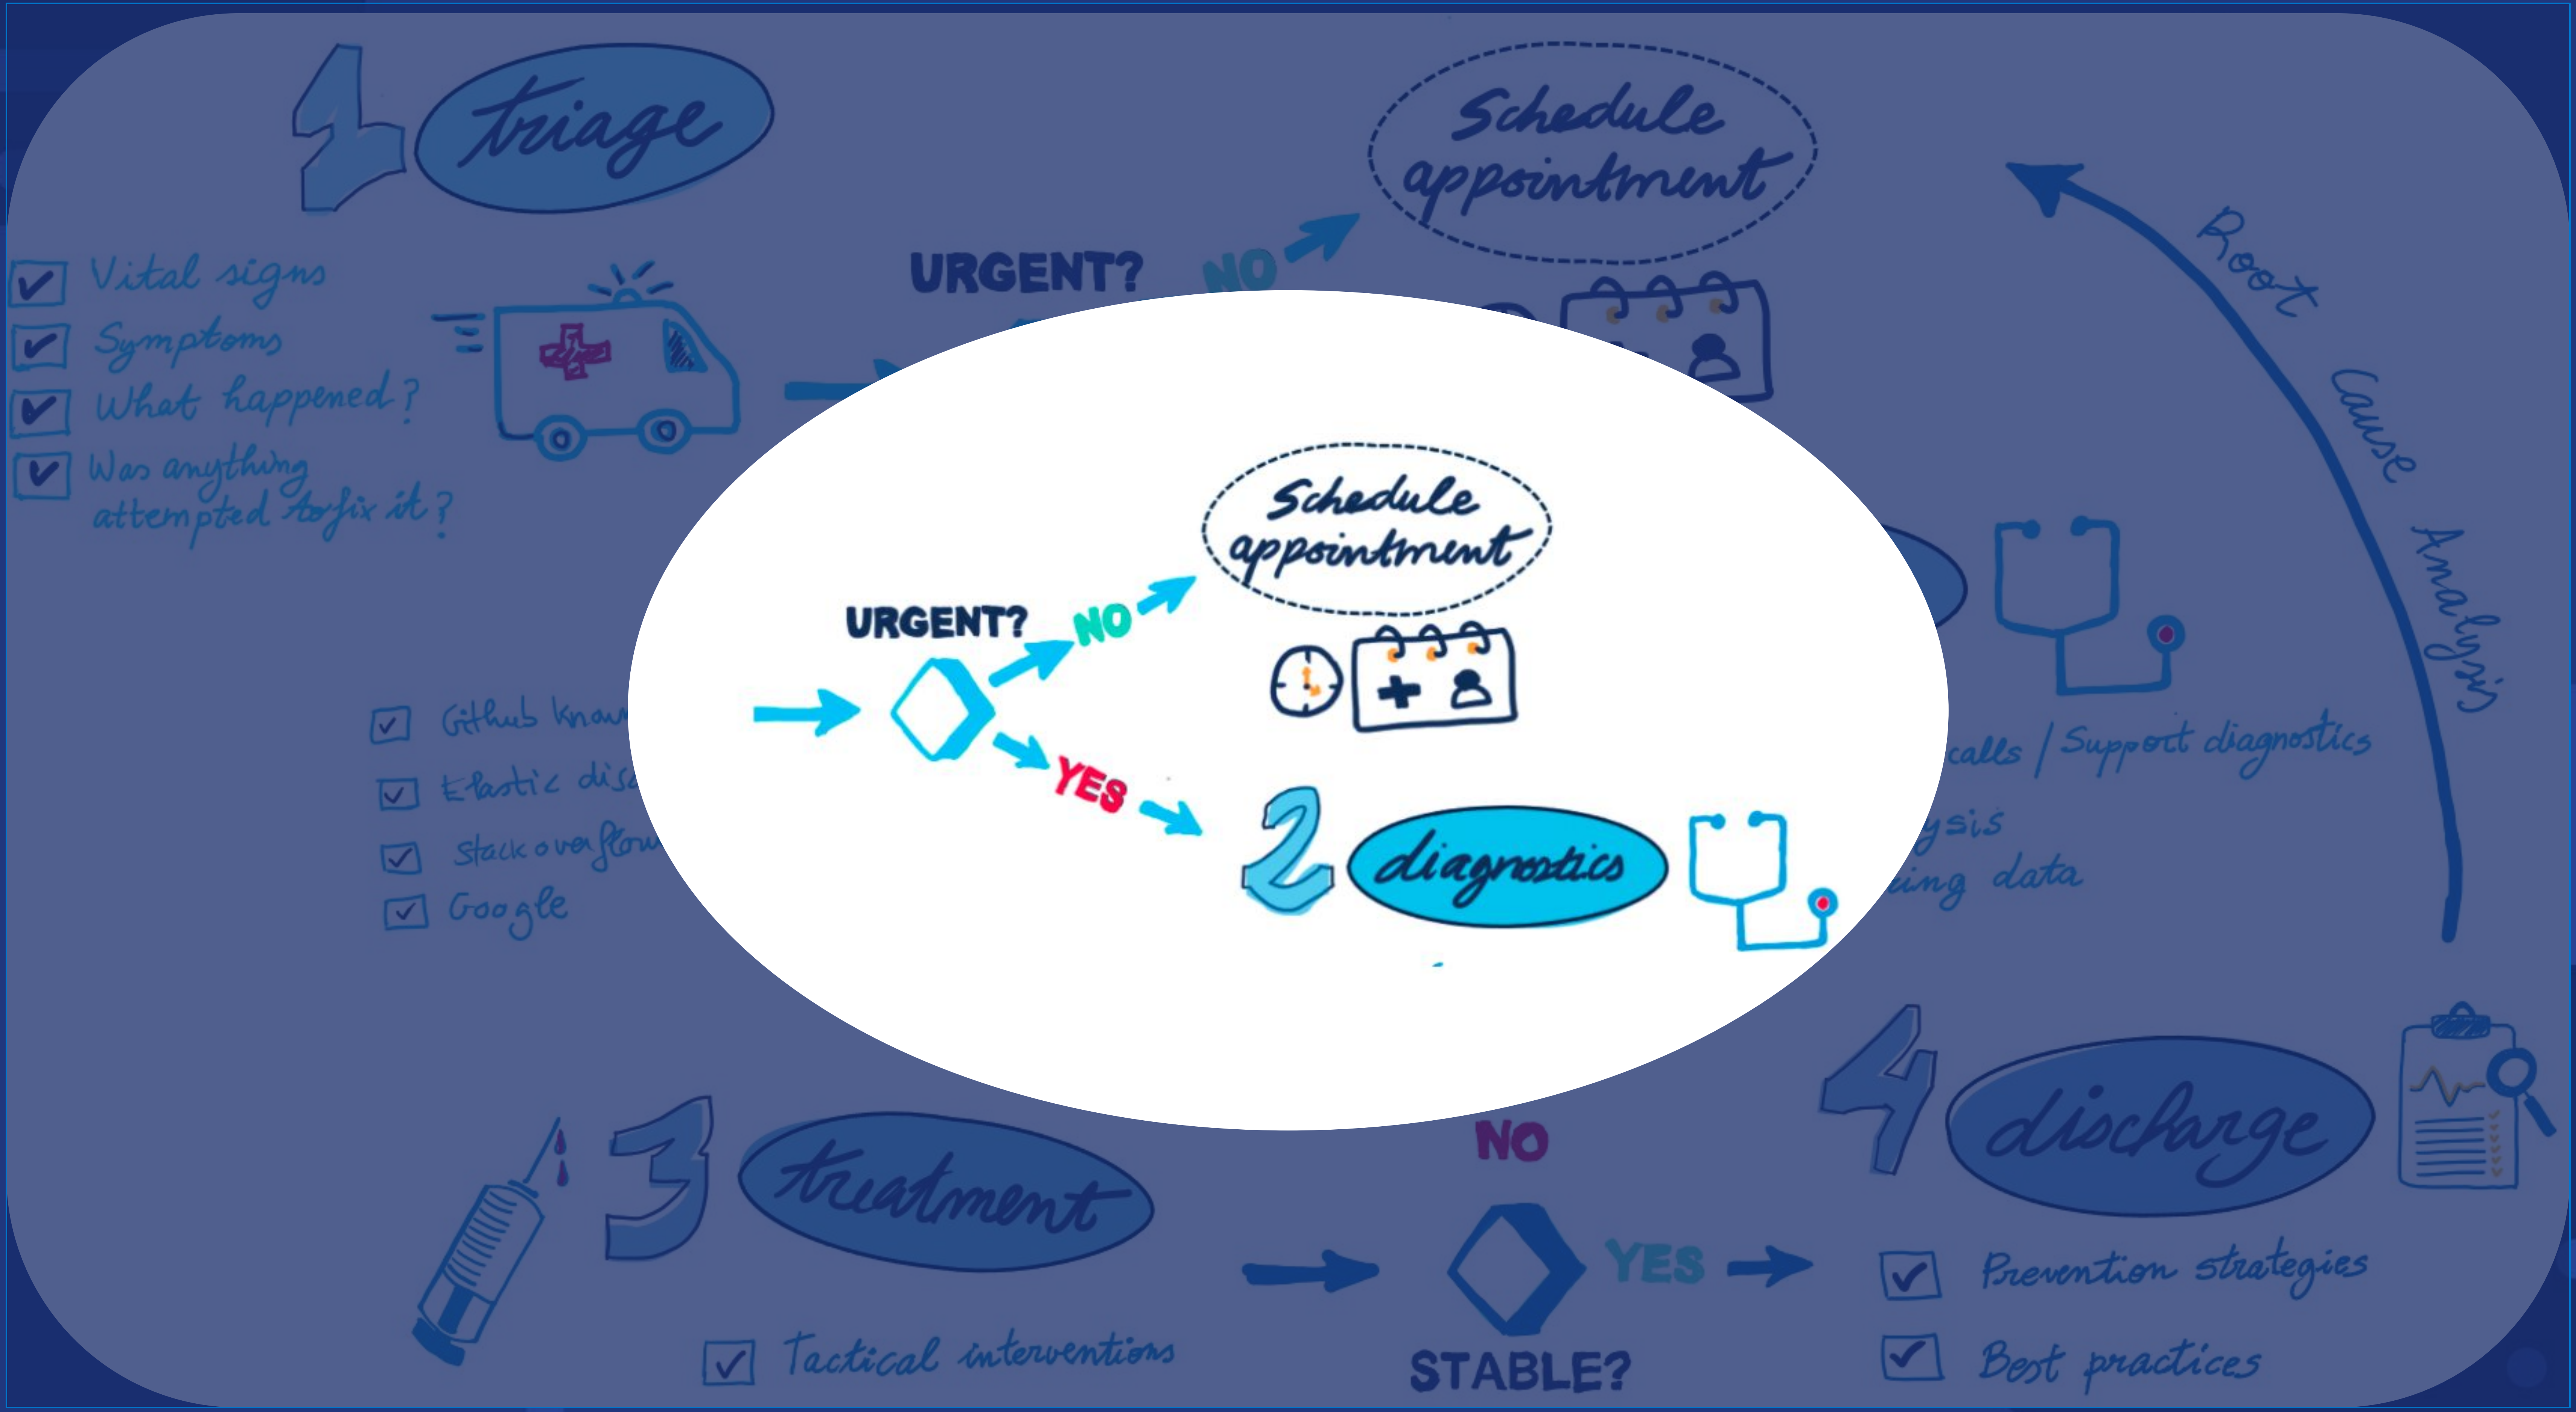
1 triage
☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix it?
URGENT?
NO
YES
Schedule appointment
2 diagnostics
☑ Github know
☑ Elastic dis
☑ stackoverflow
☑ Google
calls / Support diagnostics
ysis
ing data
Root Cause Analysis
3 treatment
☑ Tactical interventions
NO
STABLE?
YES
4 discharge
☑ Prevention strategies
☑ Best practices

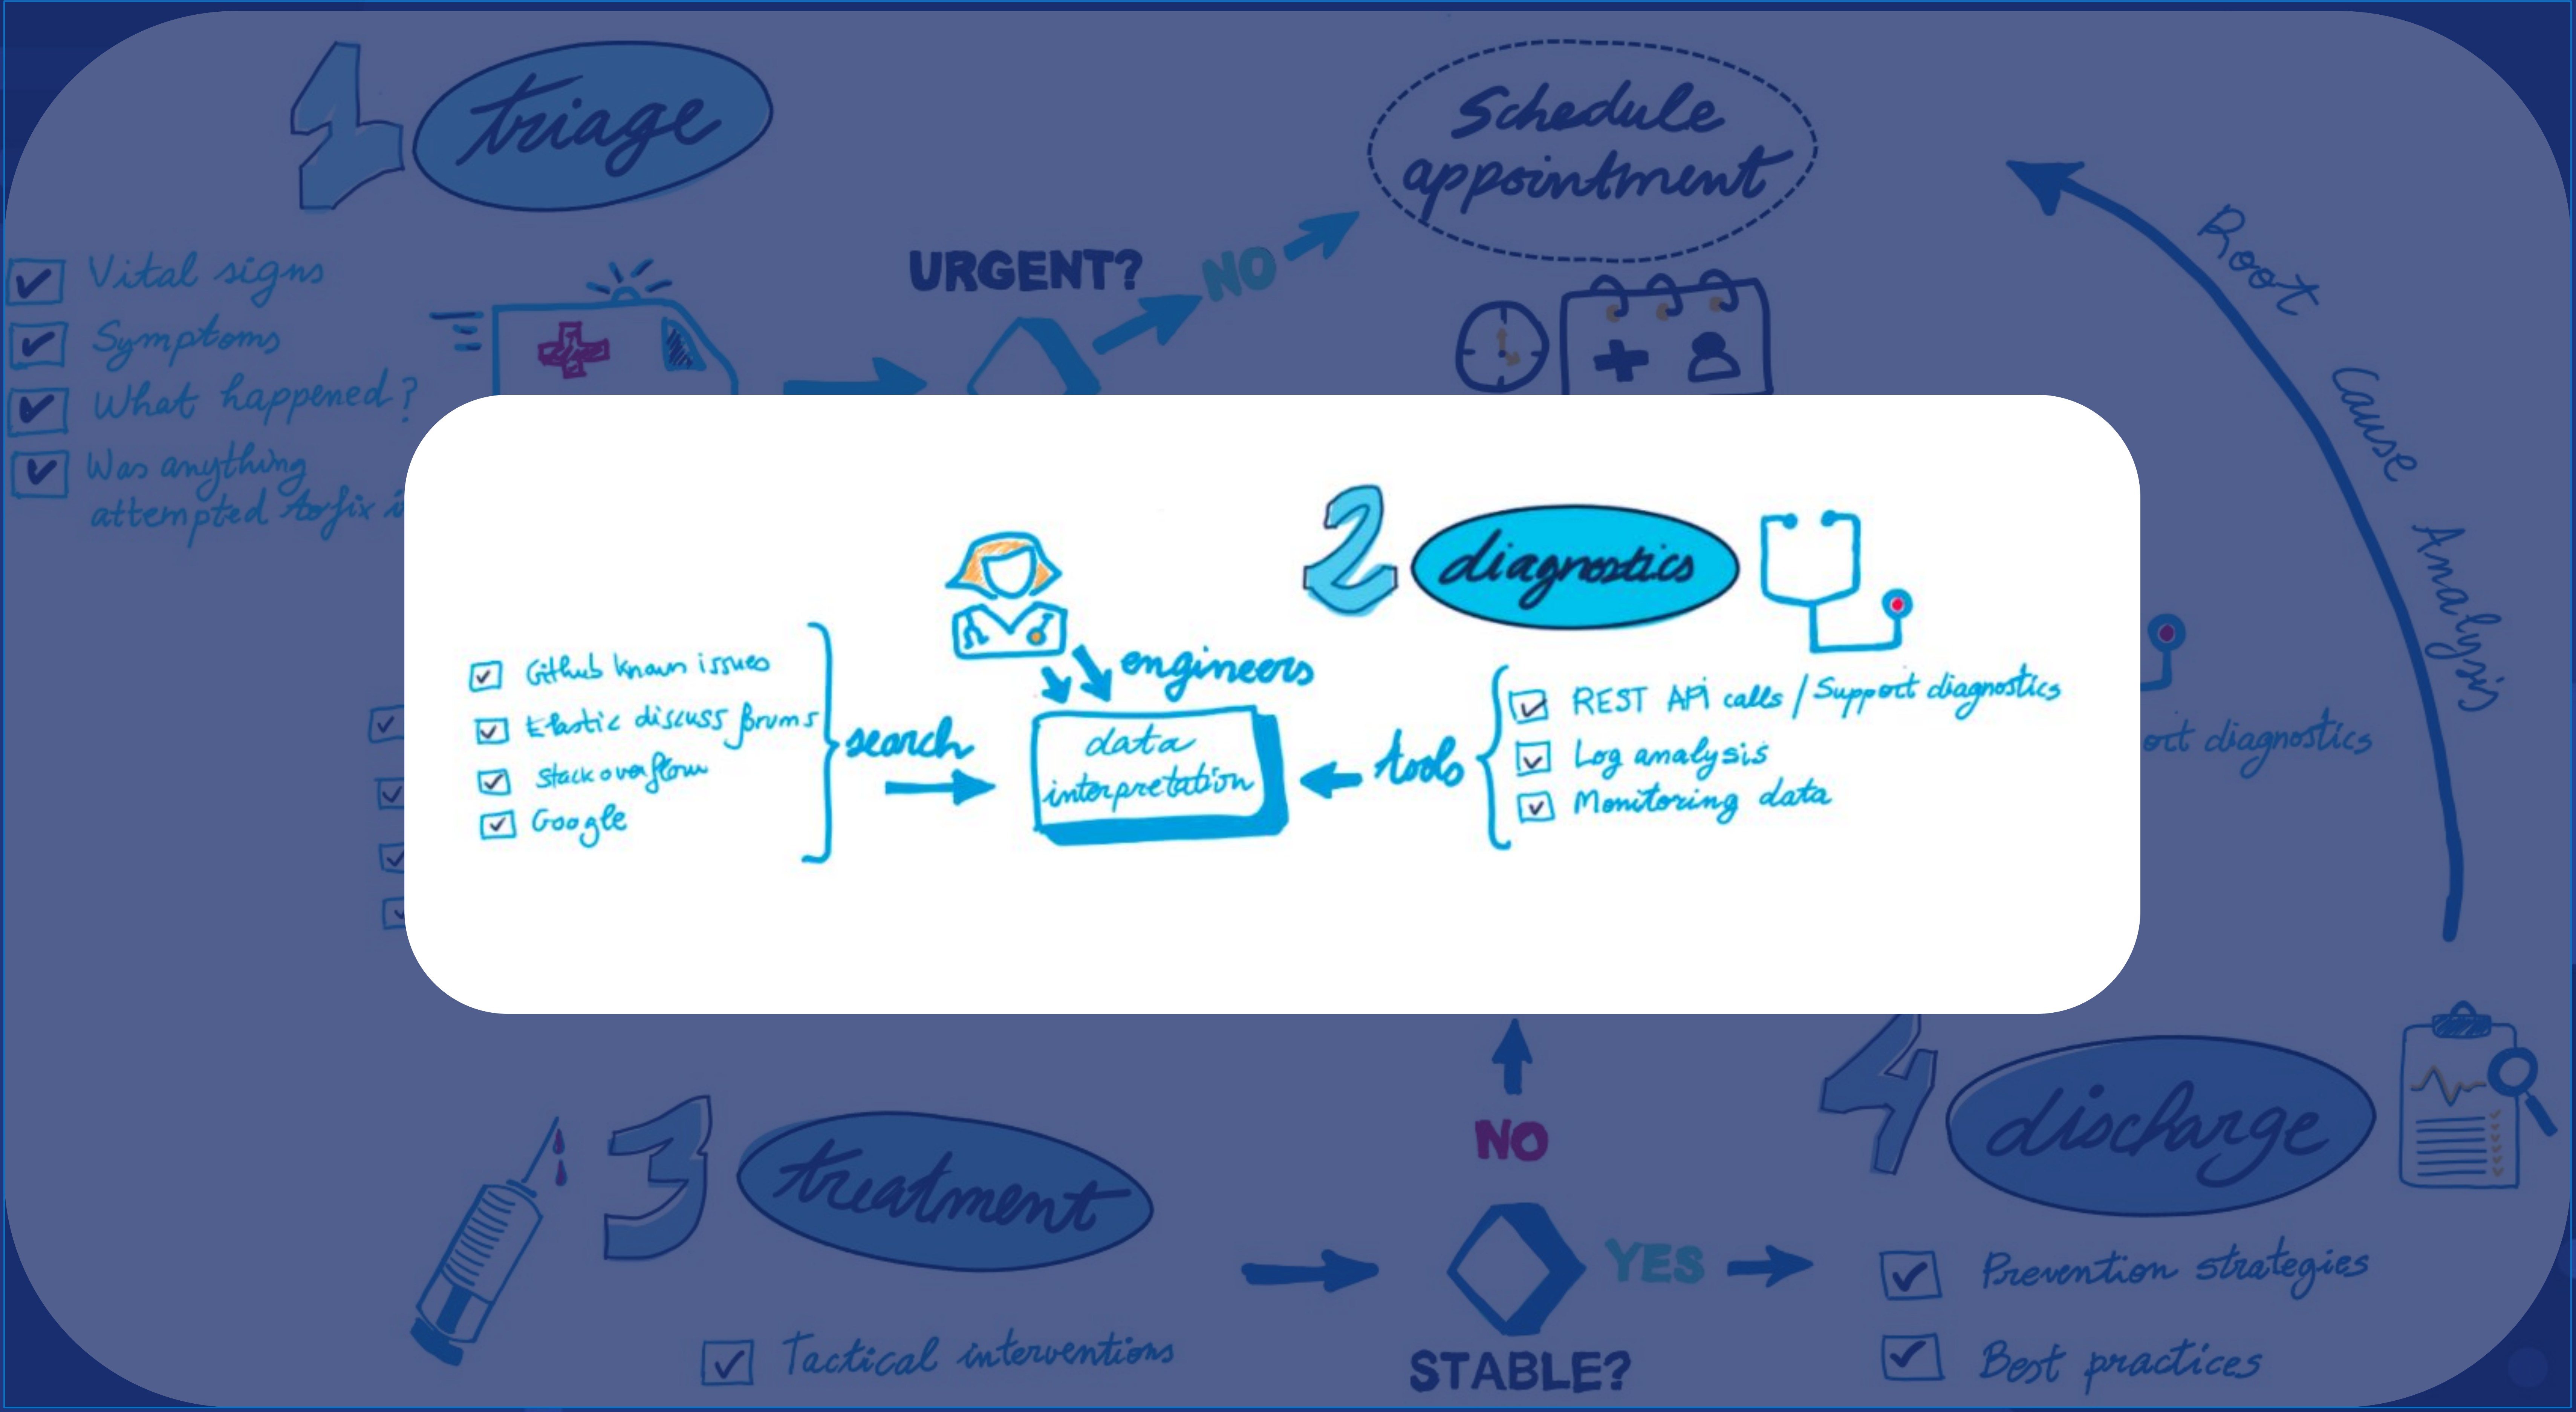
2
diagnostics
☑ Github known issues
☑ Elastic discuss forums
☑ stackoverflow
☑ Google
search
engineers
data interpretation
tools
☑ REST API calls / Support diagnostics
☑ Log analysis
☑ Monitoring data
1
triage
☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix i
URGENT?
NO
Schedule appointment
Root Cause Analysis
3
treatment
☑ Tactical interventions
NO
YES
STABLE?
4
discharge
☑ Prevention strategies
☑ Best practices

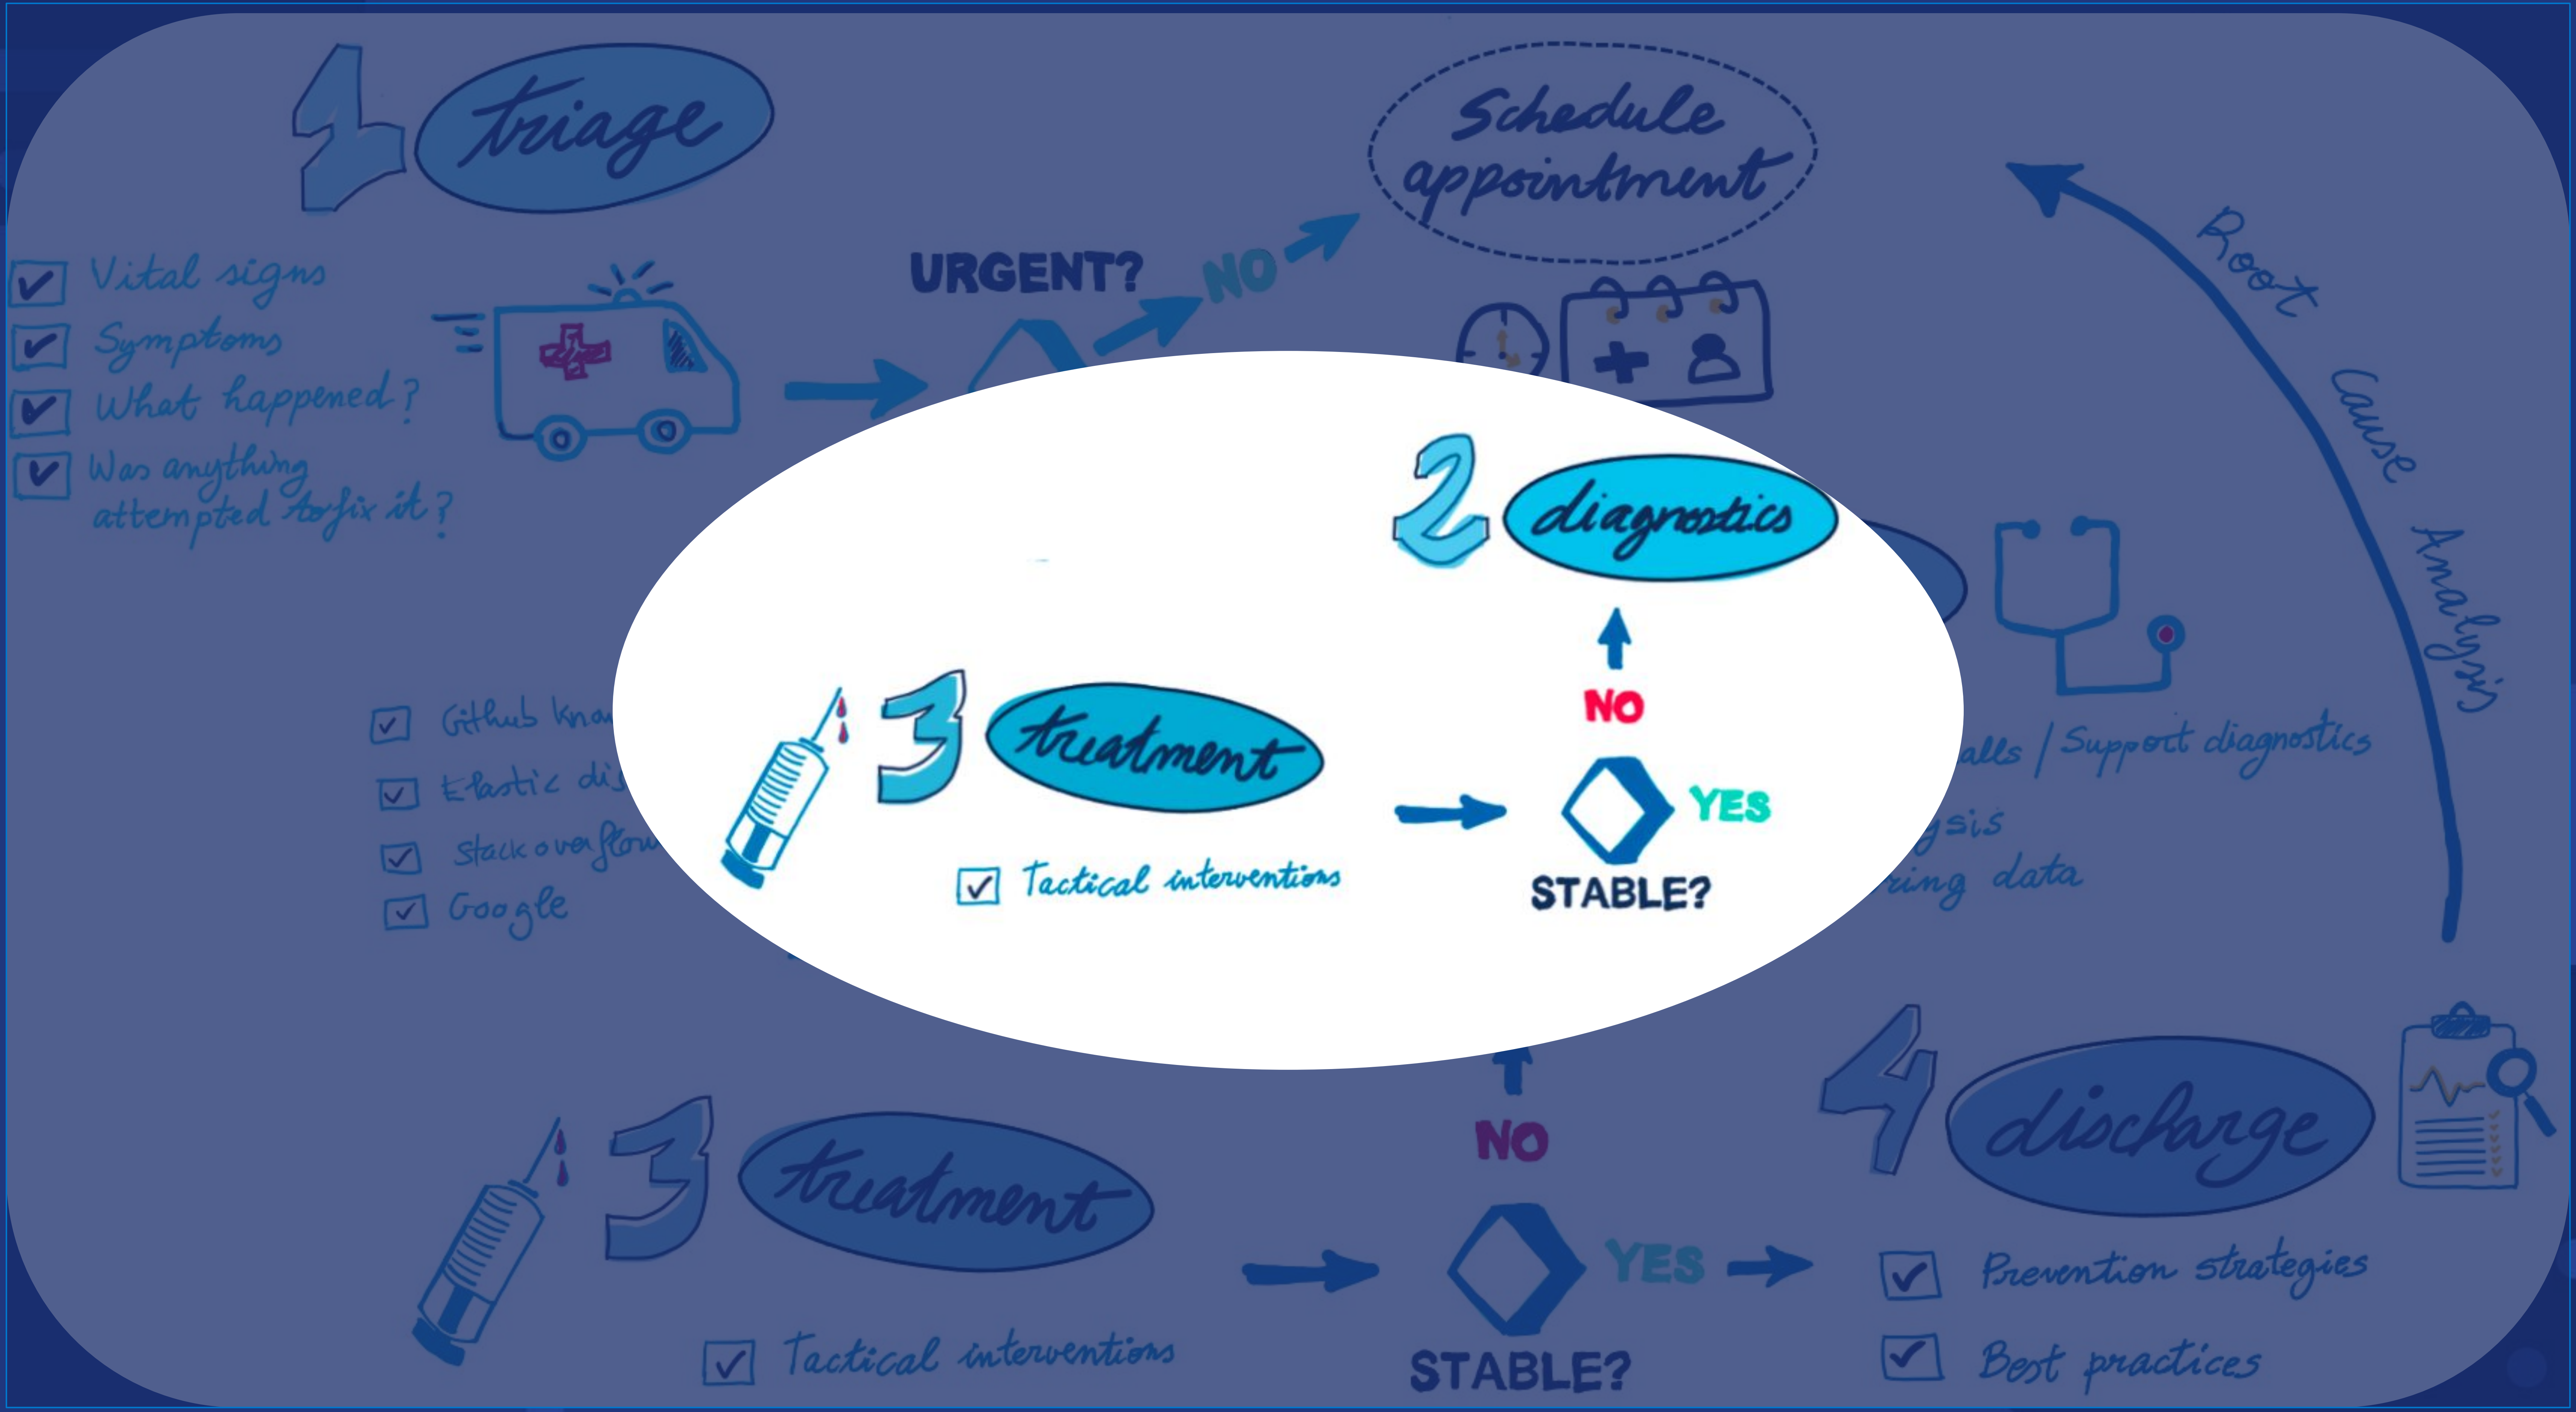
1
triage
☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix it?
URGENT?
NO
Schedule appointment
Root Cause Analysis
☑ Github kno
☑ Elastic di
☑ stackoverflow
☑ Google
alls / Support diagnostics
ysis
ring data
2
diagnostics
NO
3
treatment
☑ Tactical interventions
YES
STABLE?
3
treatment
☑ Tactical interventions
NO
YES
STABLE?
4
discharge
☑ Prevention strategies
☑ Best practices

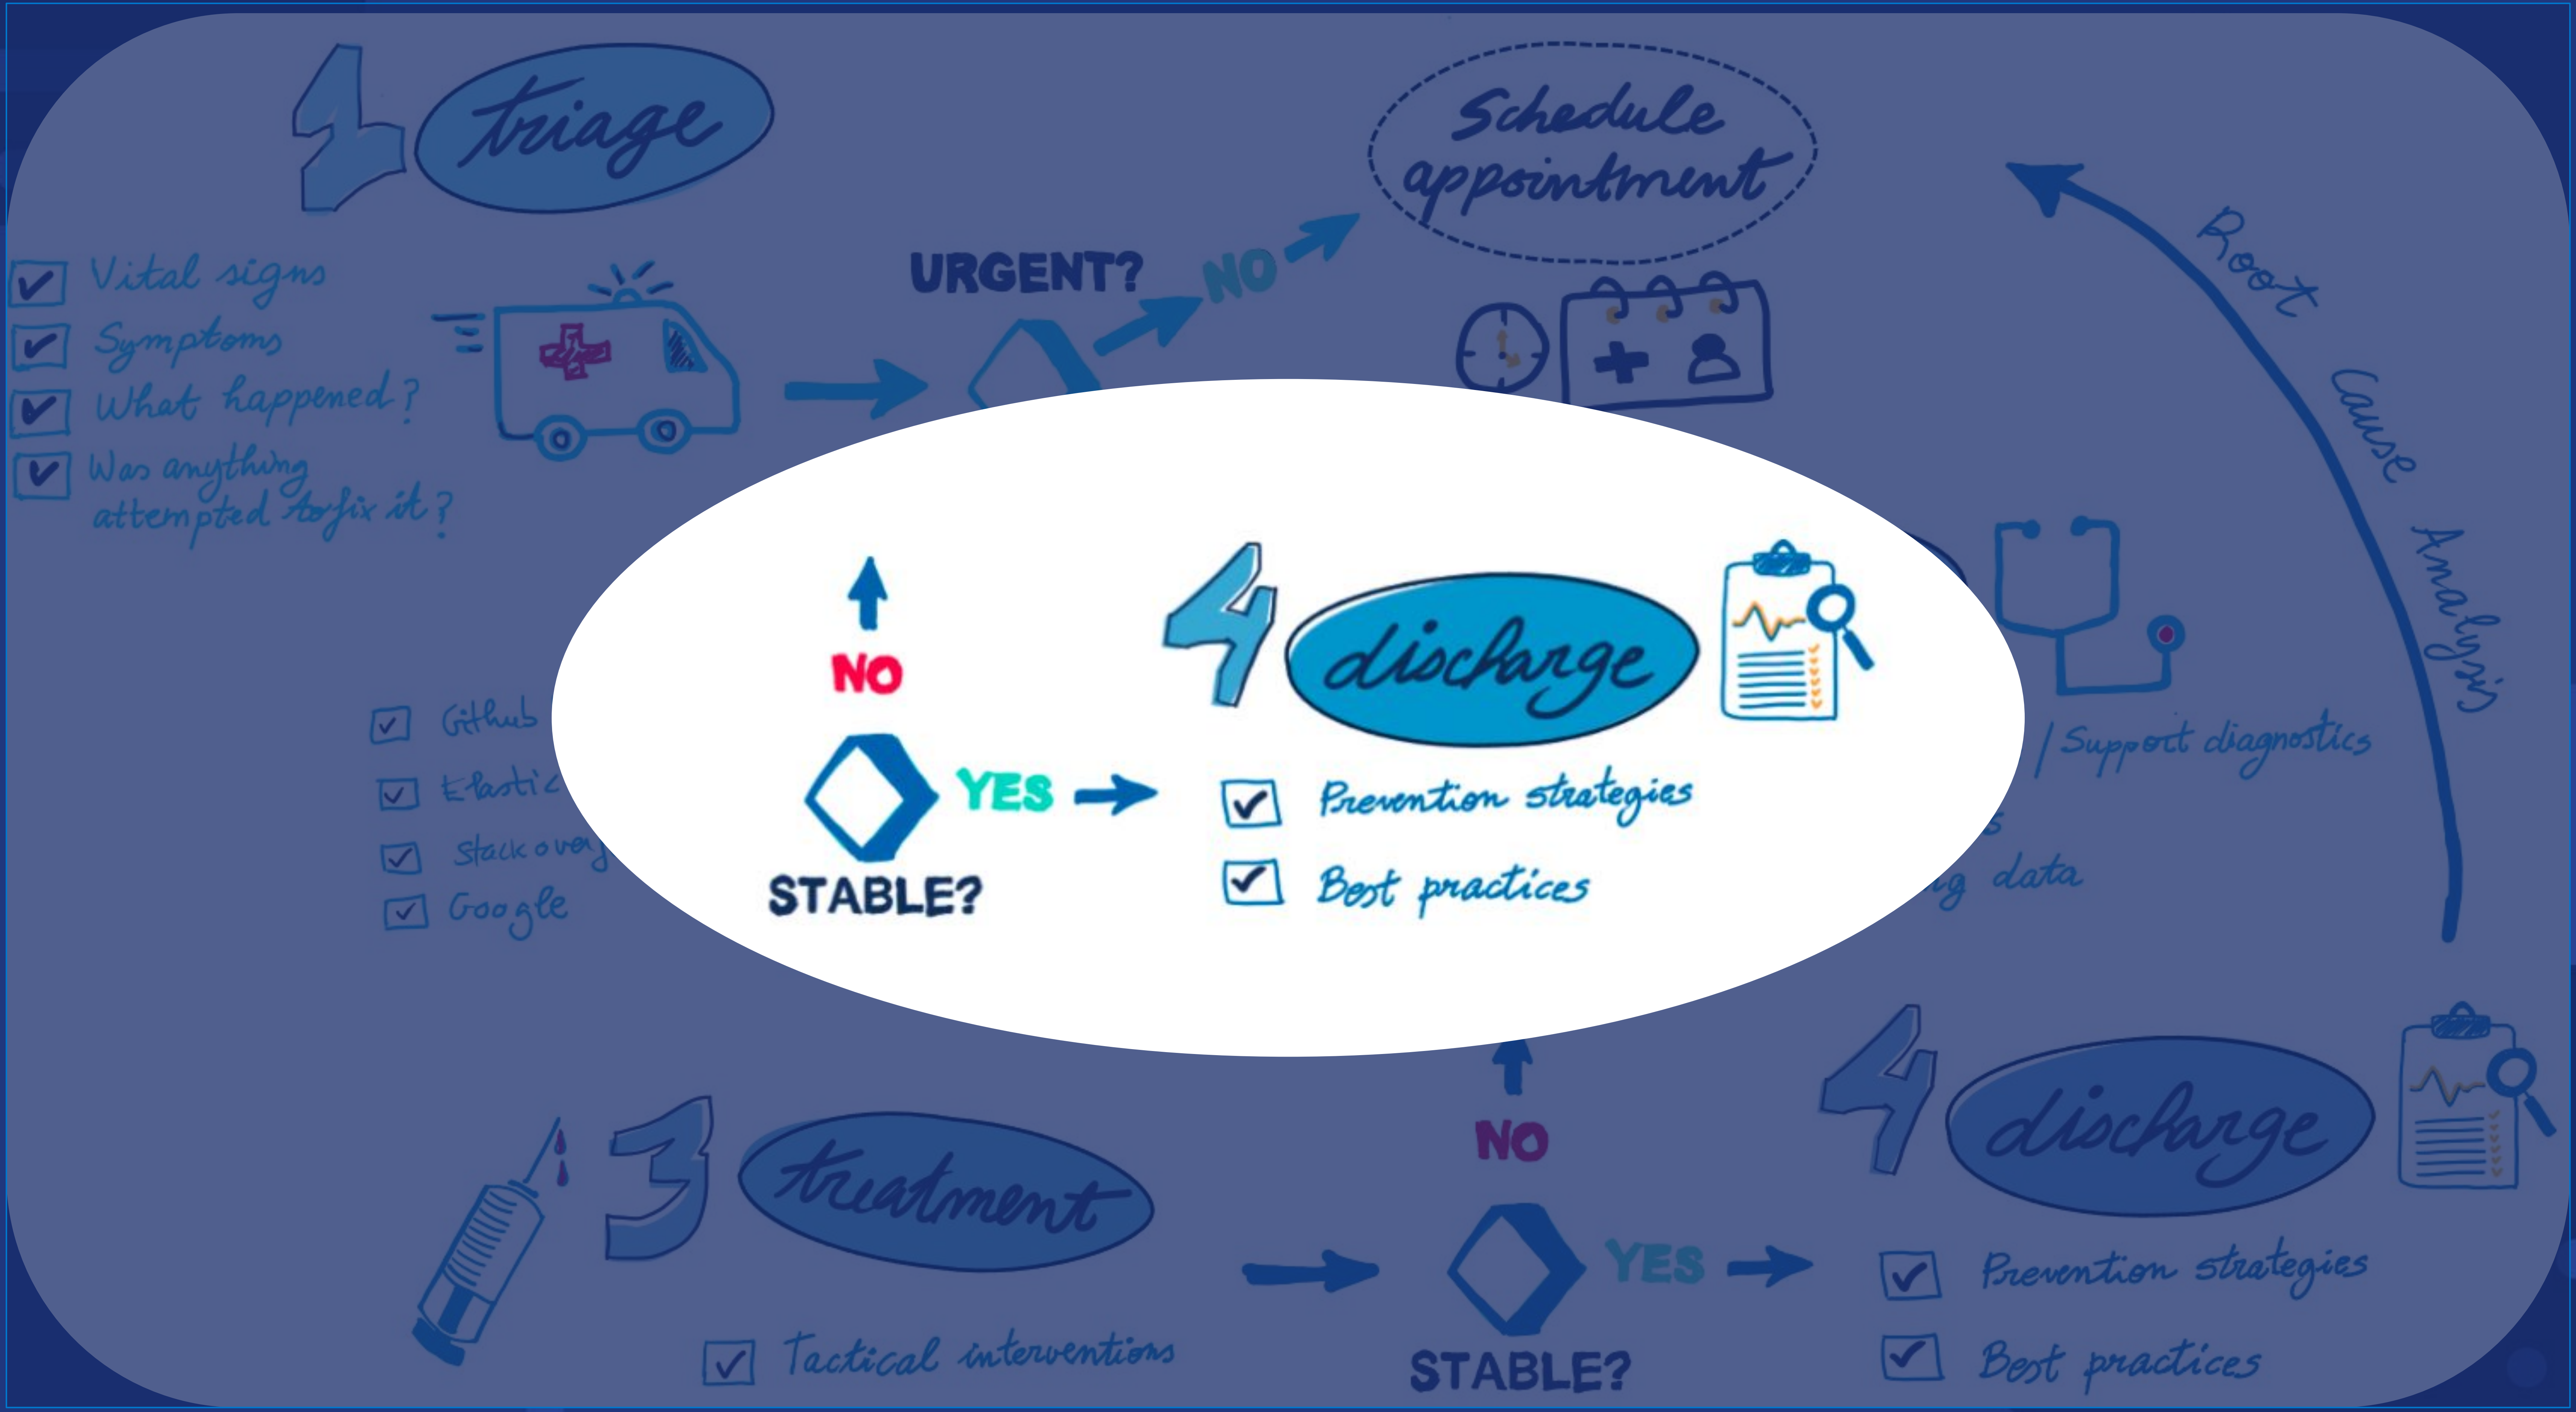
1 triage
Vital signs
Symptoms
What happened?
Was anything attempted to fix it?
URGENT?
NO
Schedule appointment
Root Cause Analysis
NO
4 discharge
YES
STABLE?
Prevention strategies
Best practices
Github
Elastic
stack over
Google
/ Support diagnostics
data
3 treatment
Tactical interventions
NO
YES
STABLE?
4 discharge
Prevention strategies
Best practices

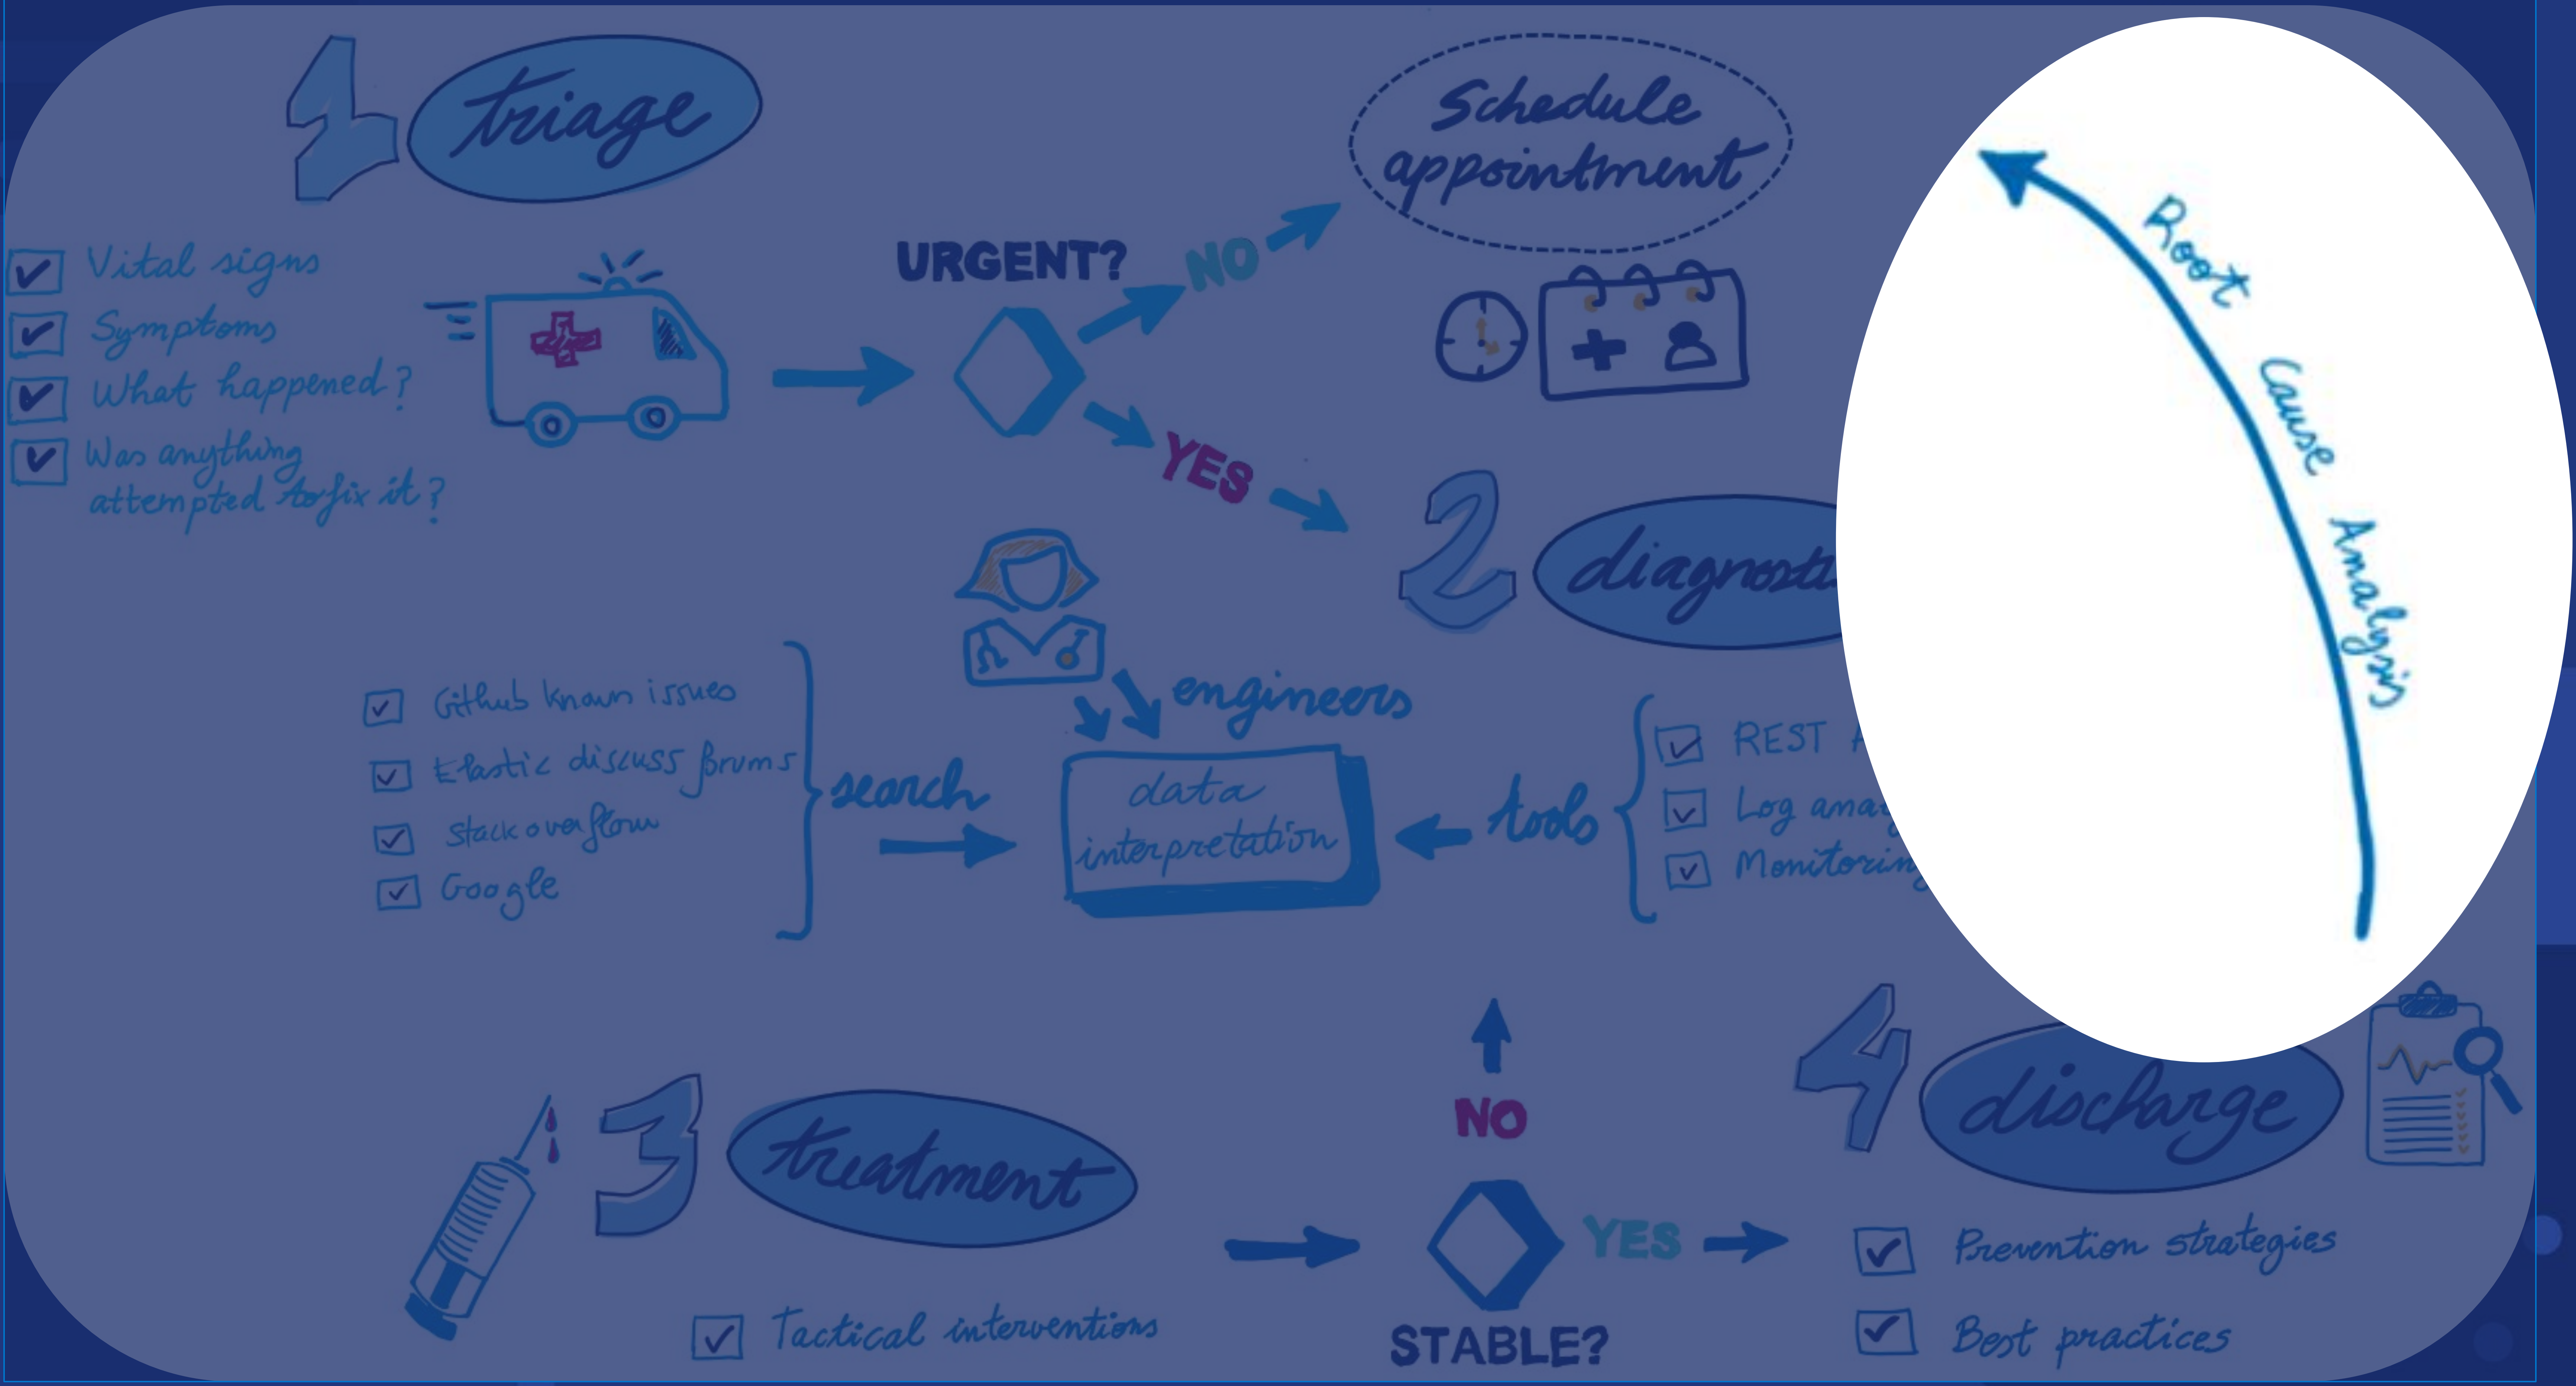
1 triage
☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix it?
URGENT?
NO
Schedule appointment
YES
2 diagnostic
engineers
☑ Github known issues
☑ Elastic discuss forums
☑ stackoverflow
☑ Google
search
data interpretation
tools
☑ REST A
☑ Log ana
☑ Monitorin
Root Cause Analysis
3 treatment
☑ Tactical interventions
STABLE?
NO
YES
4 discharge
☑ Prevention strategies
☑ Best practices

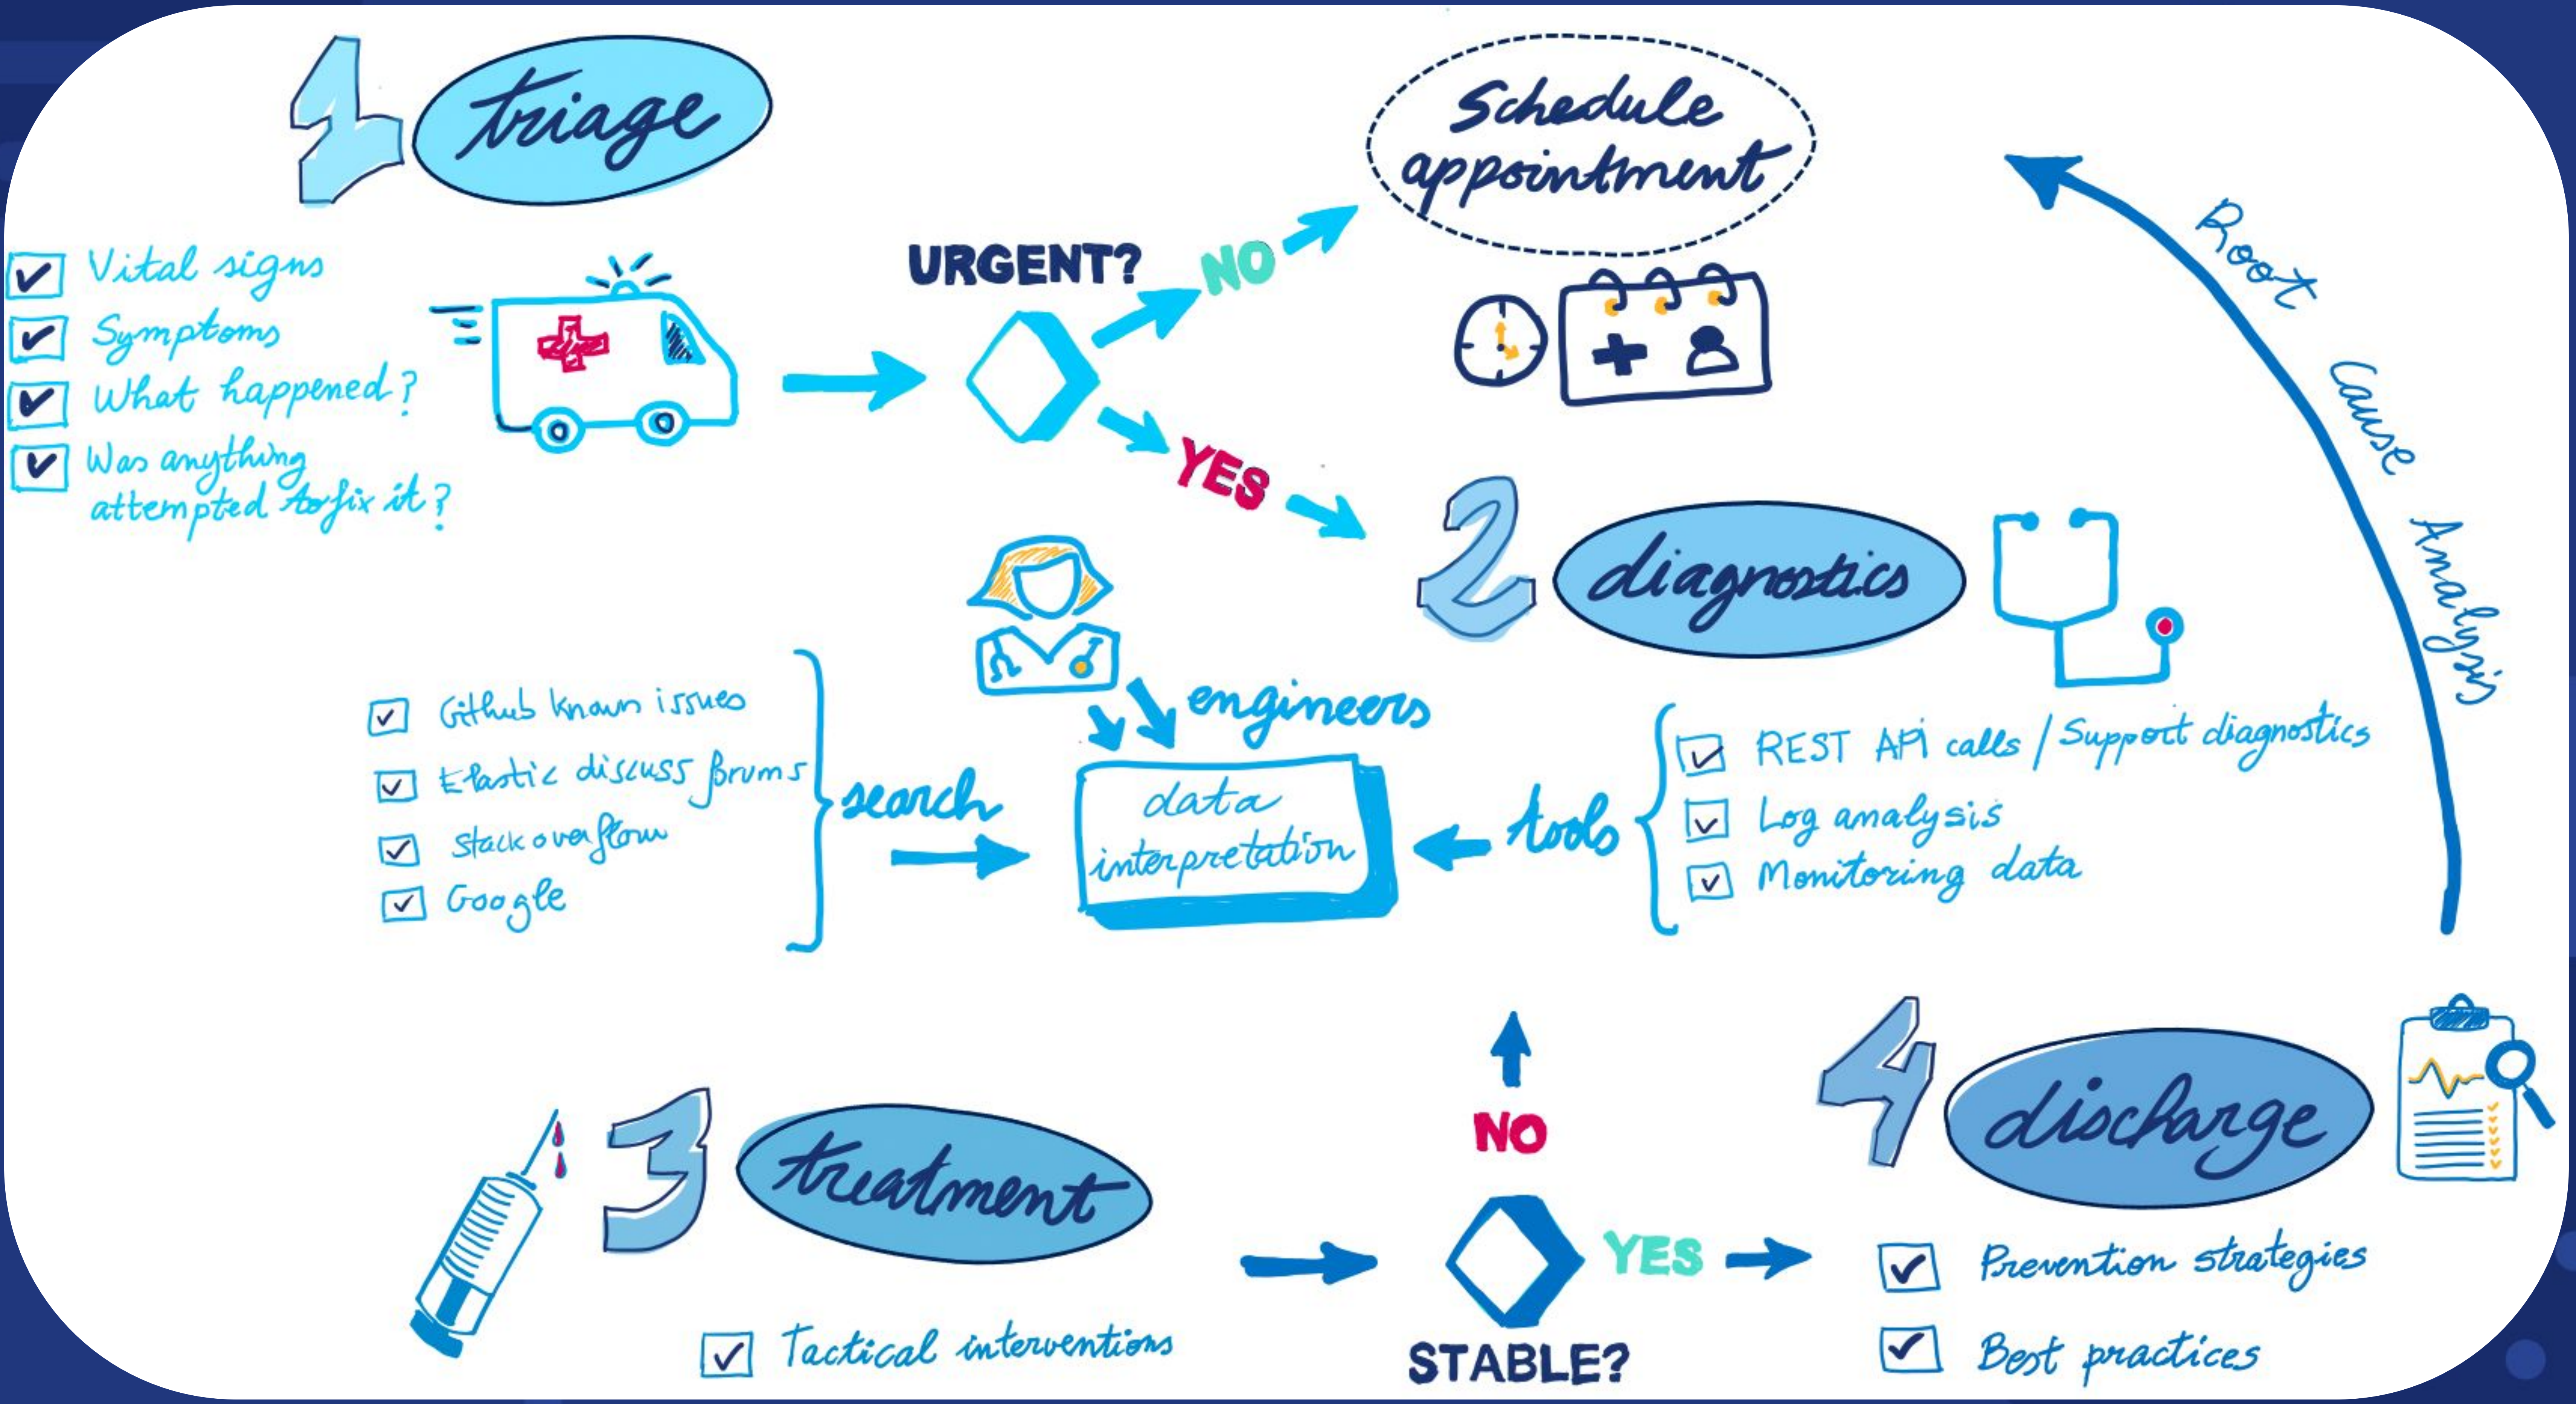
1 triage
☑ Vital signs
☑ Symptoms
☑ What happened?
☑ Was anything attempted to fix it?
URGENT?
NO
Schedule appointment
YES
2 diagnostics
engineers
☑ Github known issues
☑ Elastic discuss forums
☑ stack overflow
☑ Google
search
data interpretation
tools
☑ REST API calls / Support diagnostics
☑ Log analysis
☑ Monitoring data
Root Cause Analysis
3 treatment
☑ Tactical interventions
STABLE?
NO
YES
4 discharge
☑ Prevention strategies
☑ Best practices

Most Common Issues?

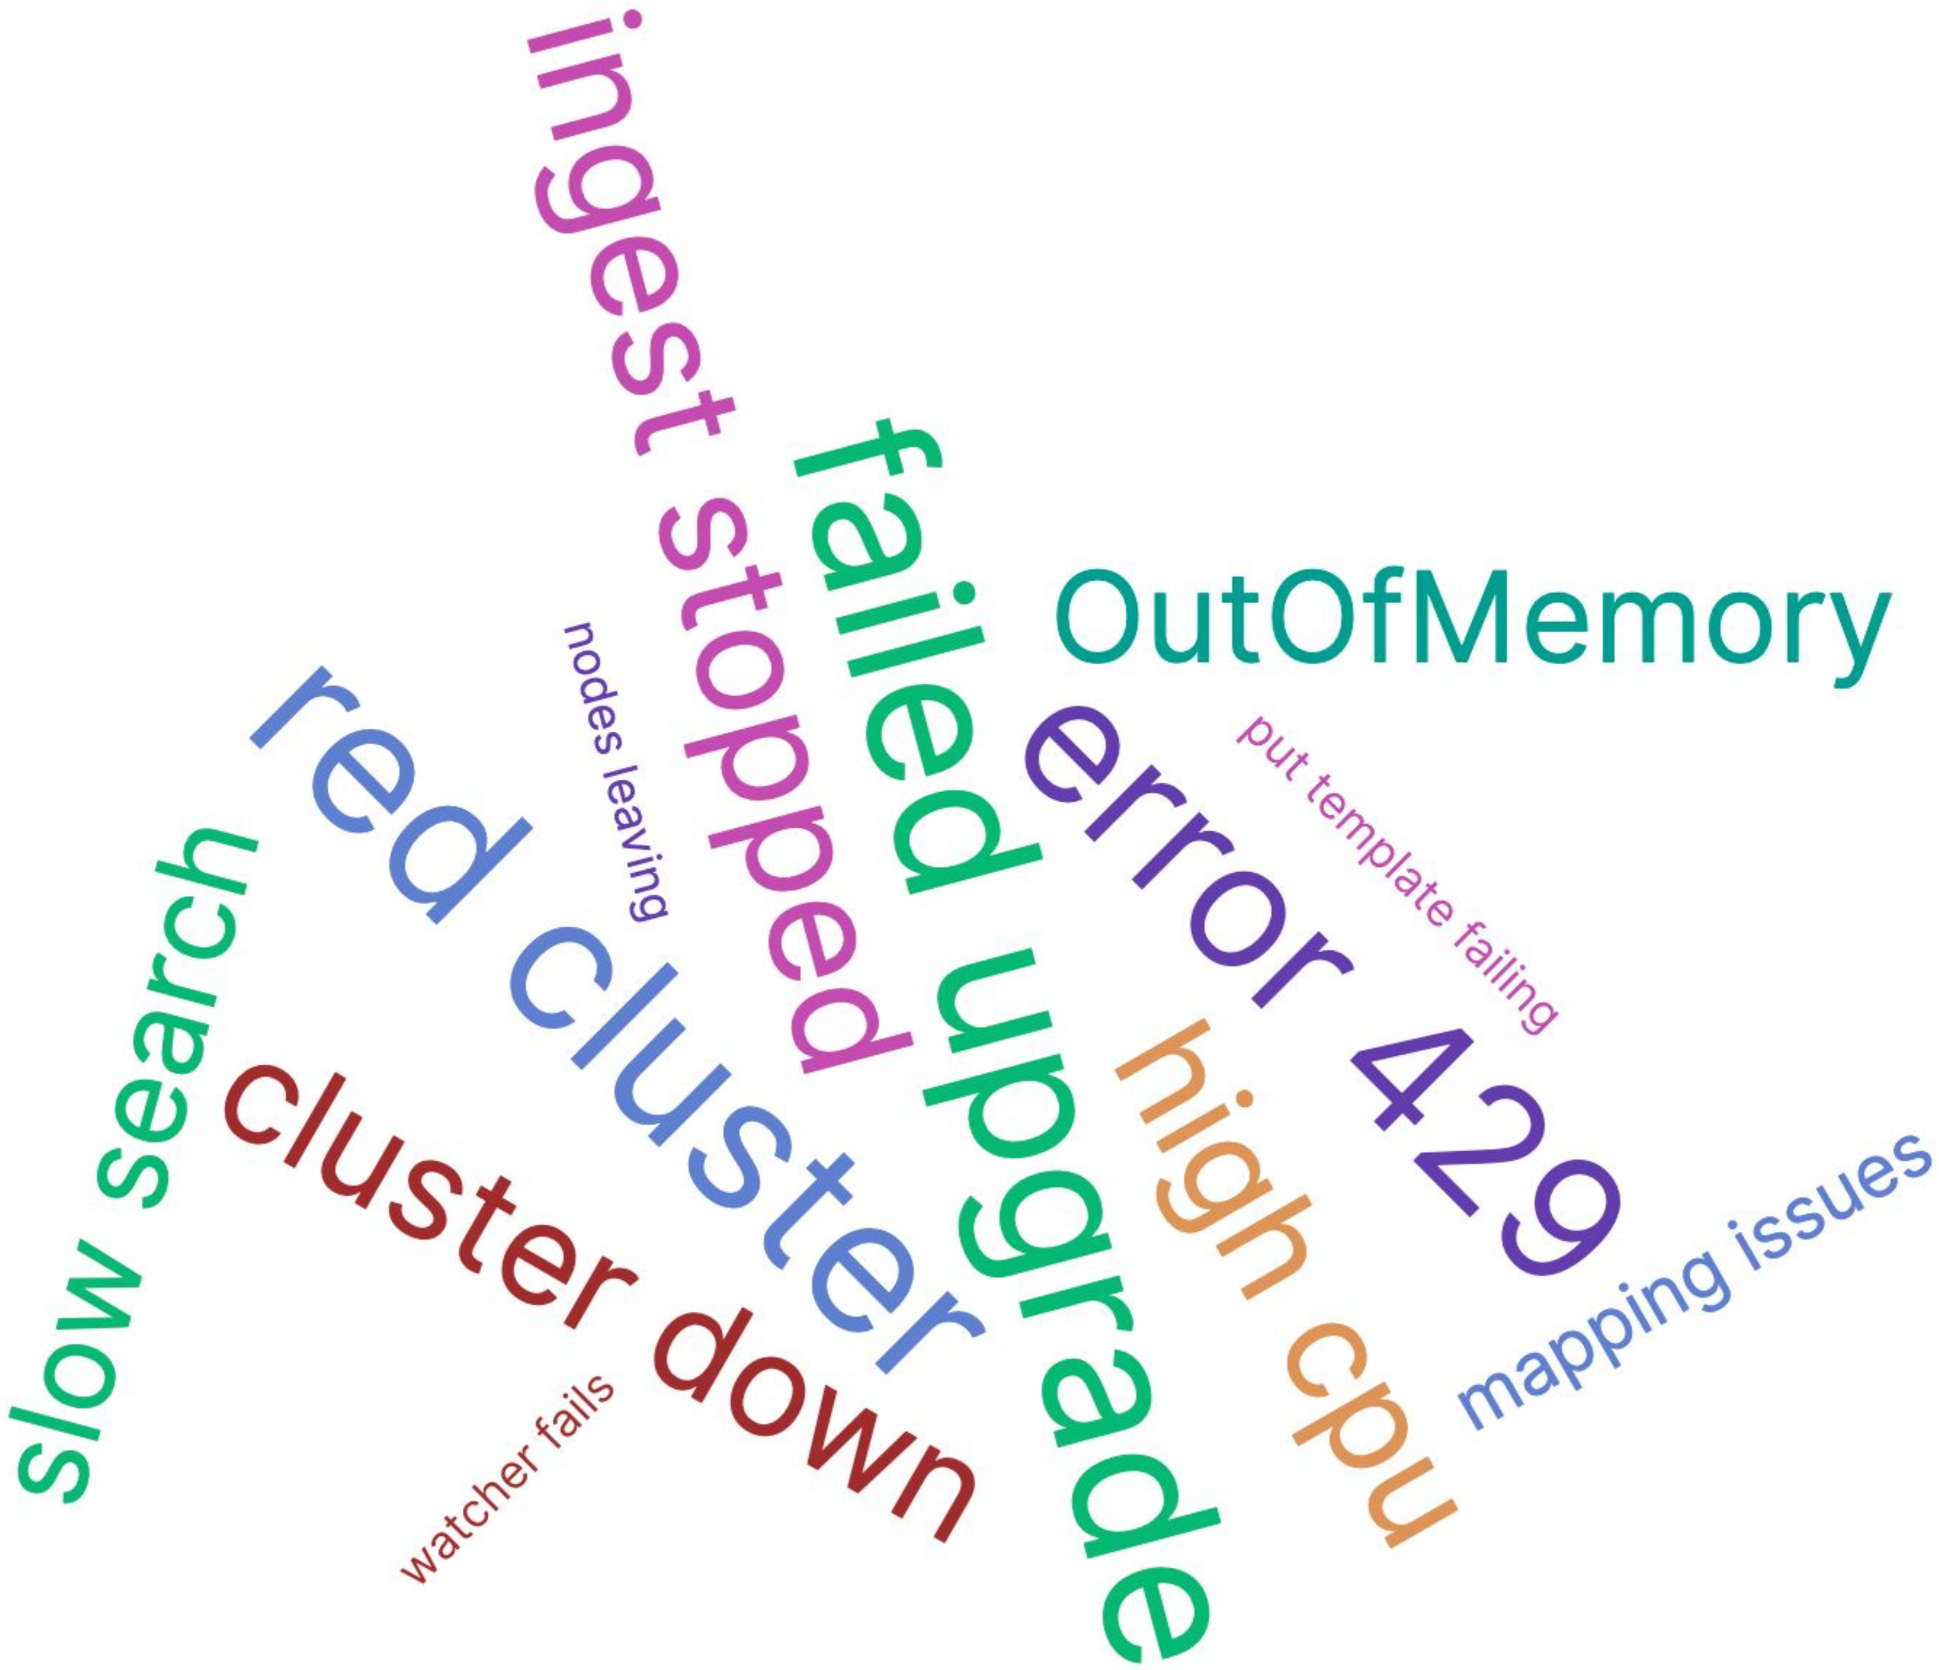

ingest stopped

failed upgrade

OutOfMemory

error 429

red cluster

nodes leaving

put template failing

high cpu

mapping issues

slow search

cluster down

watcher fails

Troubleshooting by Example

TRIAGE

[Urgent Severity] **Red cluster**

Vital signs

- → Cluster in red health
- → No ingest into any indices

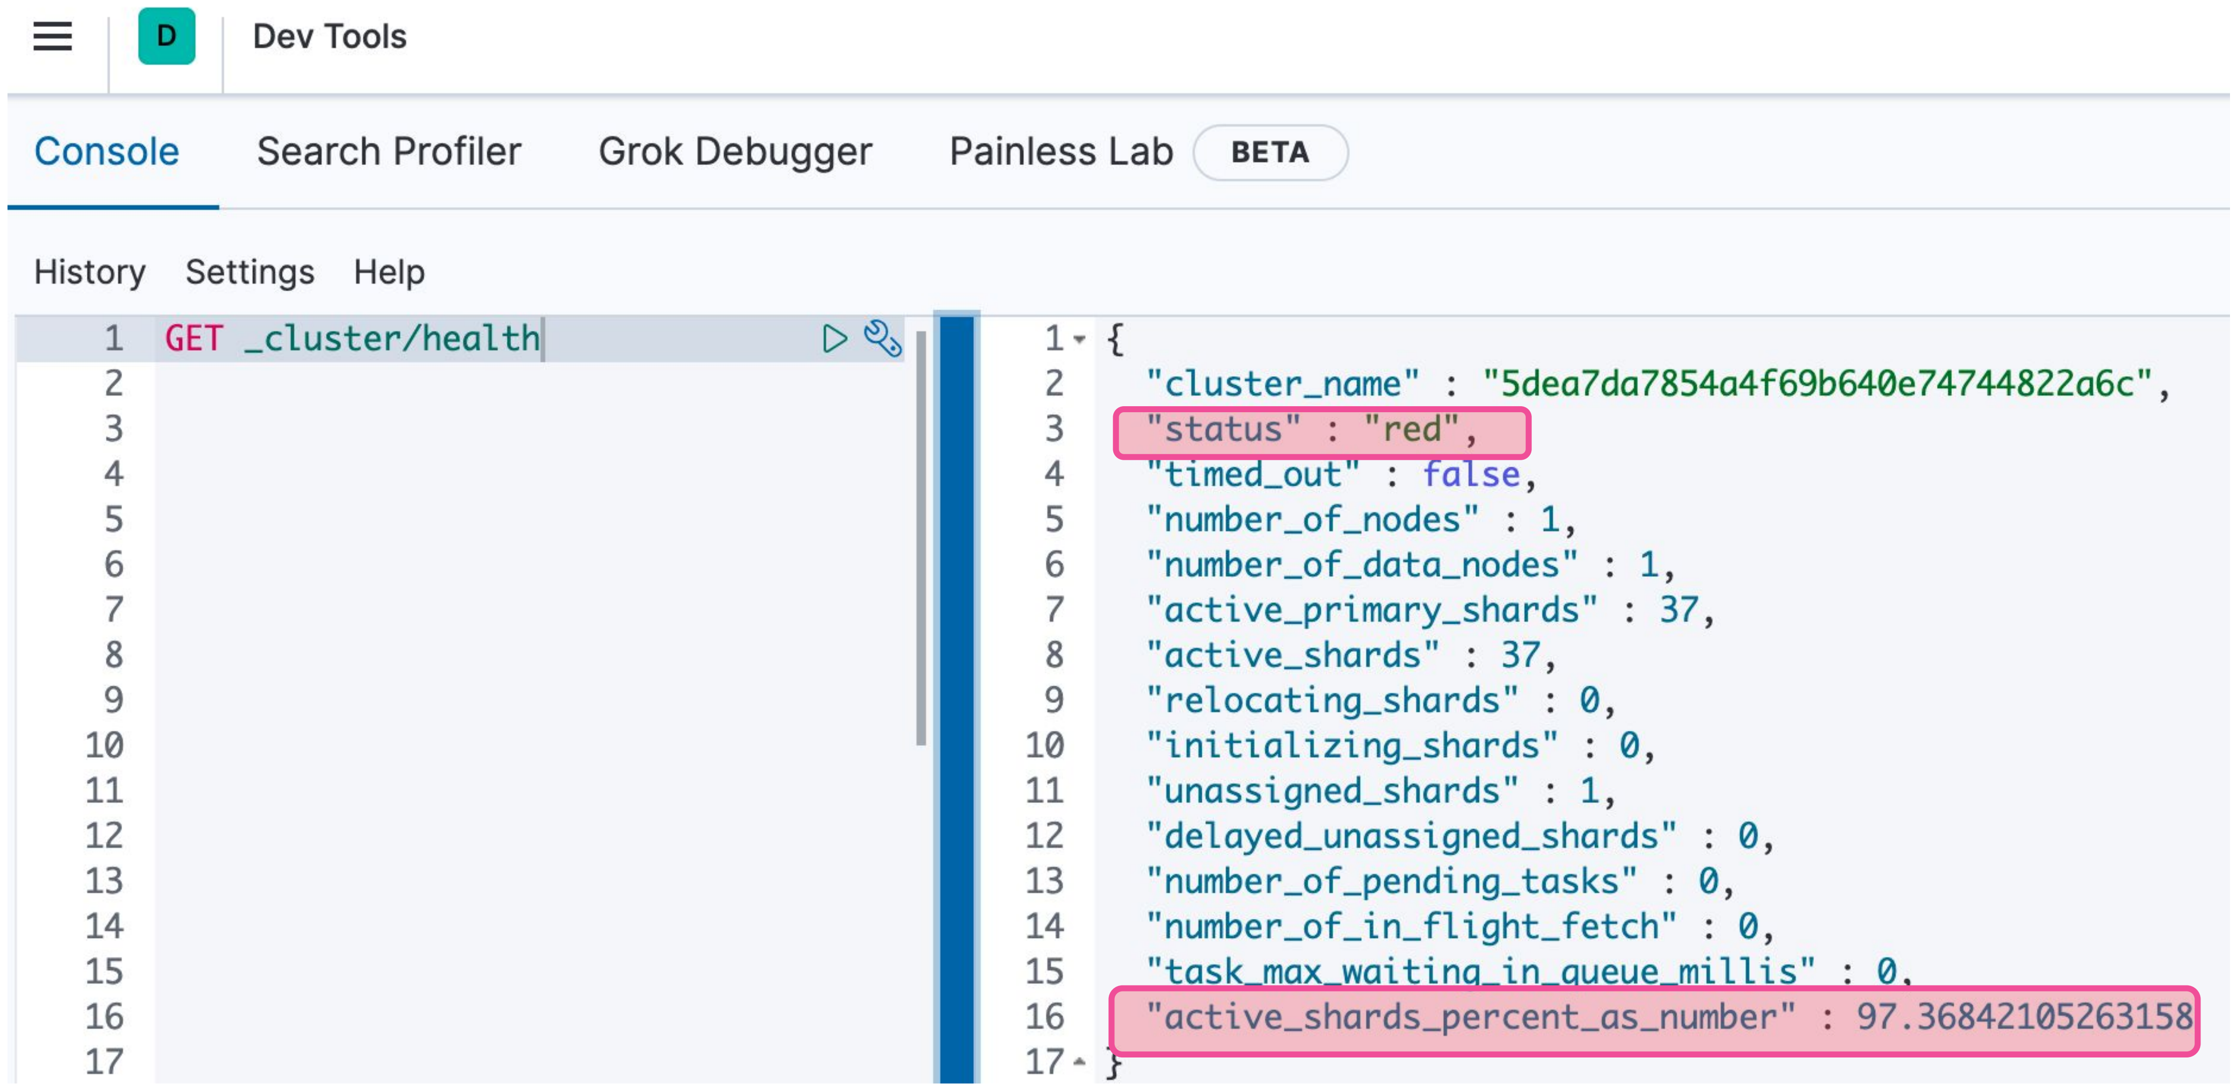

```
GET _cluster/health
```

```
{
  "cluster_name" : "5dea7da7854a4f69b640e74744822a6c",
  "status" : "red",
  "timed_out" : false,
  "number_of_nodes" : 1,
  "number_of_data_nodes" : 1,
  "active_primary_shards" : 37,
  "active_shards" : 37,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 1,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 97.36842105263158
}
```

Symptoms

- → Beats fail to ingest
- → Cluster is responsive, search and REST API still work

elastic

TRIAGE

[Urgent Severity] **Red cluster**

What happened?

→ Out of the blue, no changes

Any attempts to fix it?

→ No

Next steps

→ Share a support diagnostics that will provide REST API calls

https://www.elastic.co/blog/why-does-elastic-support-keep-asking-for-diagnostic-files

https://github.com/elastic/support-diagnostics/blob/main/src/main/resources/elastic-rest.yml

https://github.com/elastic/support-diagnostics

```
> ./diagnostics.sh --host https://localhost -u elastic -p --port 9200 --ssl --type api --noVerify
```

elastic

[Urgent Severity] **Red cluster**

Why is the cluster red?

→ REST API calls - CAT Indices API

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/rest-apis.html

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cat-indices.html

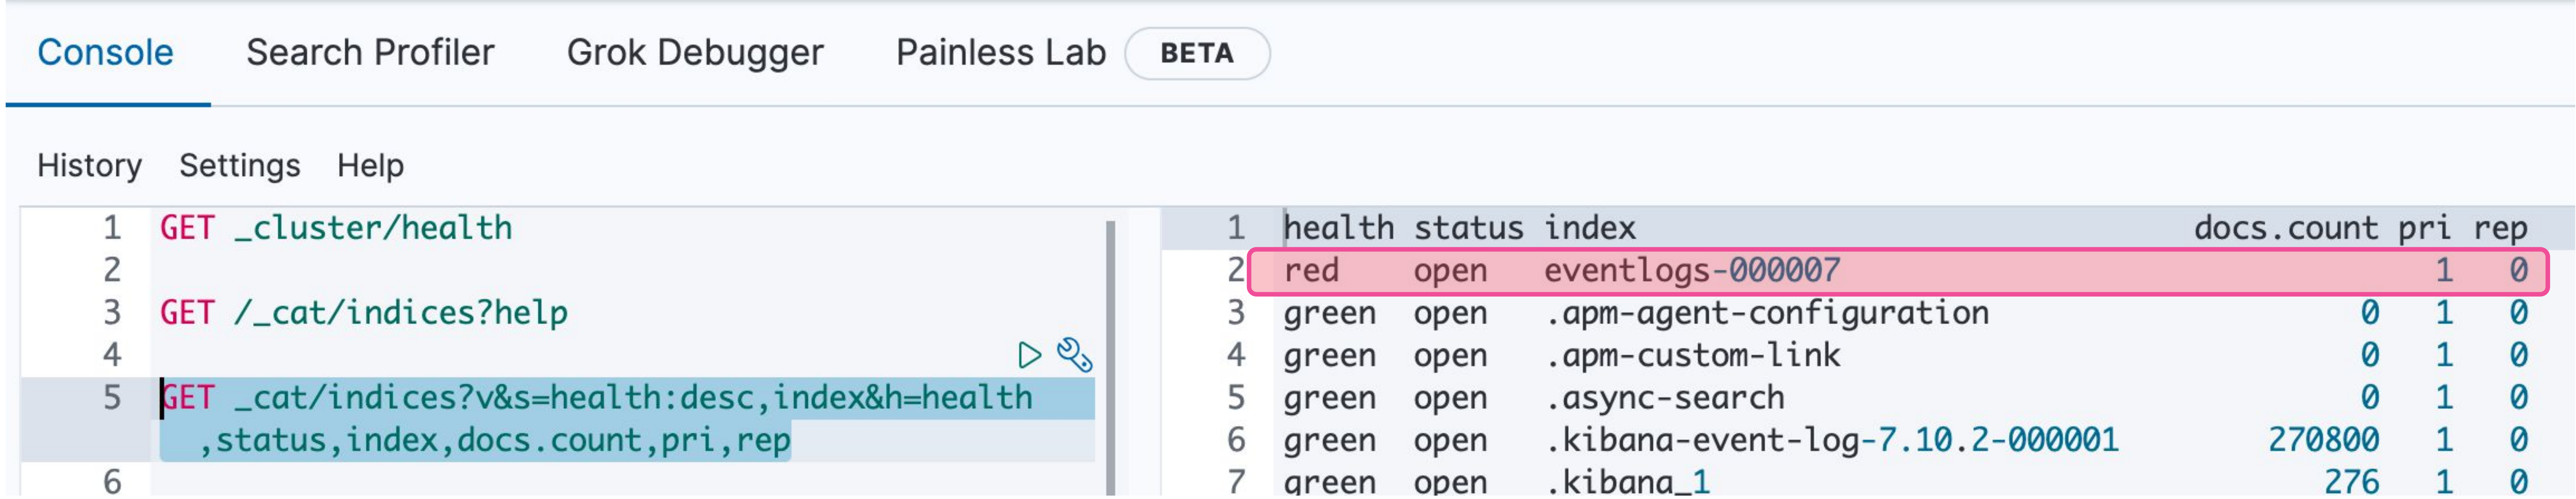

```
GET _cluster/health

GET /_cat/indices?help

GET _cat/indices?v&s=health:desc,index&h=health
  ,status,index,docs.count,pri,rep

```

```
health status index                            docs.count pri rep
red    open   eventlogs-000007                                1   0
green  open   .apm-agent-configuration                  0   1   0
green  open   .apm-custom-link                          0   1   0
green  open   .async-search                             0   1   0
green  open   .kibana-event-log-7.10.2-000001      270800   1   0
green  open   .kibana_1                               276   1   0
```

DIAGNOSTIC

[Urgent Severity] **Red cluster**

Why is an index red?

→ Check shards that are not started:

INITIALIZING or **UNASSIGNED**

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cat-shards.html

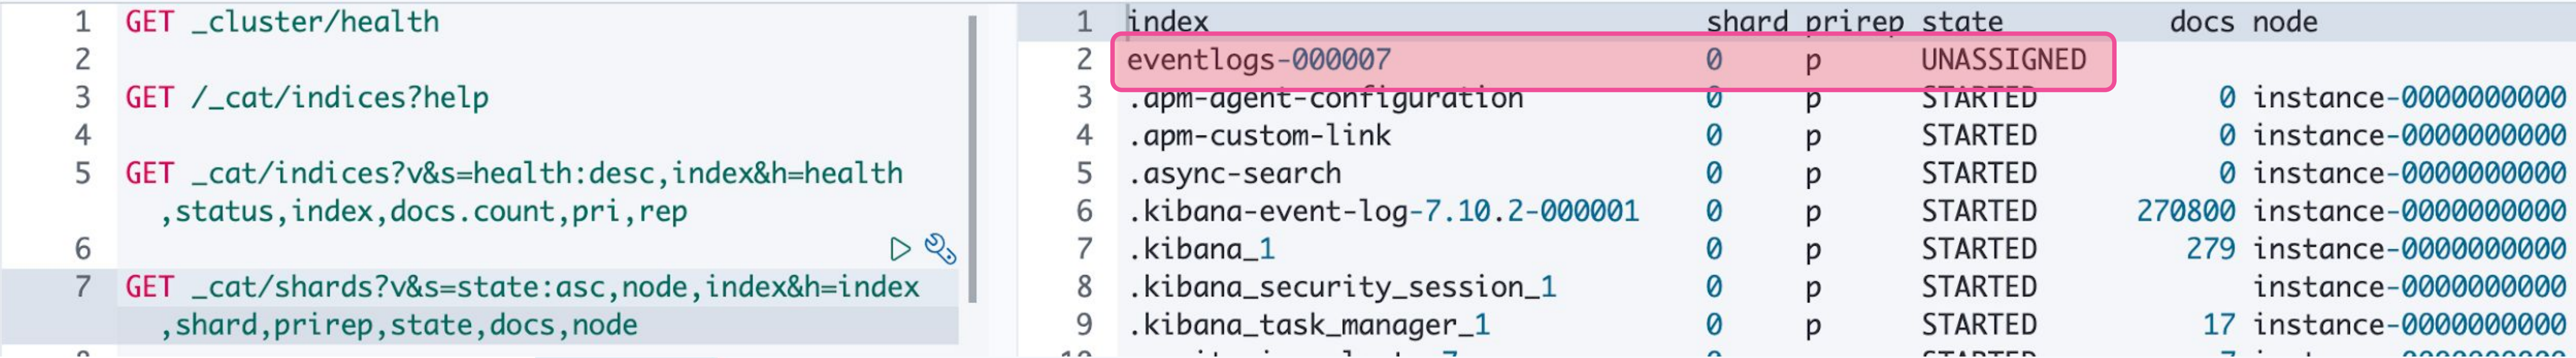

```
GET _cluster/health

GET /_cat/indices?help

GET _cat/indices?v&s=health:desc,index&h=health
  ,status,index,docs.count,pri,rep

GET _cat/shards?v&s=state:asc,node,index&h=index
  ,shard,prirep,state,docs,node
```

```
index                               shard prirep state        docs node
eventlogs-000007                    0     p      UNASSIGNED
.apm-agent-configuration            0     p      STARTED         0 instance-0000000000
.apm-custom-link                    0     p      STARTED         0 instance-0000000000
.async-search                       0     p      STARTED         0 instance-0000000000
.kibana-event-log-7.10.2-000001     0     p      STARTED    270800 instance-0000000000
.kibana_1                           0     p      STARTED       279 instance-0000000000
.kibana_security_session_1          0     p      STARTED           instance-0000000000
.kibana_task_manager_1              0     p      STARTED        17 instance-0000000000
```

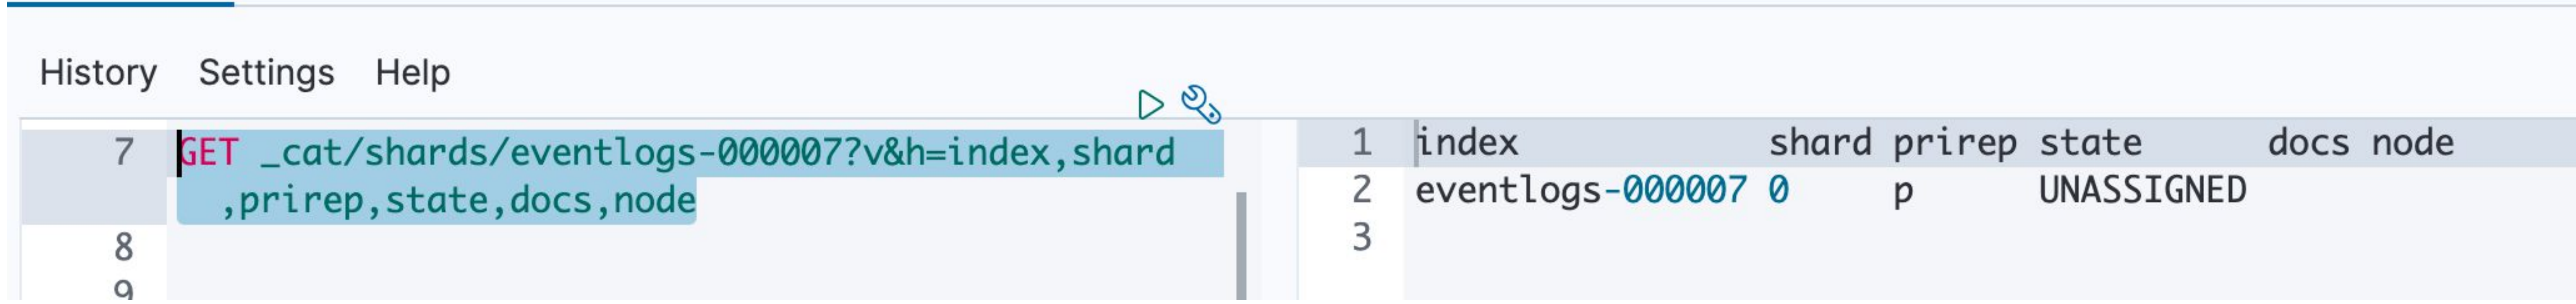

```
GET _cat/shards/eventlogs-000007?v&h=index,shard
  ,prirep,state,docs,node
```

```
index            shard prirep state      docs node
eventlogs-000007 0     p      UNASSIGNED
```

DIAGNOSTIC

[Urgent Severity] **Red cluster**

Why is a shard UNASSIGNED?

→ Cluster allocation explain API

https://www.elastic.co/guide/en/elasticsearch/reference/current/cluster-allocation-explain.html

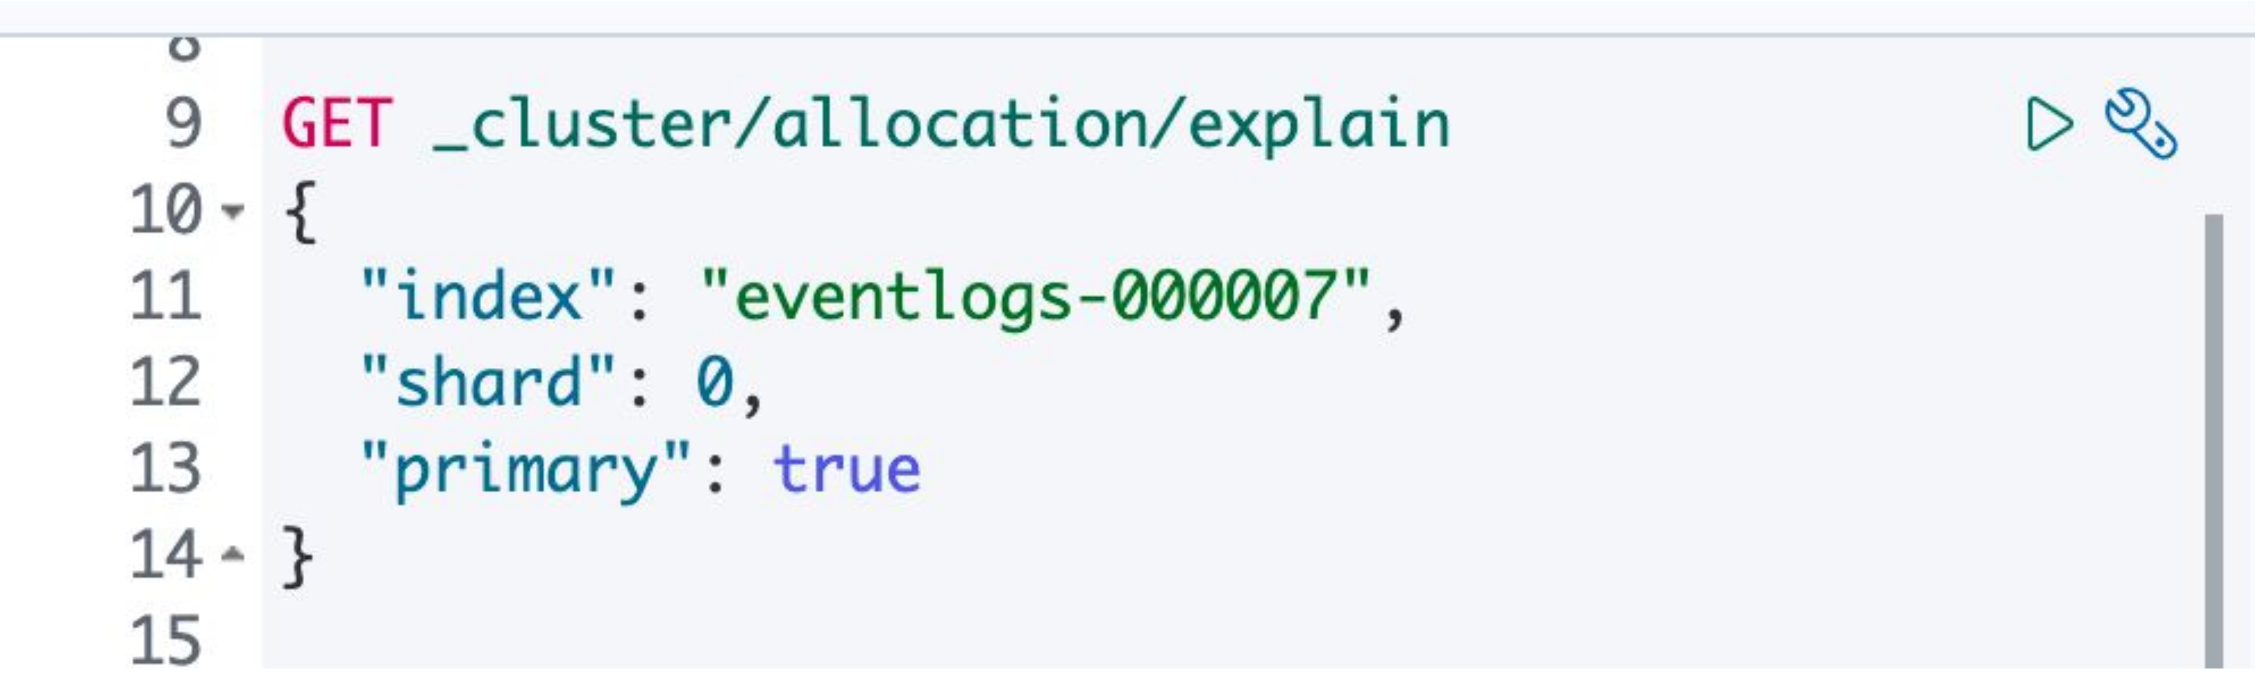

```
GET _cluster/allocation/explain
{
  "index": "eventlogs-000007",
  "shard": 0,
  "primary": true
}
```

DIAGNOSTIC

[Urgent Severity] **Red cluster**

Why is a shard UNASSIGNED?

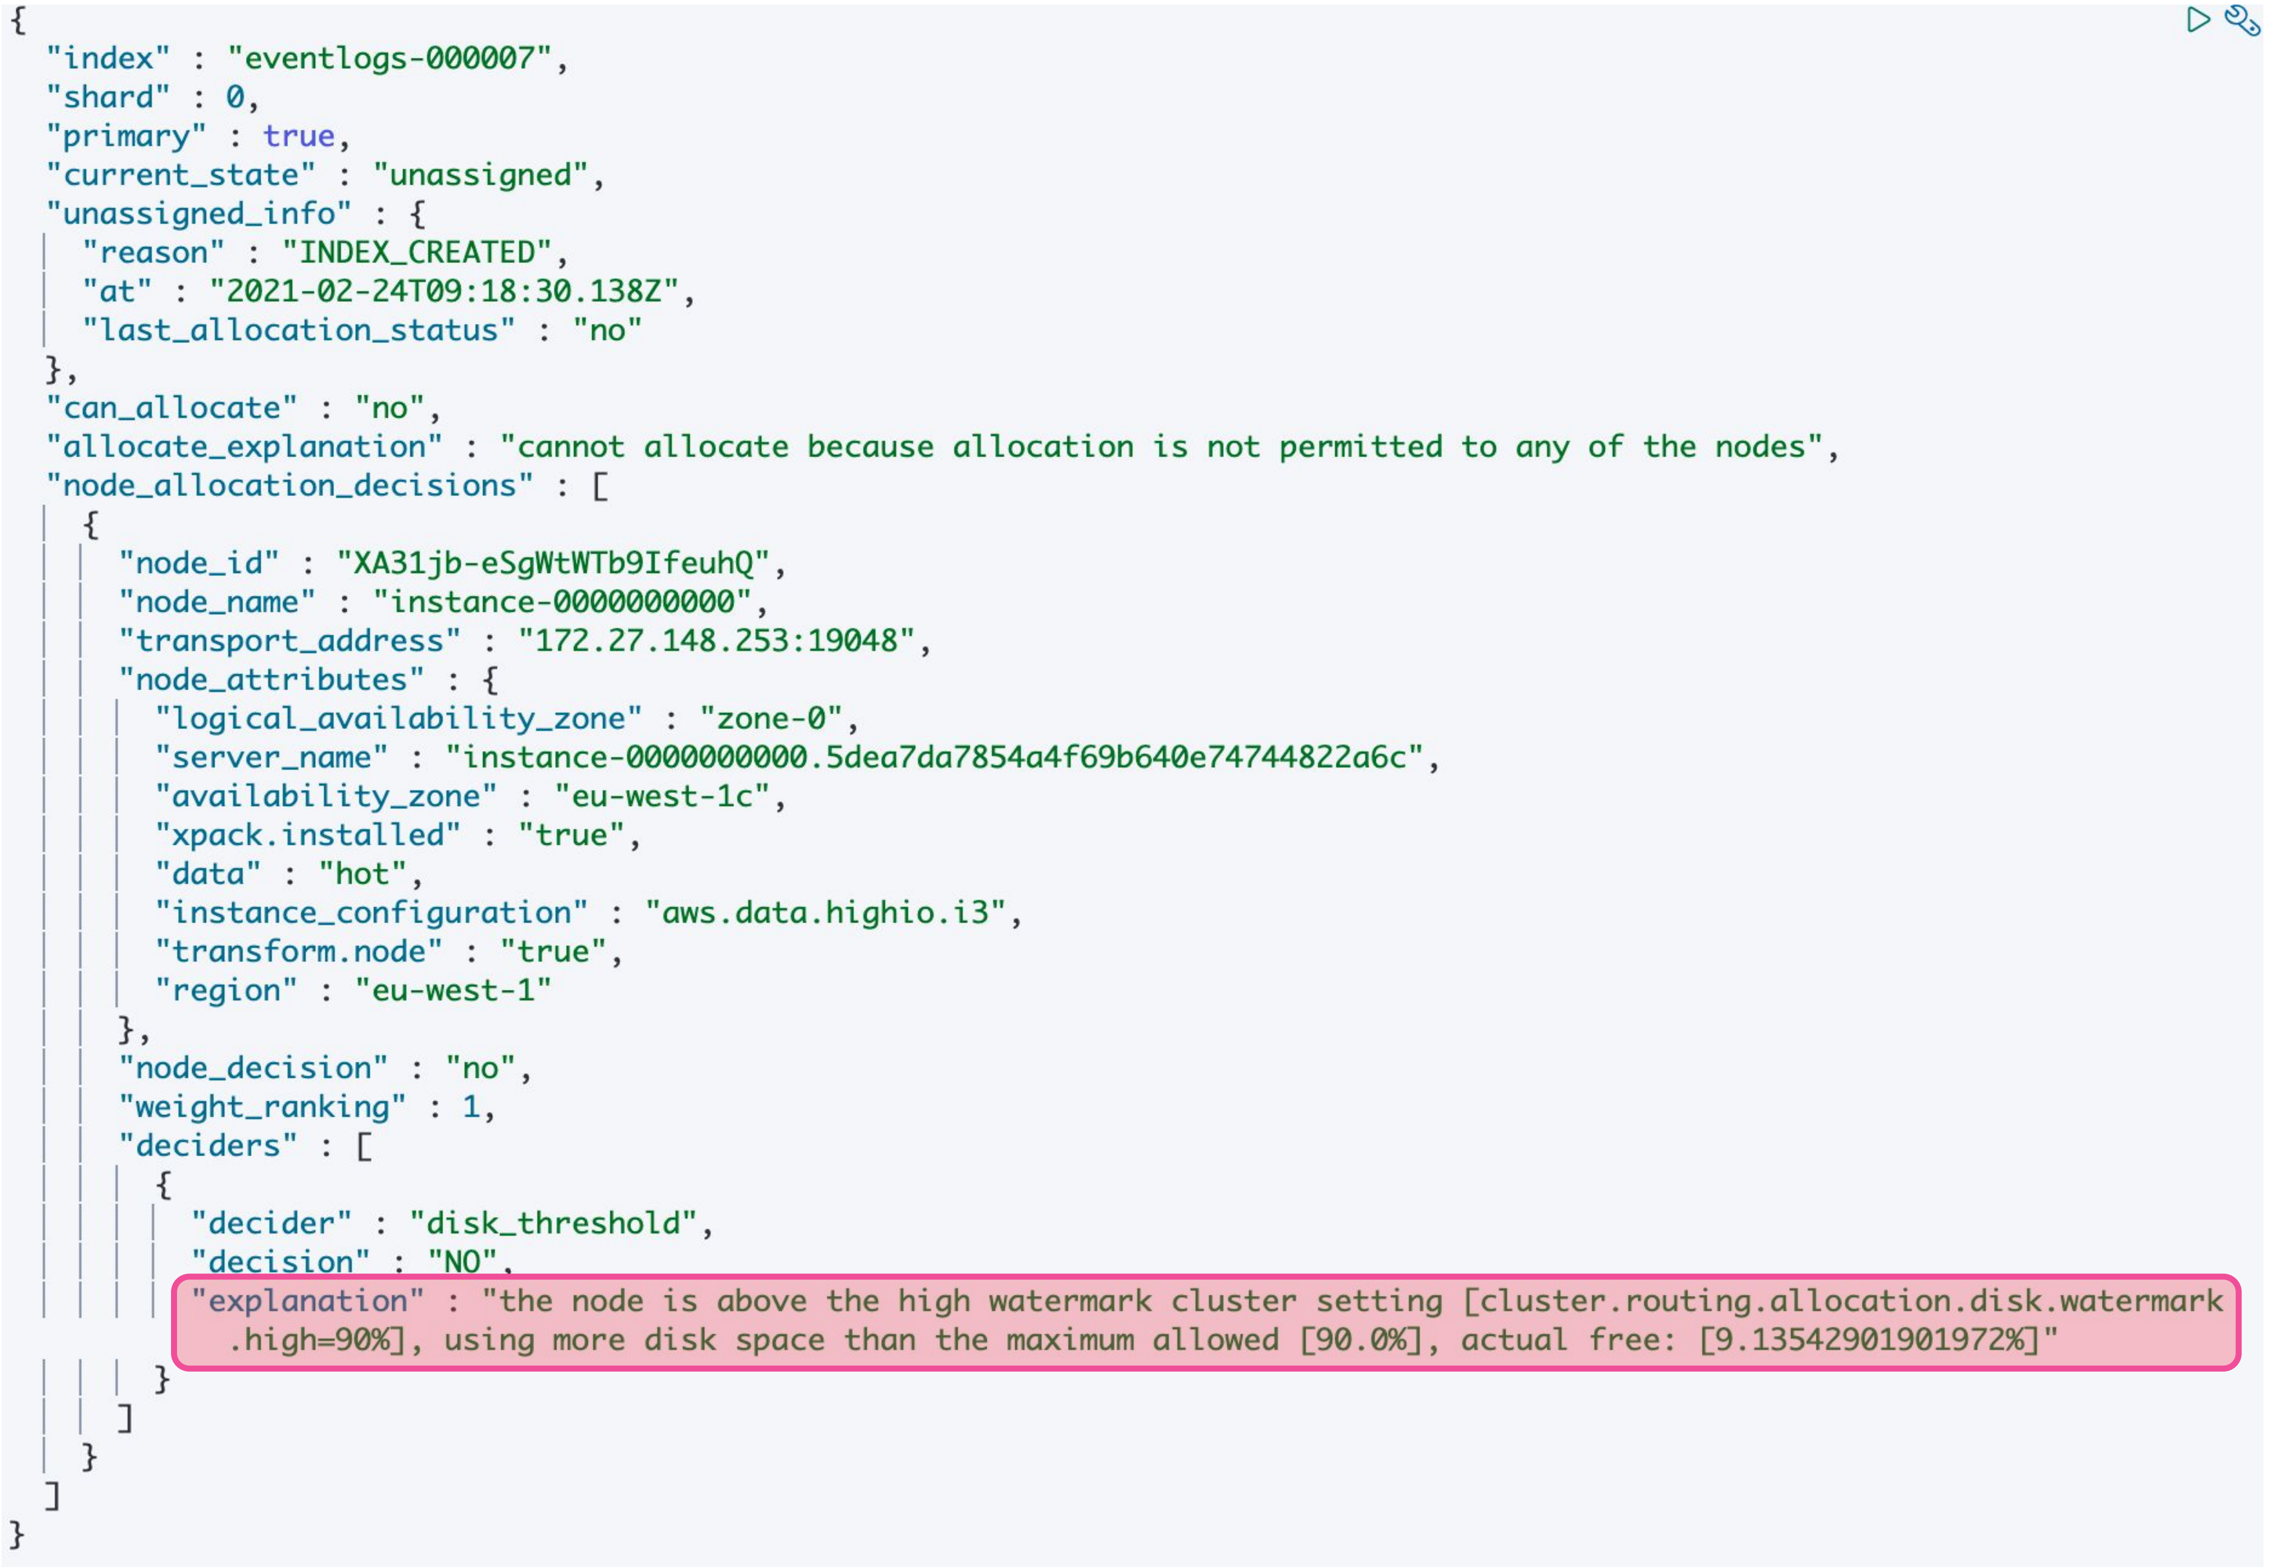

```
{
  "index" : "eventlogs-000007",
  "shard" : 0,
  "primary" : true,
  "current_state" : "unassigned",
  "unassigned_info" : {
    "reason" : "INDEX_CREATED",
    "at" : "2021-02-24T09:18:30.138Z",
    "last_allocation_status" : "no"
  },
  "can_allocate" : "no",
  "allocate_explanation" : "cannot allocate because allocation is not permitted to any of the nodes",
  "node_allocation_decisions" : [
    {
      "node_id" : "XA31jb-eSgWtWTb9IfeuhQ",
      "node_name" : "instance-0000000000",
      "transport_address" : "172.27.148.253:19048",
      "node_attributes" : {
        "logical_availability_zone" : "zone-0",
        "server_name" : "instance-0000000000.5dea7da7854a4f69b640e74744822a6c",
        "availability_zone" : "eu-west-1c",
        "xpack.installed" : "true",
        "data" : "hot",
        "instance_configuration" : "aws.data.highio.i3",
        "transform.node" : "true",
        "region" : "eu-west-1"
      },
      "node_decision" : "no",
      "weight_ranking" : 1,
      "deciders" : [
        {
          "decider" : "disk_threshold",
          "decision" : "NO",
          "explanation" : "the node is above the high watermark cluster setting [cluster.routing.allocation.disk.watermark
.high=90%], using more disk space than the maximum allowed [90.0%], actual free: [9.13542901901972%]"
        }
      ]
    }
  ]
}
```

DIAGNOSTIC

[Urgent Severity] **Red cluster**

Have we used all the cluster storage?

→ Use CAT Allocation API

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cat-allocation.html

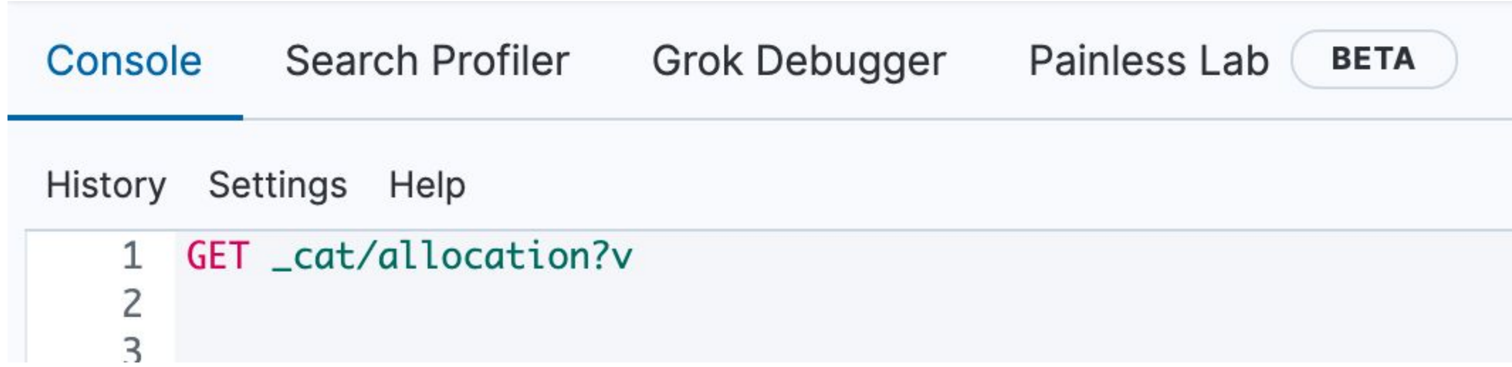

```
GET _cat/allocation?v
```

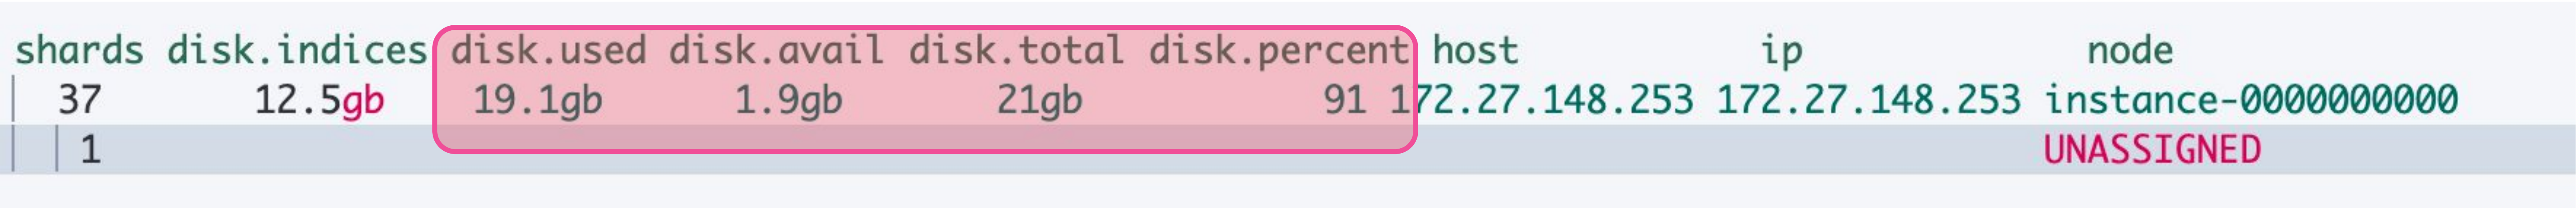

```
shards disk.indices disk.used disk.avail disk.total disk.percent host           ip             node
    37       12.5gb    19.1gb      1.9gb       21gb           91 172.27.148.253 172.27.148.253 instance-0000000000
     1                                                                                         UNASSIGNED
```

elastic

DIAGNOSTIC

[Urgent Severity] **Red cluster**

Interpret data

→ Cluster reached its disk high watermark

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/modules-cluster.html#disk-based-shard-allocation

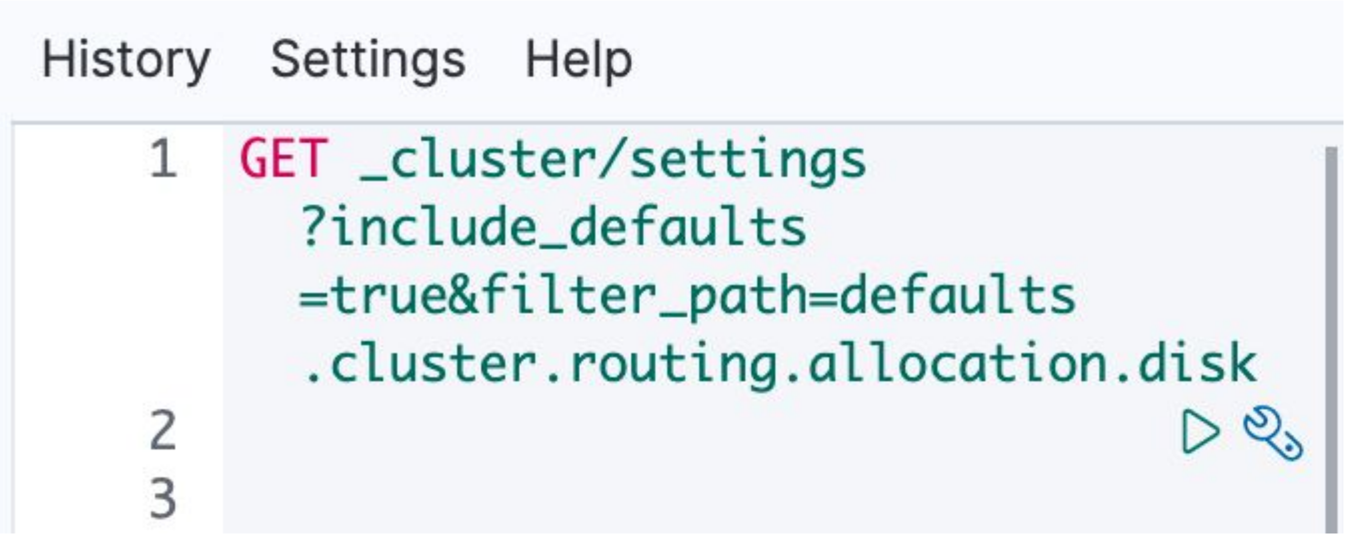

```
GET _cluster/settings
  ?include_defaults
  =true&filter_path=defaults
  .cluster.routing.allocation.disk
```

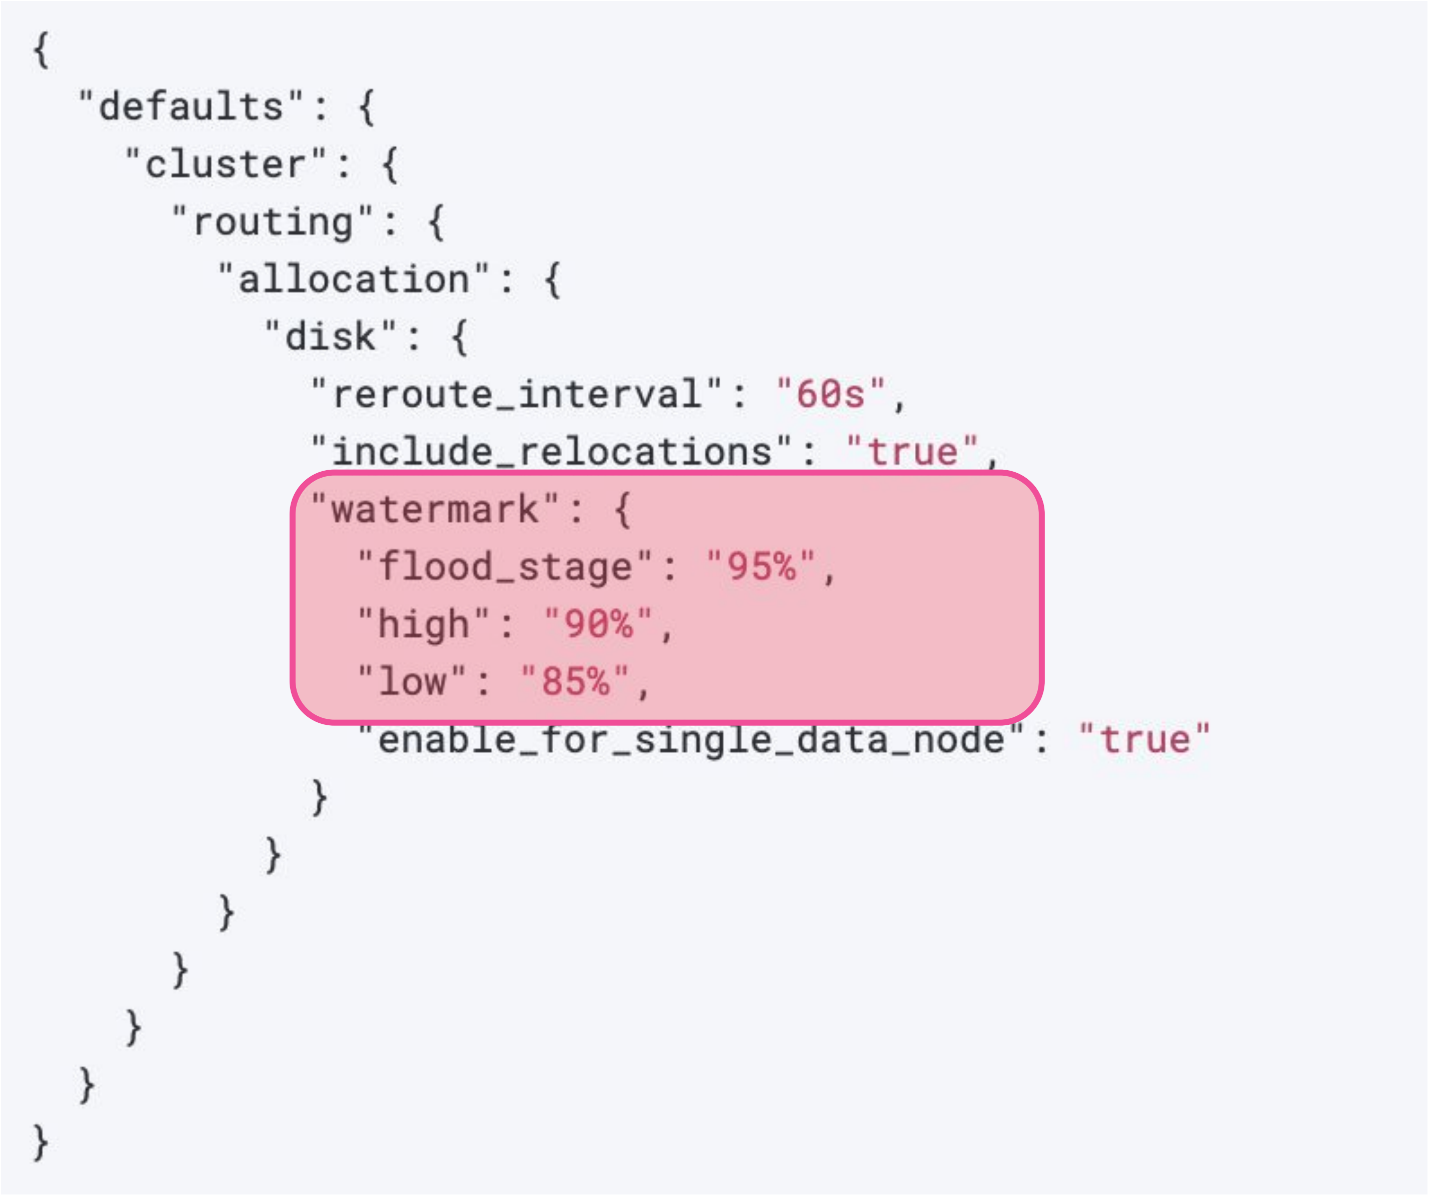

```
{
  "defaults": {
    "cluster": {
      "routing": {
        "allocation": {
          "disk": {
            "reroute_interval": "60s",
            "include_relocations": "true",
            "watermark": {
              "flood_stage": "95%",
              "high": "90%",
              "low": "85%",
              "enable_for_single_data_node": "true"
            }
          }
        }
      }
    }
  }
}
```

elastic

DIAGNOSTIC

[Urgent Severity] Red cluster

Interpret data

→ Existing indices are blocked for write

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cluster-get-settings.html

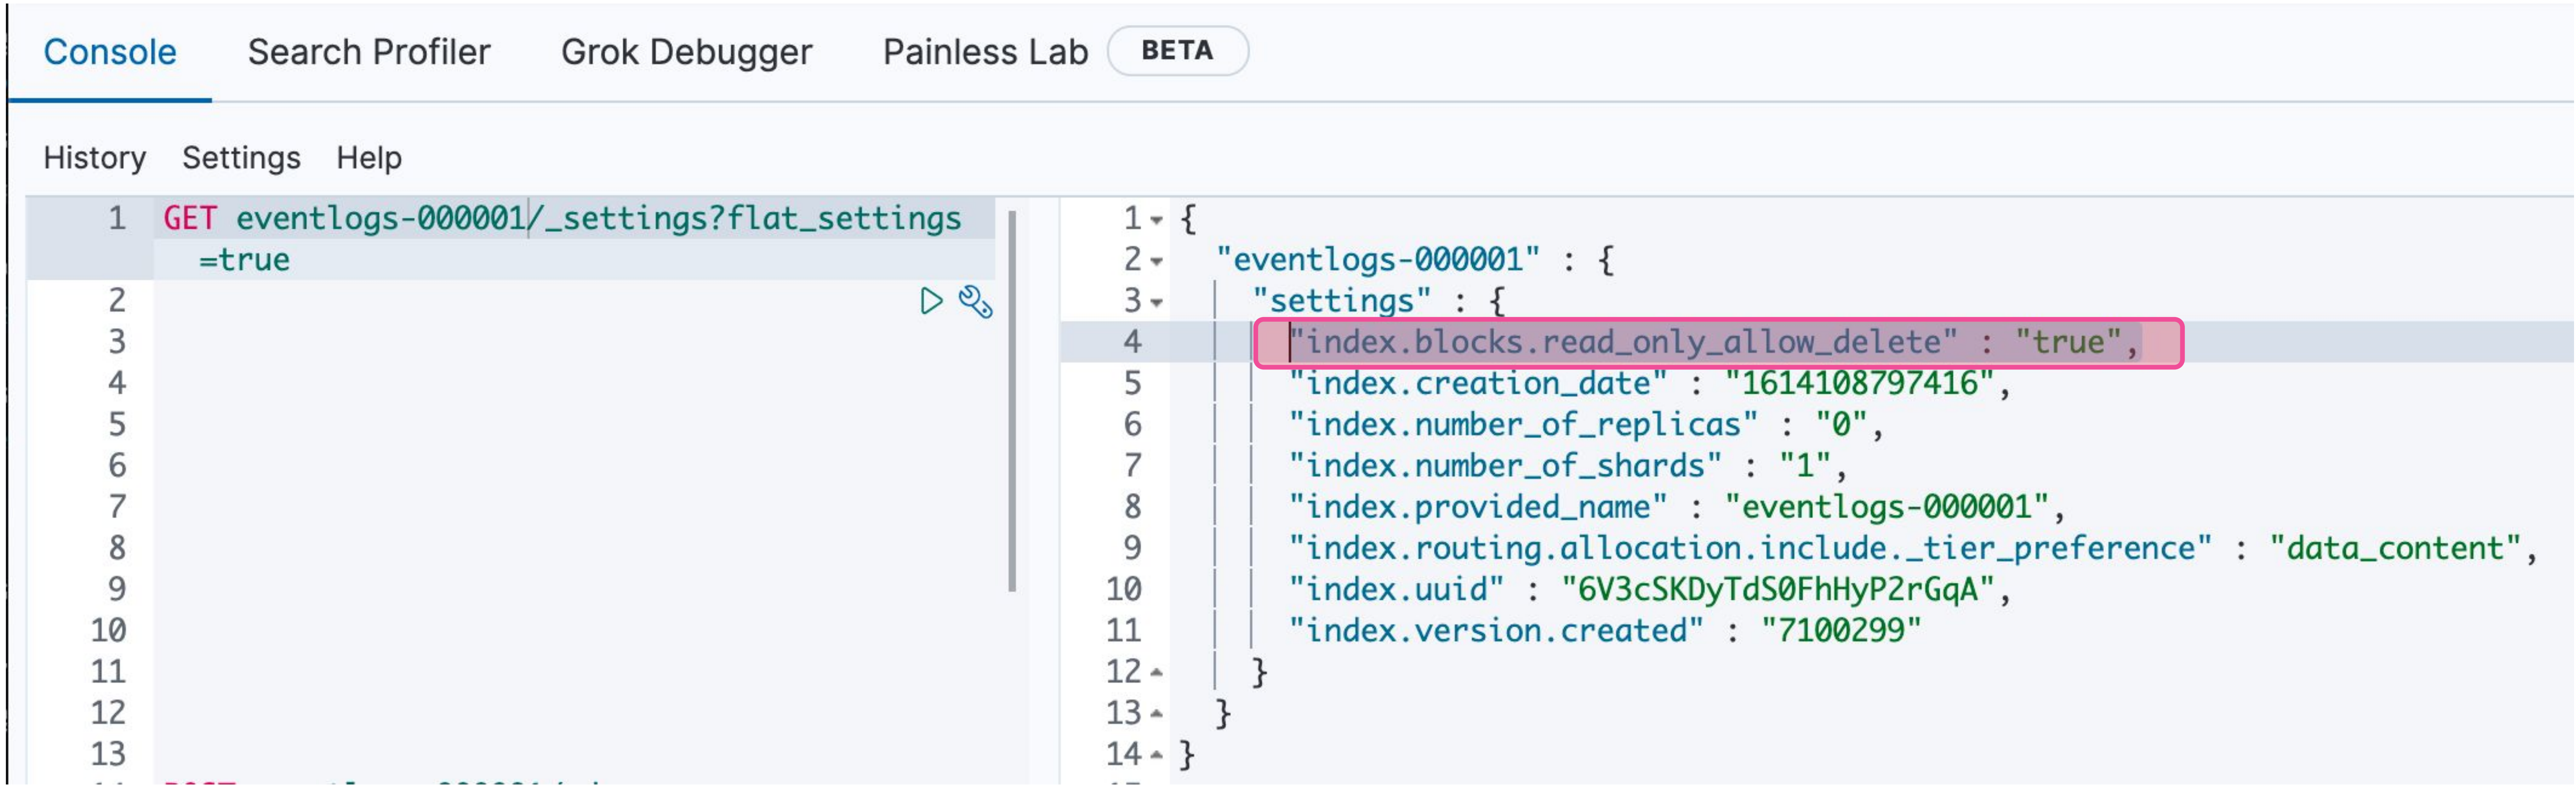

```
GET eventlogs-000001/_settings?flat_settings
  =true
```

```
{
  "eventlogs-000001" : {
    "settings" : {
      "index.blocks.read_only_allow_delete" : "true",
      "index.creation_date" : "1614108797416",
      "index.number_of_replicas" : "0",
      "index.number_of_shards" : "1",
      "index.provided_name" : "eventlogs-000001",
      "index.routing.allocation.include._tier_preference" : "data_content",
      "index.uuid" : "6V3cSKDyTdS0FhHyP2rGqA",
      "index.version.created" : "7100299"
    }
  }
}
```

TREATMENT

[Urgent Severity] Red **cluster**

Fixing the root cause

- **Delete indices** to increase available storage

  https://www.elastic.co/guide/en/elasticsearch/reference/7.11/indices-delete-index.html

  Do we have **snapshots**? We can restore later.

  https://www.elastic.co/guide/en/elasticsearch/reference/7.11/snapshot-restore.html

- **Add nodes** or increase storage capacity (easier on cloud)

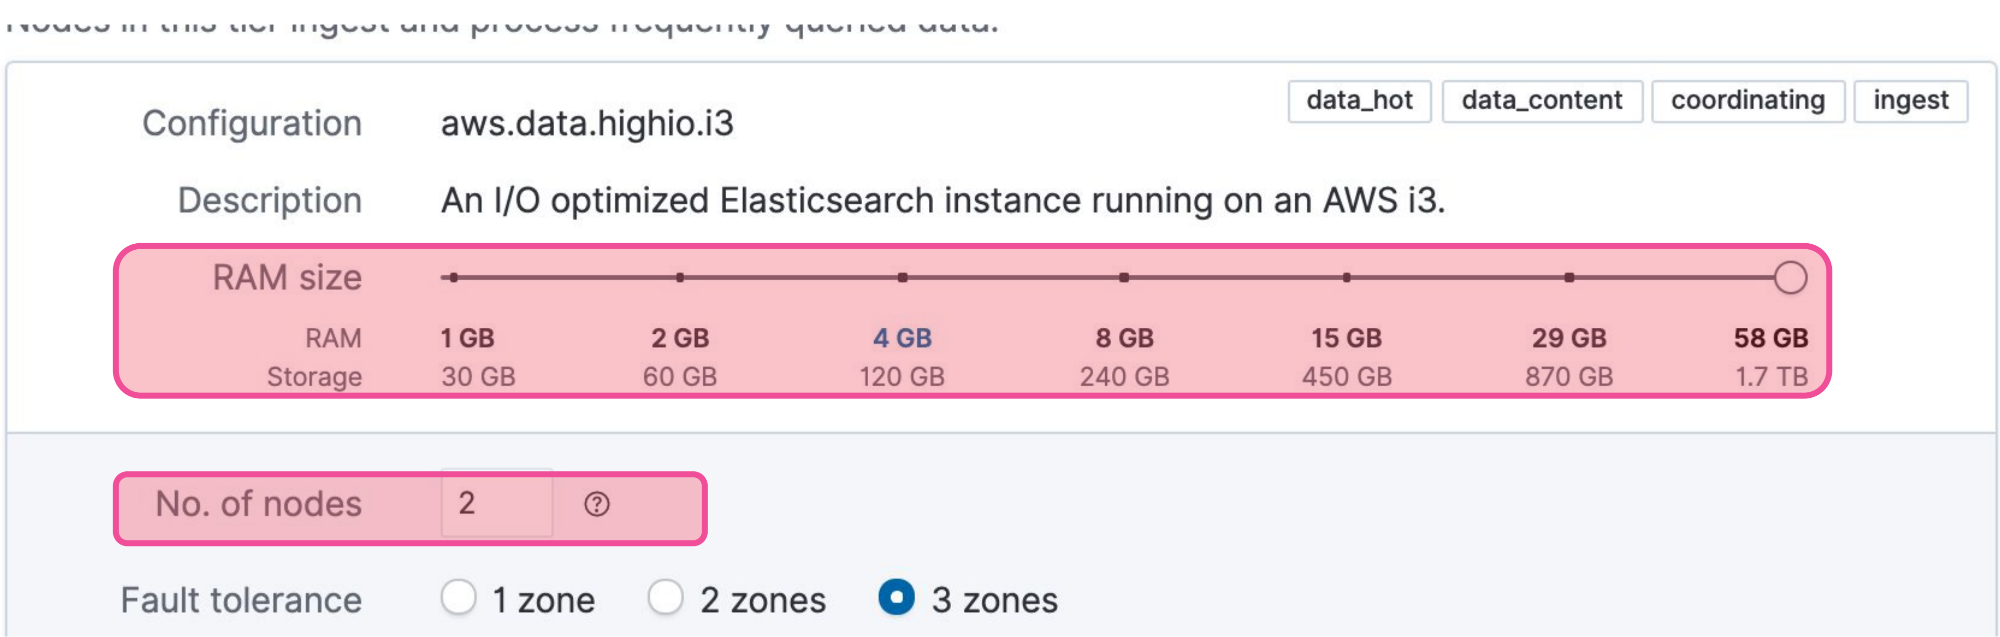

Nodes in this tier ingest and process frequently queried data.

Configuration aws.data.highio.i3

Description An I/O optimized Elasticsearch instance running on an AWS i3.

RAM size

RAM	1 GB	2 GB	4 GB	8 GB	15 GB	29 GB	58 GB
Storage	30 GB	60 GB	120 GB	240 GB	450 GB	870 GB	1.7 TB

No. of nodes 2

Fault tolerance ○ 1 zone ○ 2 zones ● 3 zones

TREATMENT

[Urgent Severity] **Red cluster**

Temporary Hotfix

→ Alter the cluster settings to temporarily allow a higher disk usage

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cluster-update-settings.html

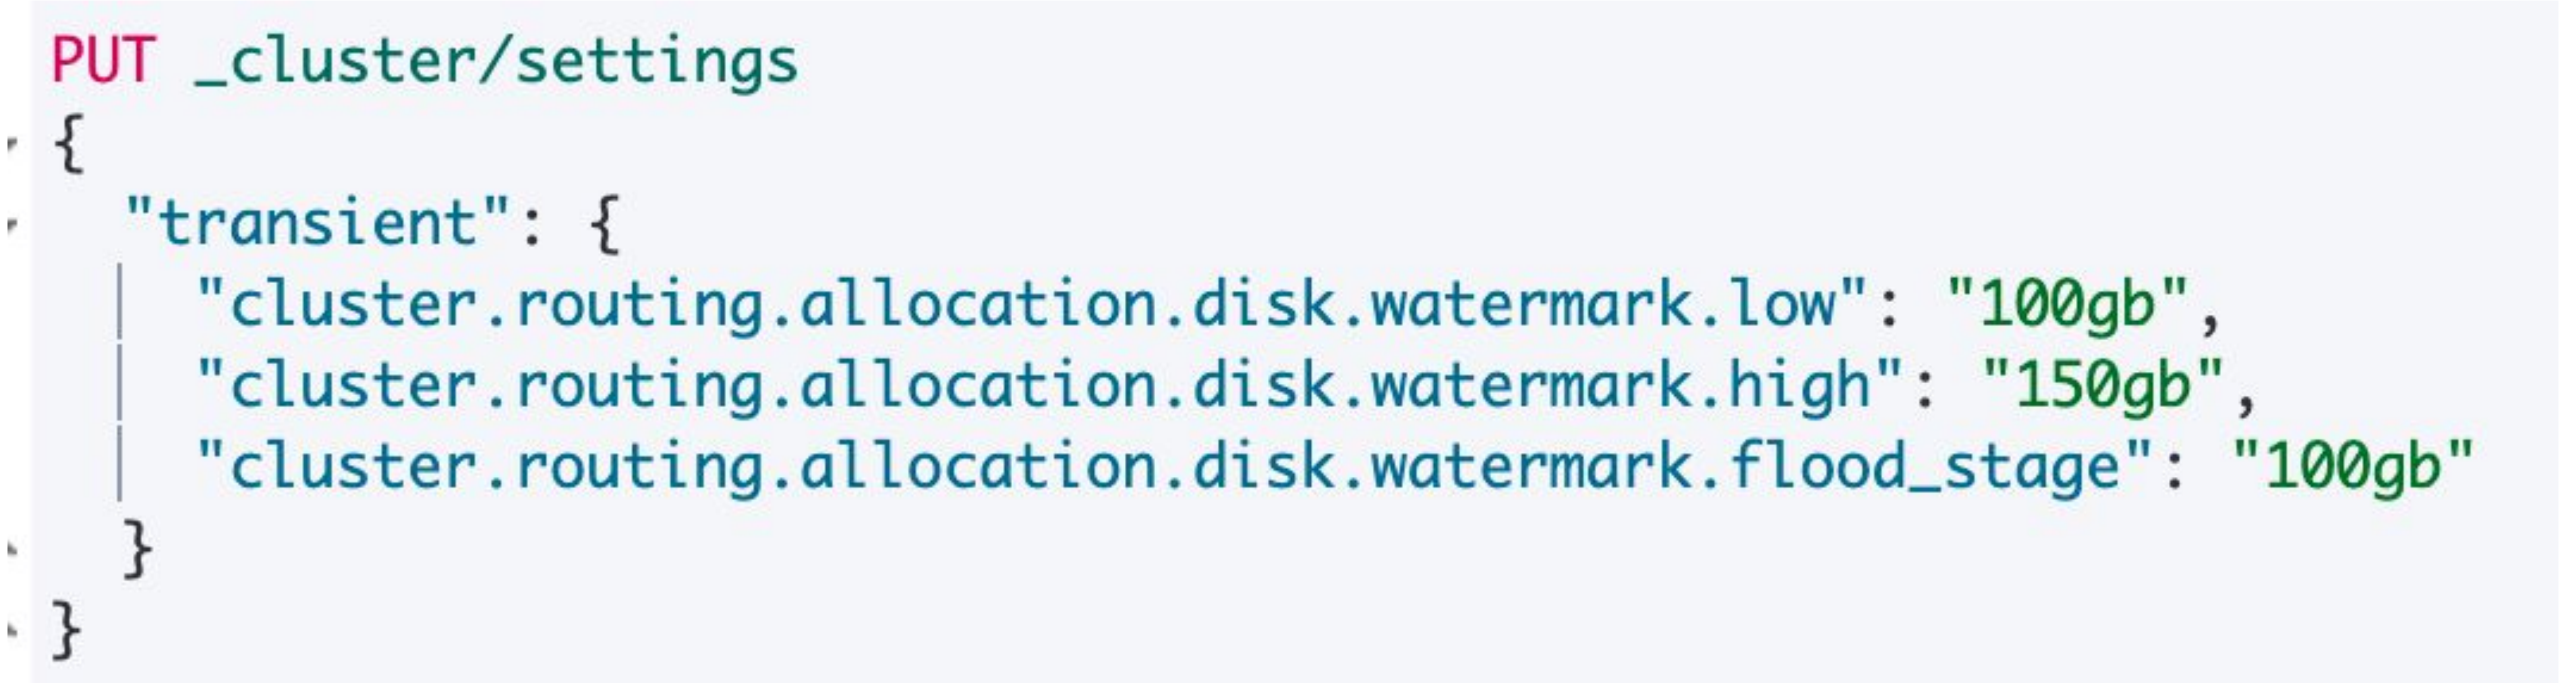

```
PUT _cluster/settings
{
  "transient": {
    "cluster.routing.allocation.disk.watermark.low": "100gb",
    "cluster.routing.allocation.disk.watermark.high": "150gb",
    "cluster.routing.allocation.disk.watermark.flood_stage": "100gb"
  }
}
```

TREATMENT

[Urgent Severity] **Red cluster**

Remove write block on the indices

→ Once we have enough disk, remove the index block if needed

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/indices-update-settings.html

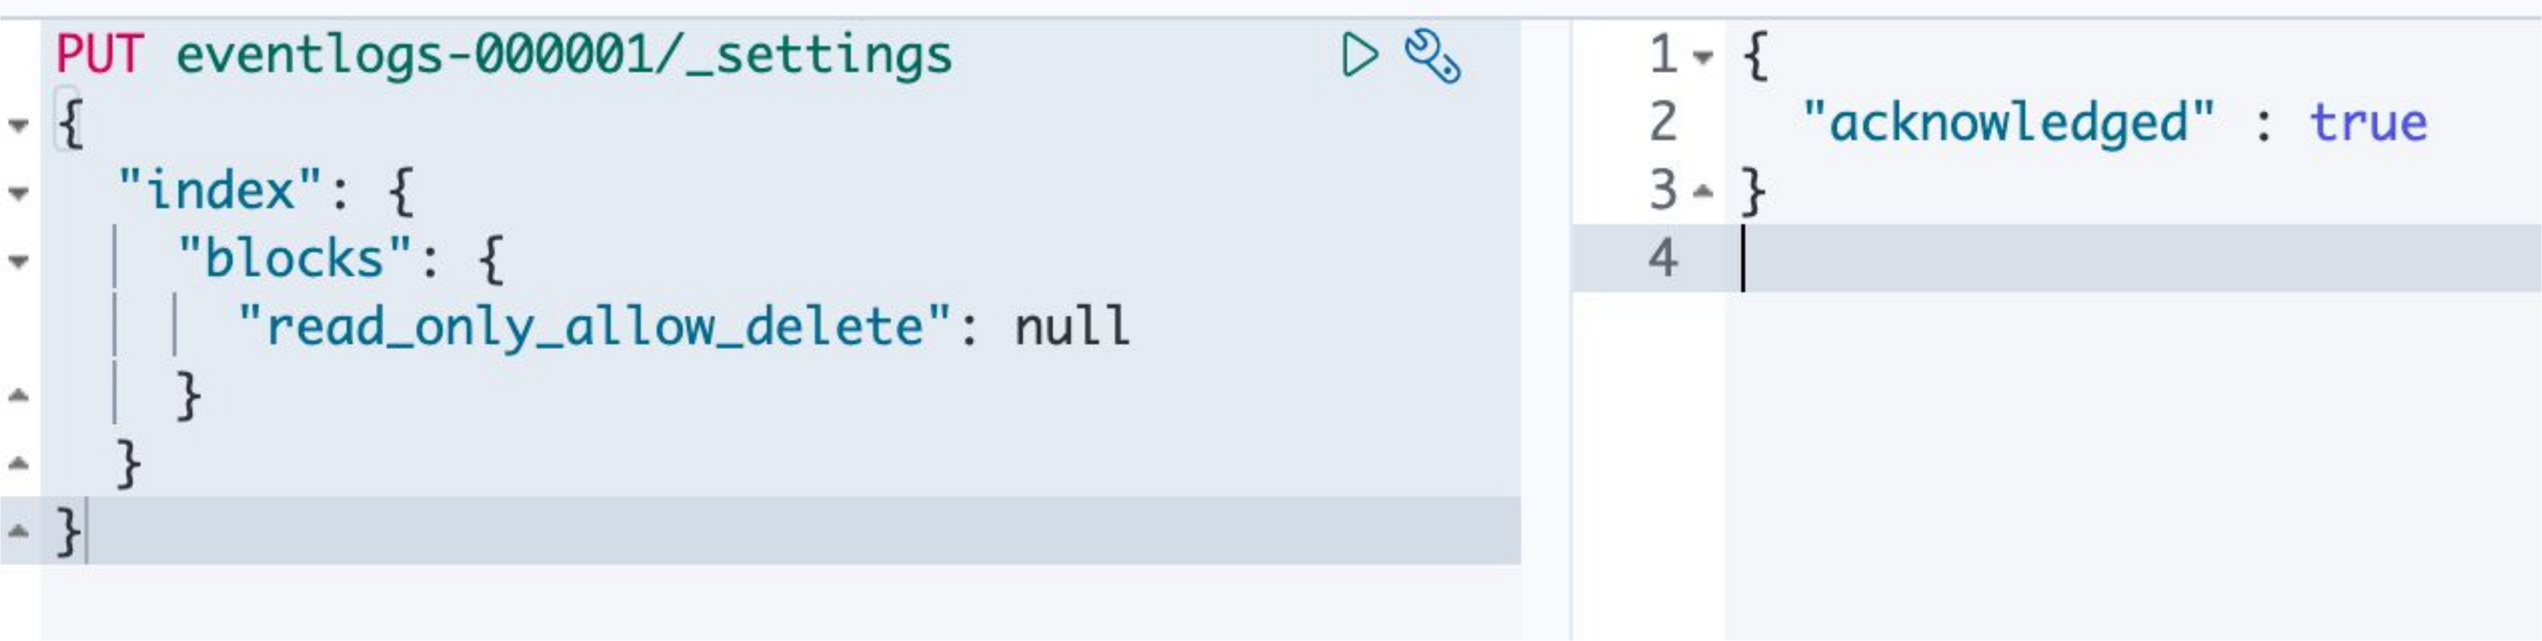

```
PUT eventlogs-000001/_settings
{
  "index": {
    "blocks": {
      "read_only_allow_delete": null
    }
  }
}
```

```
{
  "acknowledged" : true
}
```

TREATMENT

[Urgent Severity] **Red cluster**

Bonus track

→ If corrupted shards, and no snapshots, we can force allocation accepting potential data loss

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cluster-reroute.html#cluster-reroute-api-request-body

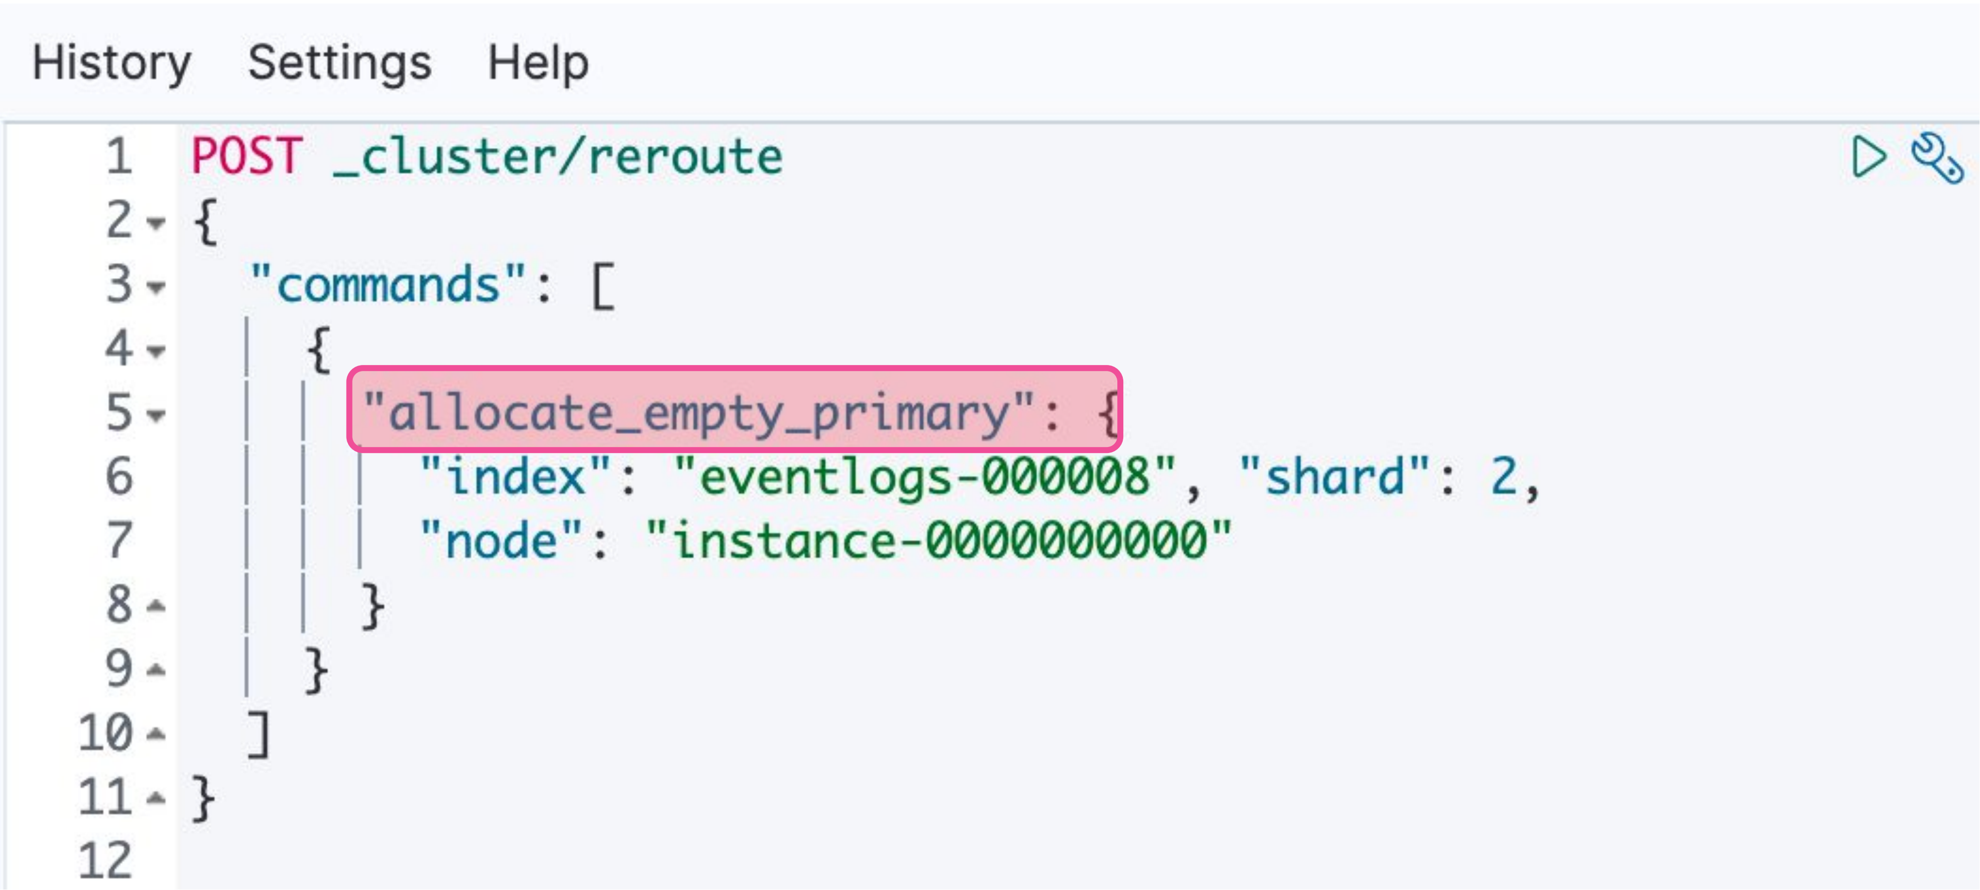

```
POST _cluster/reroute
{
  "commands": [
    {
      "allocate_empty_primary": {
        "index": "eventlogs-000008", "shard": 2,
        "node": "instance-0000000000"
      }
    }
  ]
}
```

DISCHARGE

[Urgent Severity] **Red cluster**

Takeaways

→ Proactively **monitor disk usage** on each node / **Alerts**

Aim to 75% used storage to be on the safe side (< 85%)

→ Plan for data retention / deletion with **ILM** or **Data Tiers**

Index Lifecycle Management (ILM) can help automate

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/index-lifecycle-management.html

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/data-tiers.html

→ Snapshot / Snapshot Lifecycle Management (**SLM**) for backups

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/snapshot-lifecycle-management.html

SUMMARY

[Urgent Severity] **Red cluster**

Triage

- → Cluster health is **red**
- → Stopped ingesting
- → Search works

Diagnostic

Reached high disk watermark

- → CAT APIs
- → Allocation Explain
- → Cluster and index settings

Support diagnostics

Treatment

- → Delete indices
- → Add data node/s
- → Update index settings / allow write

Discharge

- → Proactively monitor disk usage (alerts)
- → Snapshots
- → Index Lifecycle Management deletes old data and manages replicas
- → Data Tiers with Cold Tier

More Tools & Resources

Monitoring

→ Monitoring in production - dedicated cluster

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/monitoring-production.html

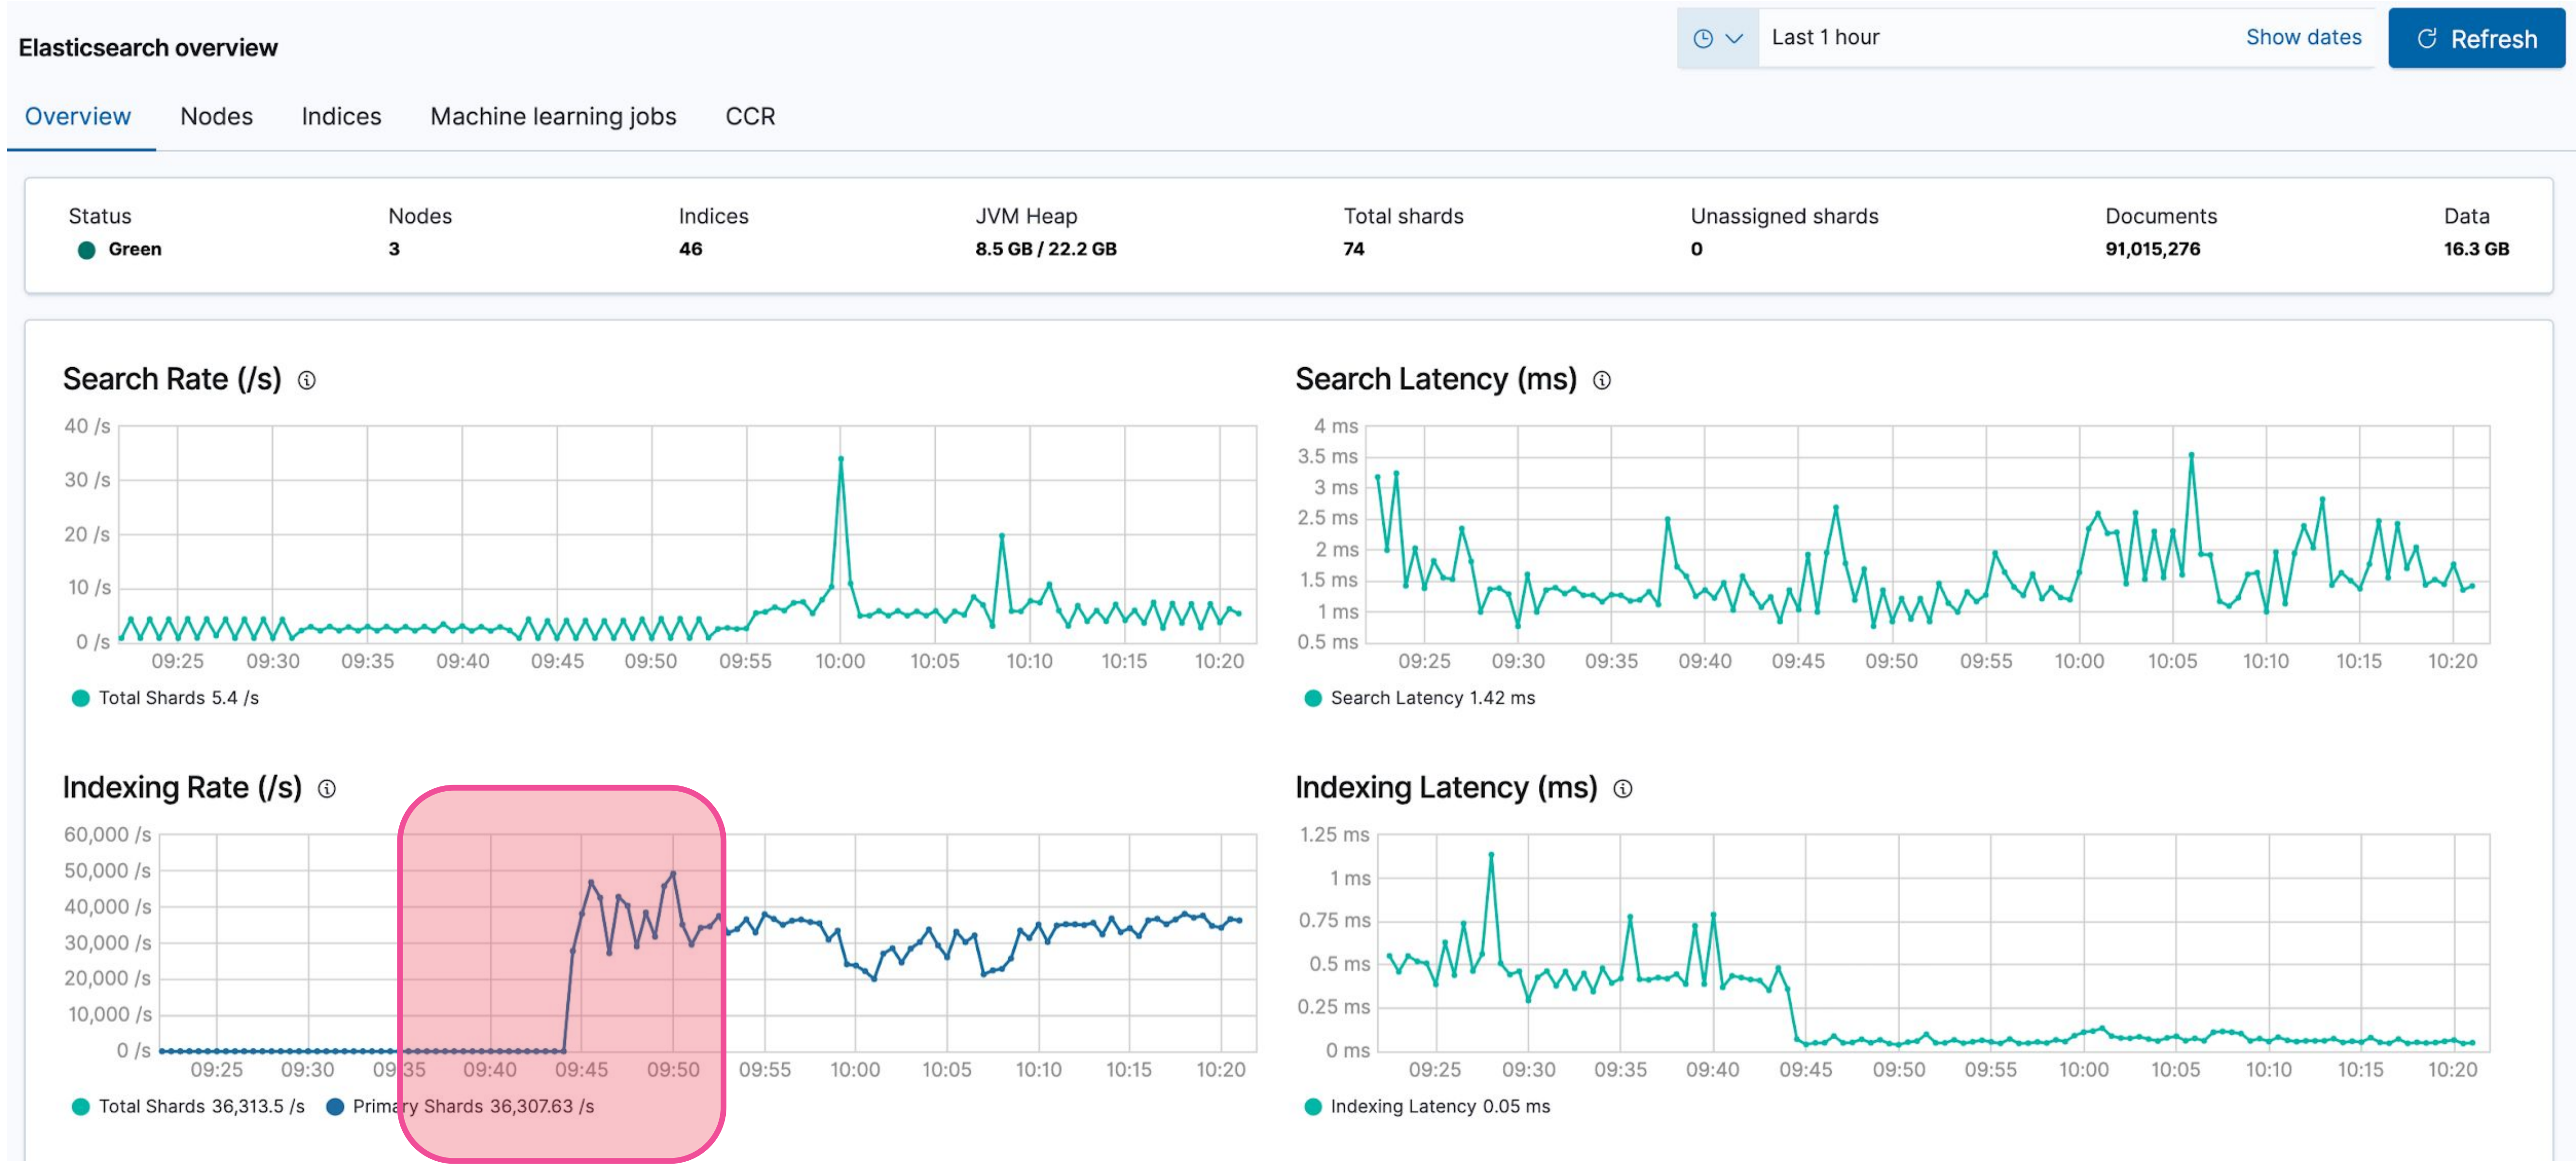

Elasticsearch overview

Last 1 hour | Show dates | Refresh

Overview | Nodes | Indices | Machine learning jobs | CCR

Status	Nodes	Indices	JVM Heap	Total shards	Unassigned shards	Documents	Data
● Green	3	46	8.5 GB / 22.2 GB	74	0	91,015,276	16.3 GB

Search Rate (/s)

● Total Shards 5.4 /s

Search Latency (ms)

● Search Latency 1.42 ms

Indexing Rate (/s)

● Total Shards 36,313.5 /s ● Primary Shards 36,307.63 /s

Indexing Latency (ms)

● Indexing Latency 0.05 ms

elastic

Monitoring

→ Nodes' memory usage

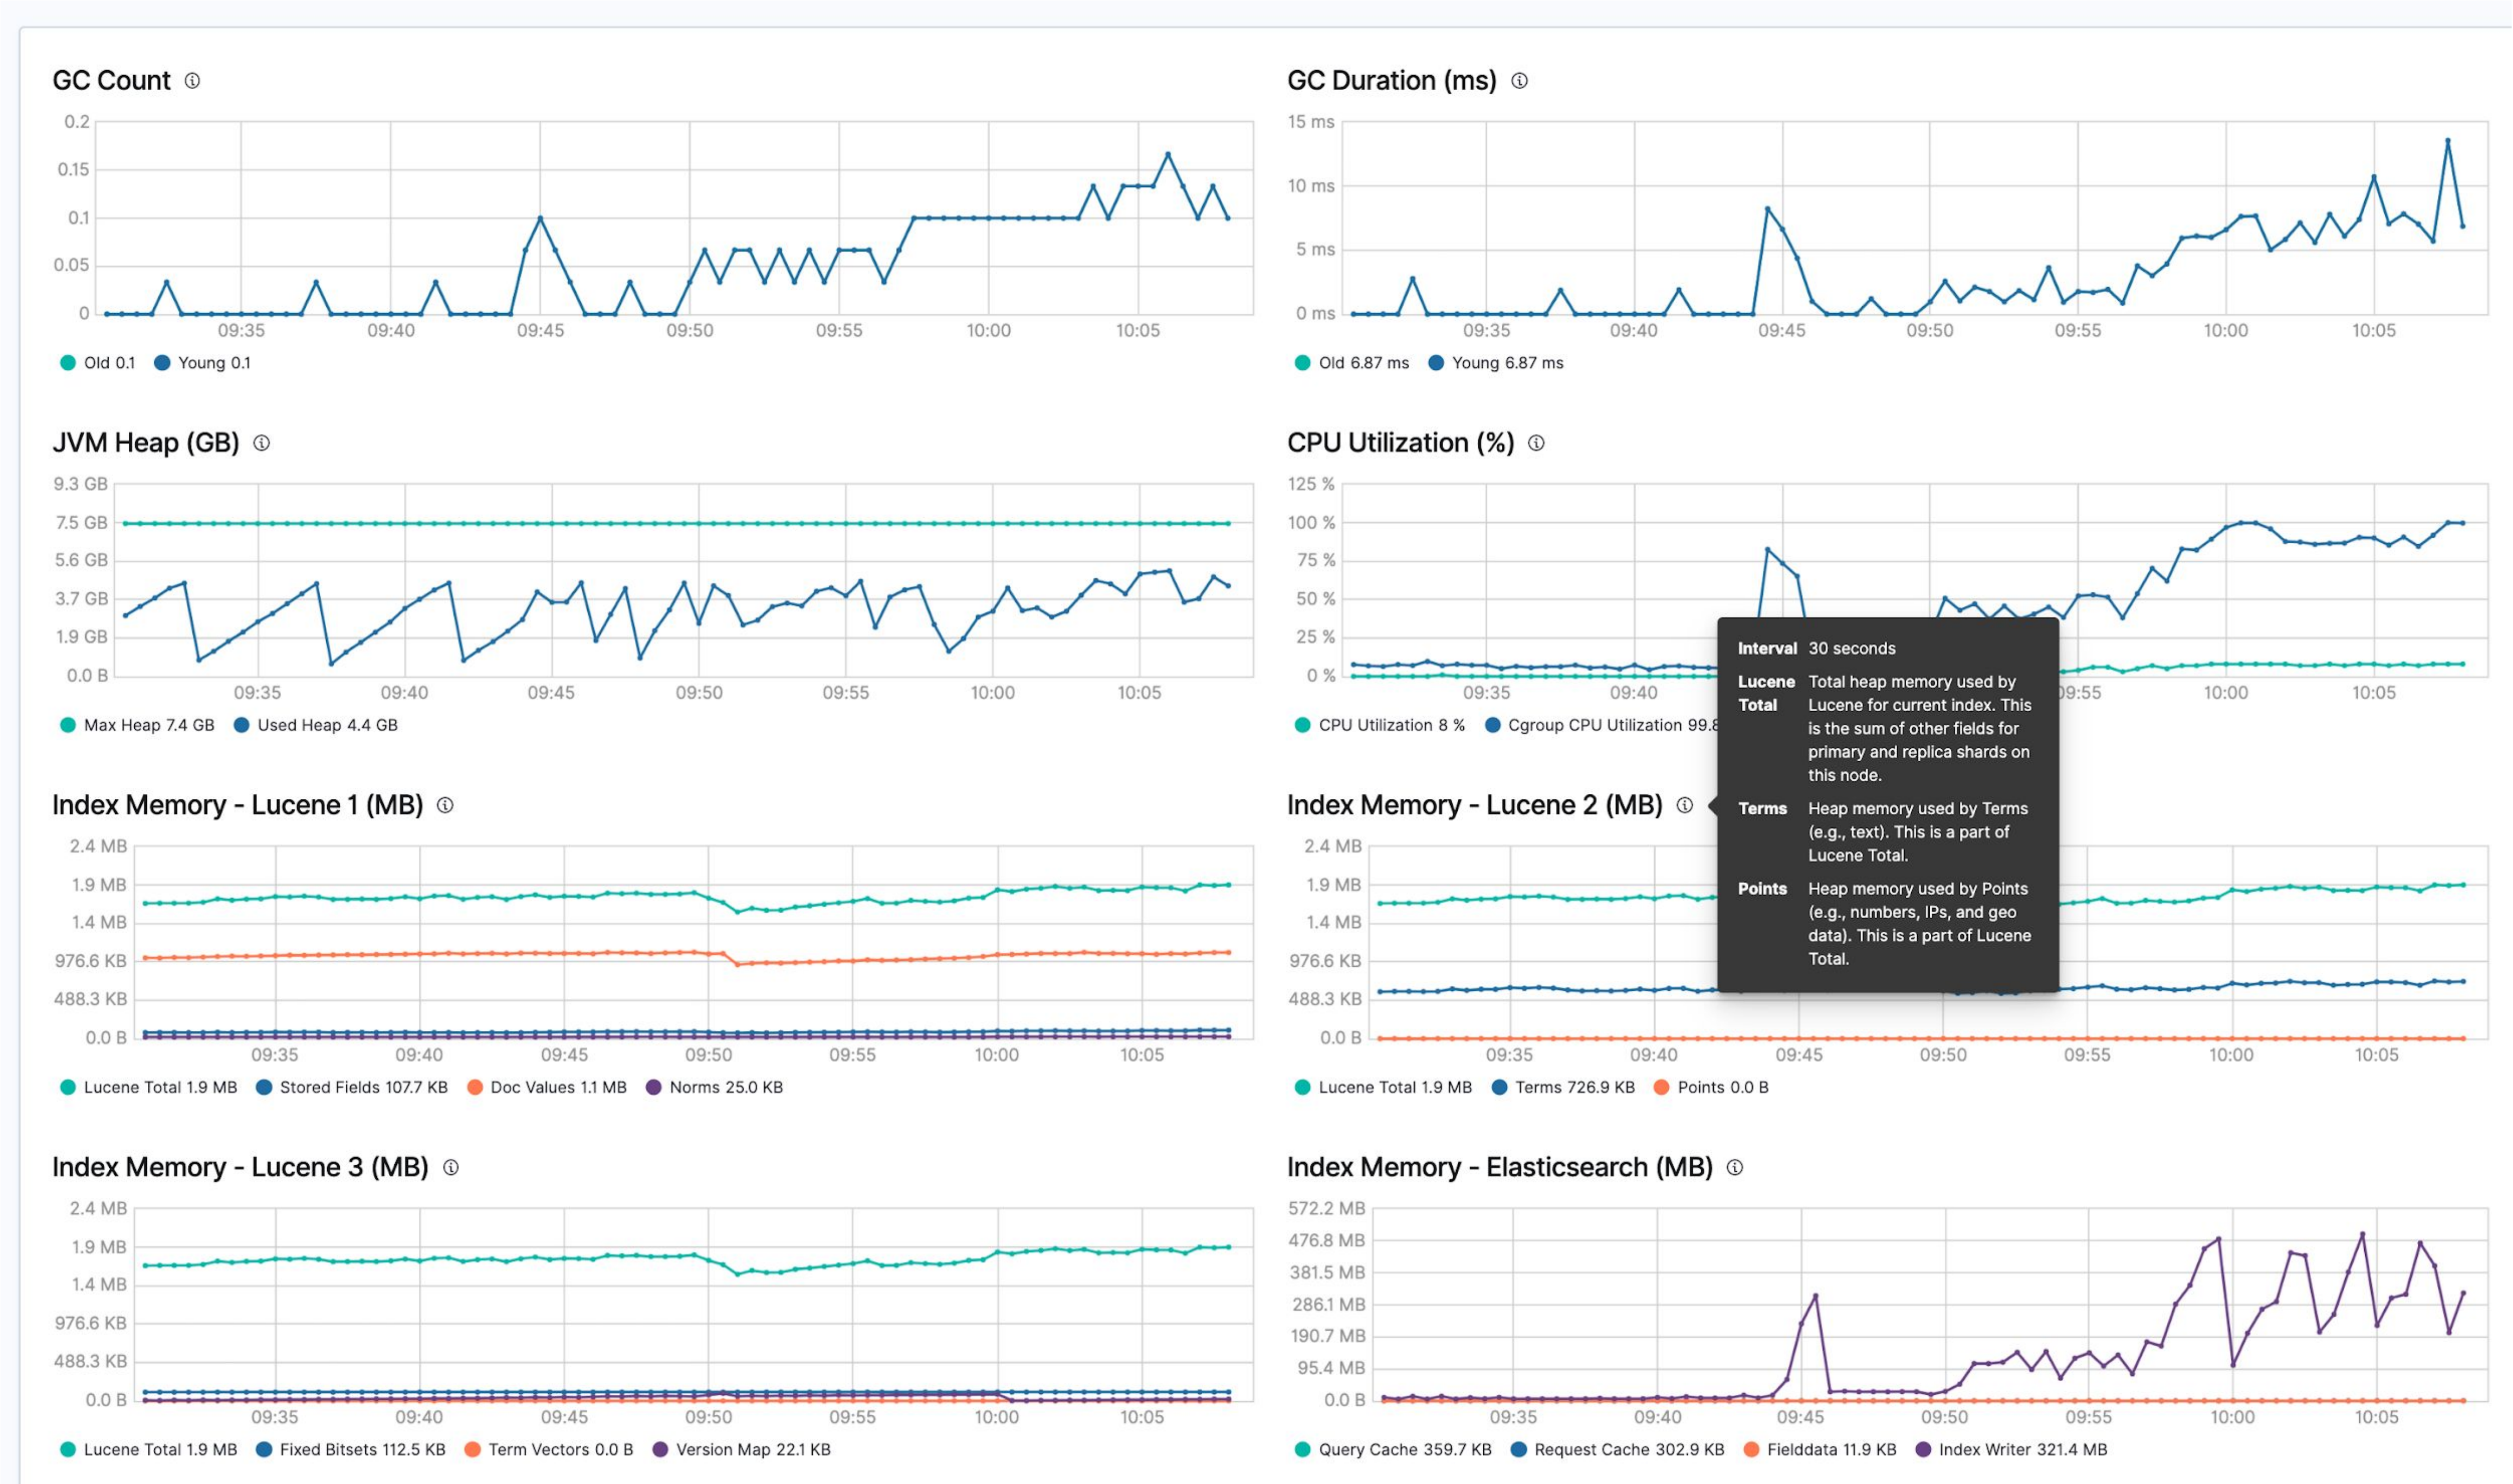

GC Count ⓘ

● Old 0.1 ● Young 0.1

GC Duration (ms) ⓘ

● Old 6.87 ms ● Young 6.87 ms

JVM Heap (GB) ⓘ

● Max Heap 7.4 GB ● Used Heap 4.4 GB

CPU Utilization (%) ⓘ

● CPU Utilization 8 % ● Cgroup CPU Utilization 99.8

Interval 30 seconds

Lucene Total: Total heap memory used by Lucene for current index. This is the sum of other fields for primary and replica shards on this node.

Terms: Heap memory used by Terms (e.g., text). This is a part of Lucene Total.

Points: Heap memory used by Points (e.g., numbers, IPs, and geo data). This is a part of Lucene Total.

Index Memory - Lucene 1 (MB) ⓘ

● Lucene Total 1.9 MB ● Stored Fields 107.7 KB ● Doc Values 1.1 MB ● Norms 25.0 KB

Index Memory - Lucene 2 (MB) ⓘ

● Lucene Total 1.9 MB ● Terms 726.9 KB ● Points 0.0 B

Index Memory - Lucene 3 (MB) ⓘ

● Lucene Total 1.9 MB ● Fixed Bitsets 112.5 KB ● Term Vectors 0.0 B ● Version Map 22.1 KB

Index Memory - Elasticsearch (MB) ⓘ

● Query Cache 359.7 KB ● Request Cache 302.9 KB ● Fielddata 11.9 KB ● Index Writer 321.4 MB

elastic

Monitoring

→ Ingest and Search queues and rejections

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cat-thread-pool.html

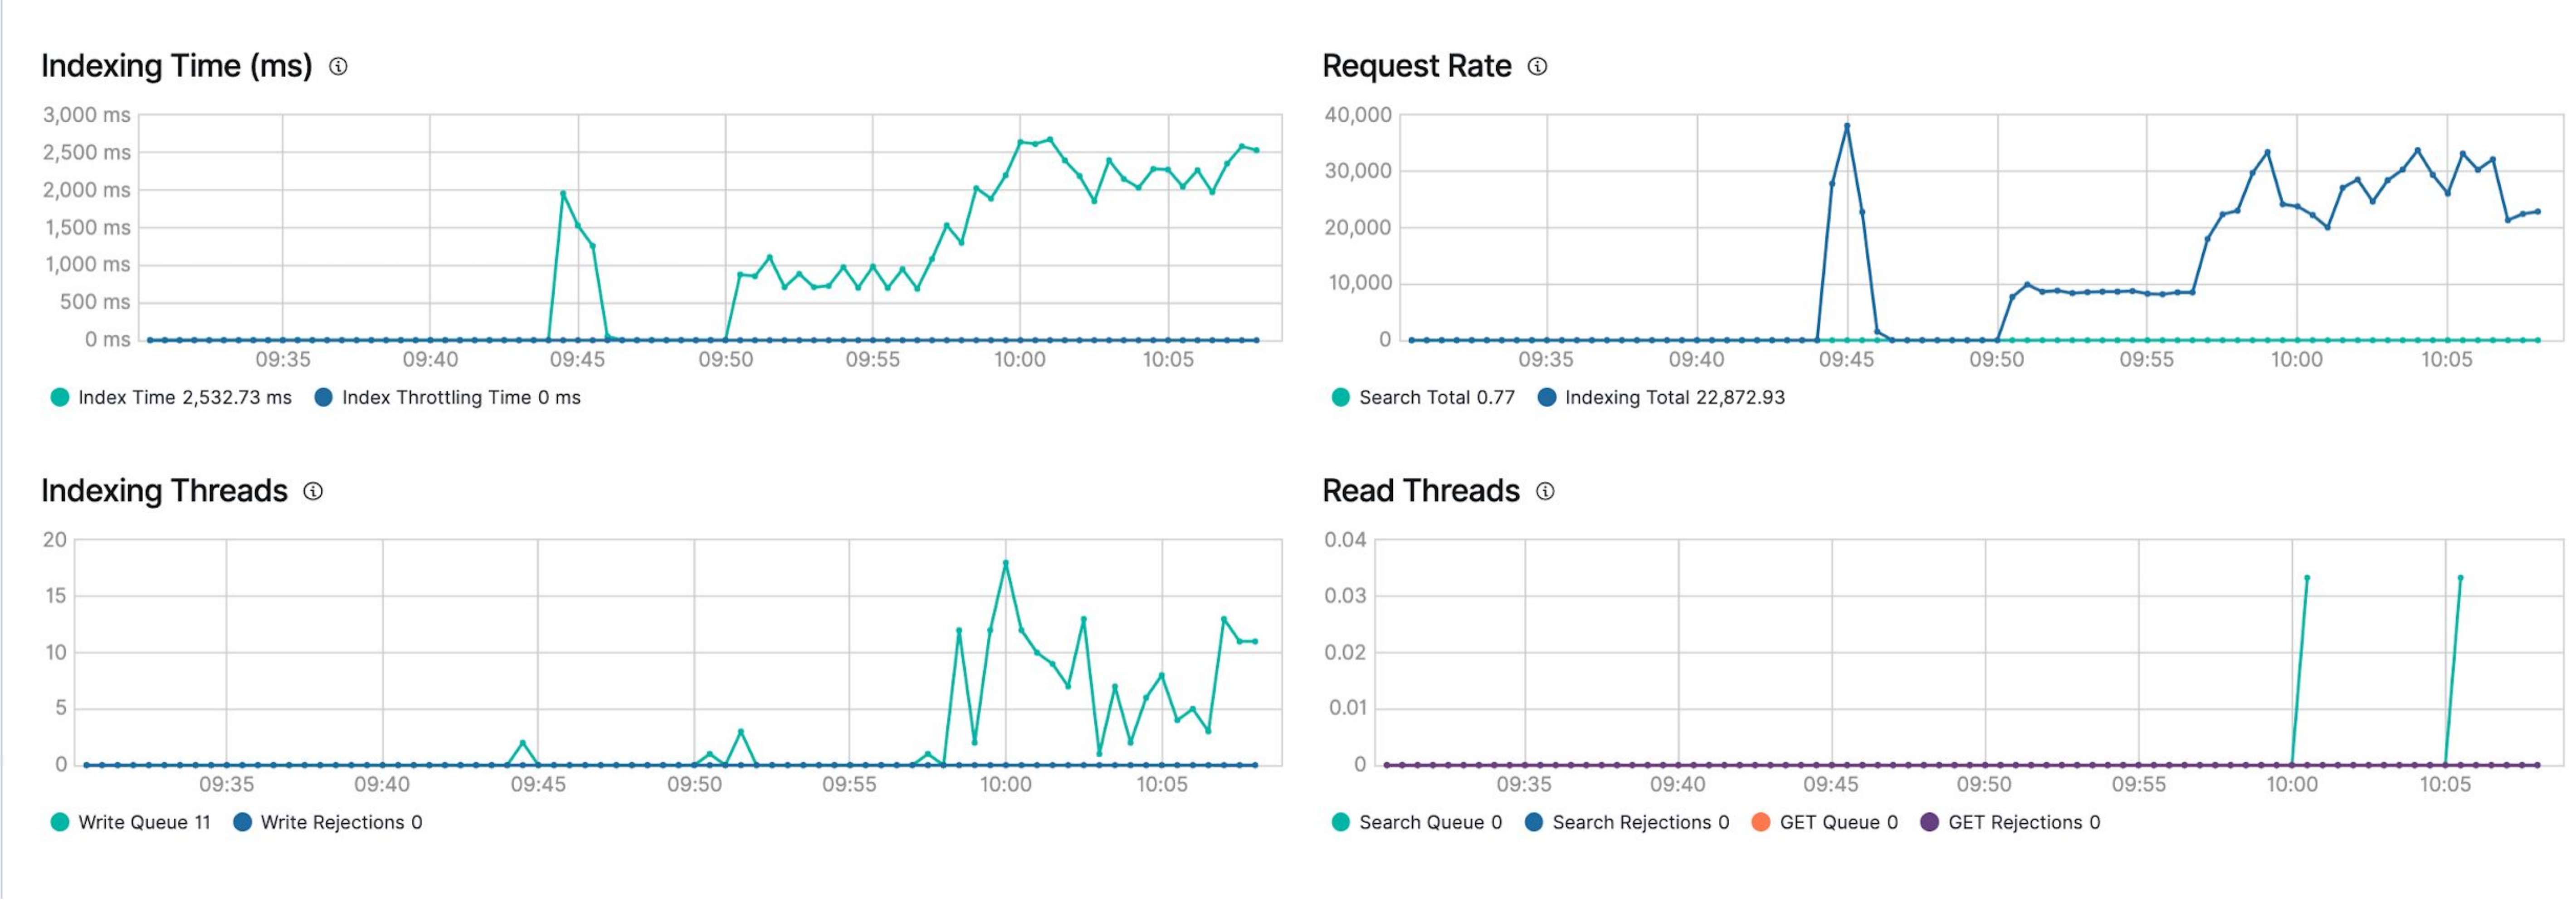

Monitoring

→ Example - High CPU usage

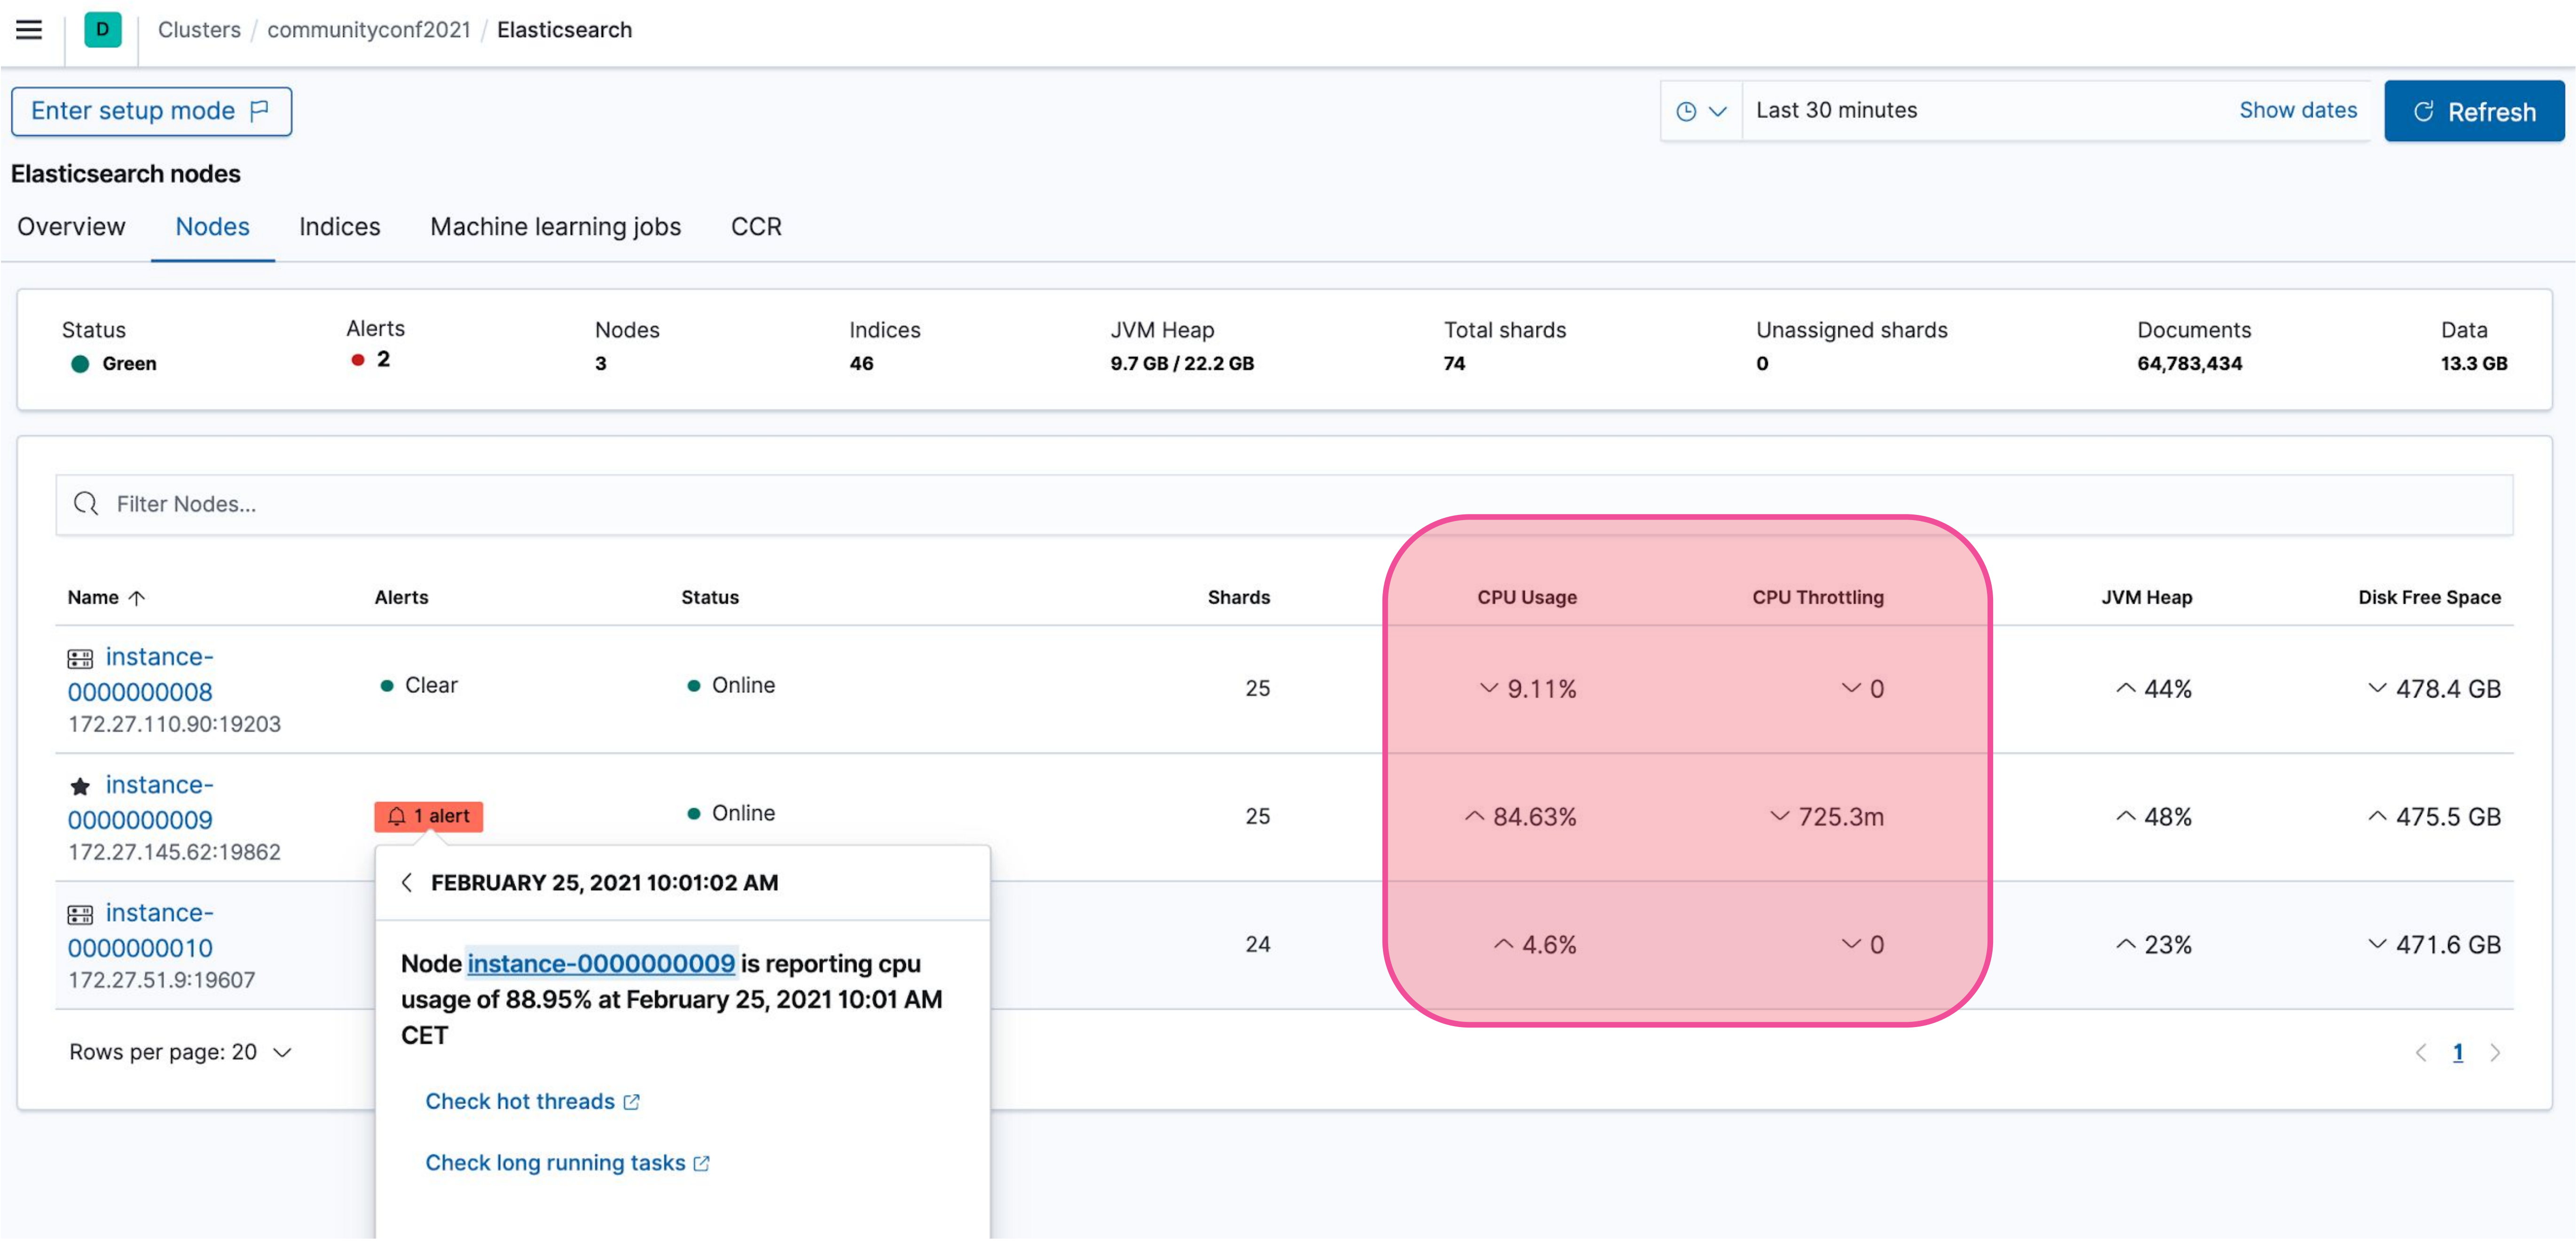

Elasticsearch nodes

Status	Alerts	Nodes	Indices	JVM Heap	Total shards	Unassigned shards	Documents	Data
● Green	● 2	3	46	9.7 GB / 22.2 GB	74	0	64,783,434	13.3 GB

Name ↑	Alerts	Status	Shards	CPU Usage	CPU Throttling	JVM Heap	Disk Free Space
instance-0000000008 172.27.110.90:19203	● Clear	● Online	25	∨ 9.11%	∨ 0	∧ 44%	∨ 478.4 GB
★ instance-0000000009 172.27.145.62:19862	1 alert	● Online	25	∧ 84.63%	∨ 725.3m	∧ 48%	∧ 475.5 GB
instance-0000000010 172.27.51.9:19607			24	∧ 4.6%	∨ 0	∧ 23%	∨ 471.6 GB

FEBRUARY 25, 2021 10:01:02 AM

Node instance-0000000009 is reporting cpu usage of 88.95% at February 25, 2021 10:01 AM CET

Check hot threads

Check long running tasks

elastic

Monitoring

→ Example - High CPU usage

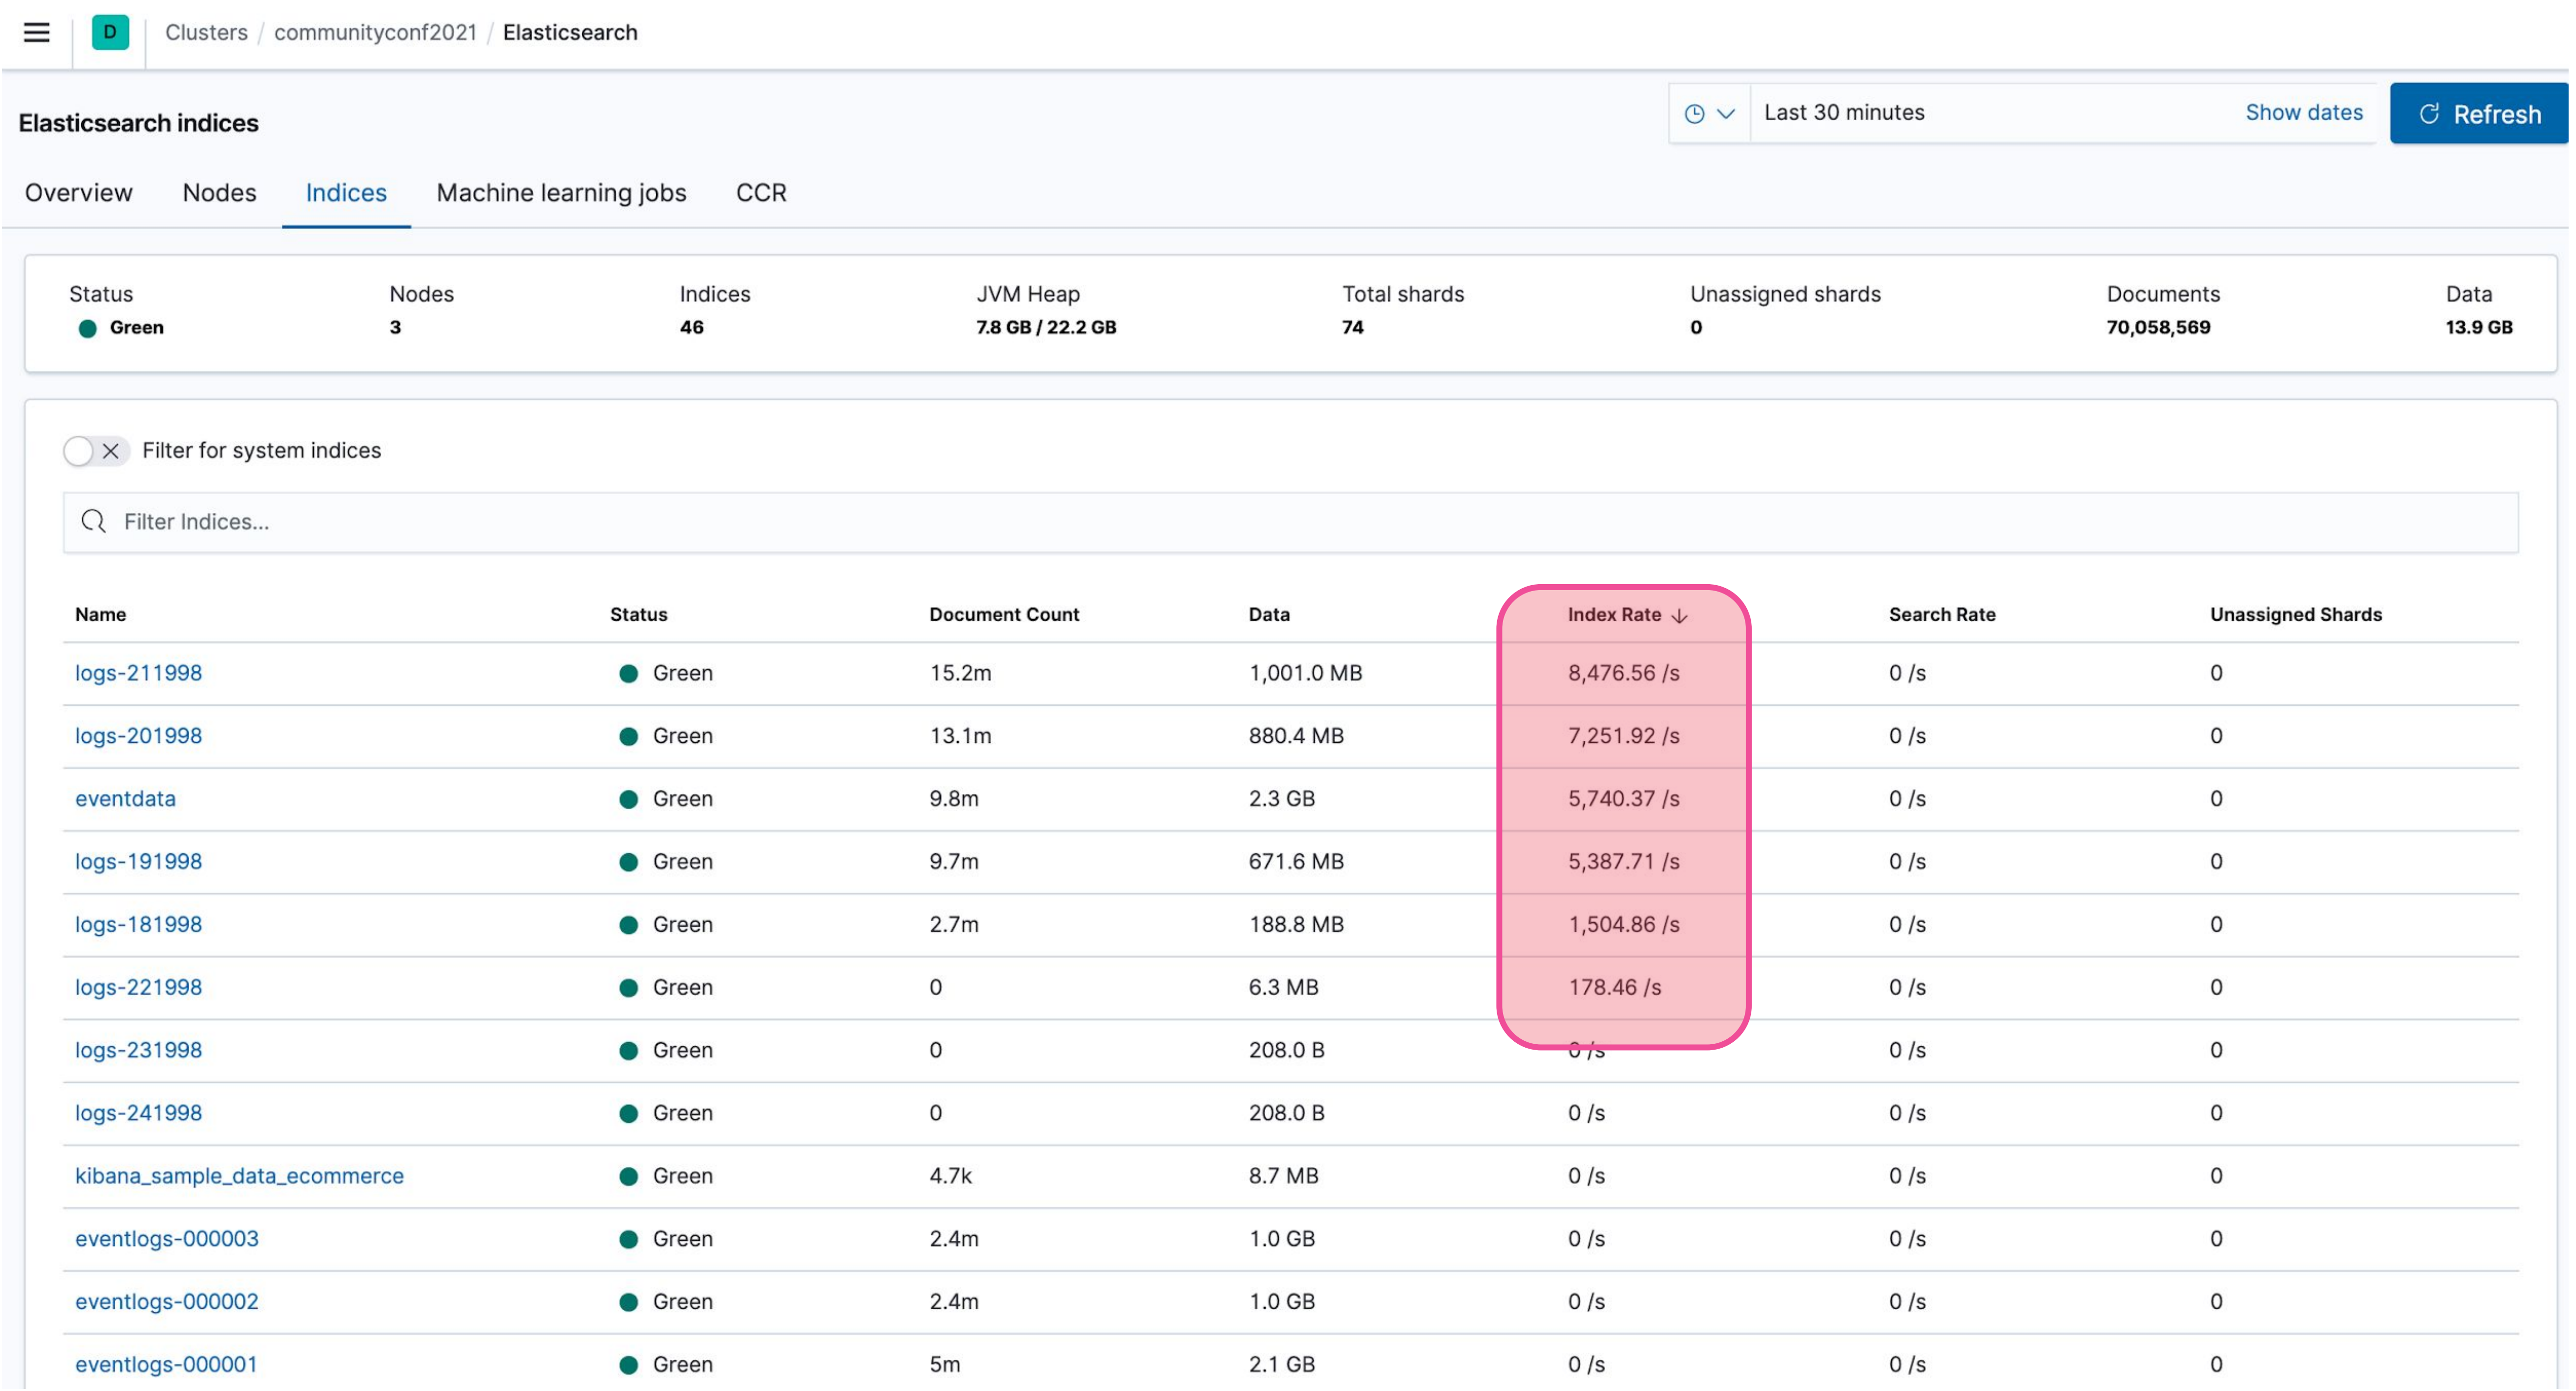

Elasticsearch indices

Status	Nodes	Indices	JVM Heap	Total shards	Unassigned shards	Documents	Data
● Green	3	46	7.8 GB / 22.2 GB	74	0	70,058,569	13.9 GB

Name	Status	Document Count	Data	Index Rate ↓	Search Rate	Unassigned Shards
logs-211998	● Green	15.2m	1,001.0 MB	8,476.56 /s	0 /s	0
logs-201998	● Green	13.1m	880.4 MB	7,251.92 /s	0 /s	0
eventdata	● Green	9.8m	2.3 GB	5,740.37 /s	0 /s	0
logs-191998	● Green	9.7m	671.6 MB	5,387.71 /s	0 /s	0
logs-181998	● Green	2.7m	188.8 MB	1,504.86 /s	0 /s	0
logs-221998	● Green	0	6.3 MB	178.46 /s	0 /s	0
logs-231998	● Green	0	208.0 B	0 /s	0 /s	0
logs-241998	● Green	0	208.0 B	0 /s	0 /s	0
kibana_sample_data_ecommerce	● Green	4.7k	8.7 MB	0 /s	0 /s	0
eventlogs-000003	● Green	2.4m	1.0 GB	0 /s	0 /s	0
eventlogs-000002	● Green	2.4m	1.0 GB	0 /s	0 /s	0
eventlogs-000001	● Green	5m	2.1 GB	0 /s	0 /s	0

And CAT APIs again!

➔ Example - High CPU usage

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cat-shards.html

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/cluster-nodes-hot-threads.html

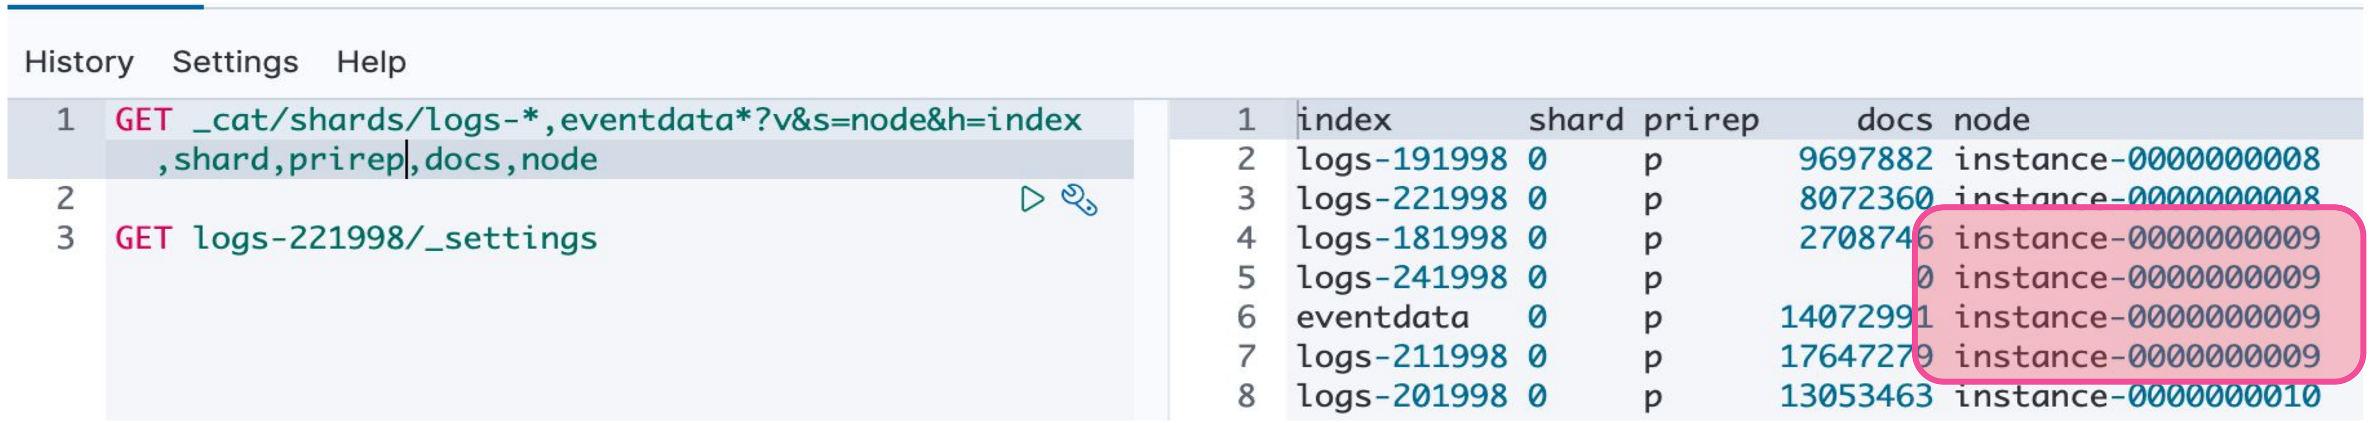

```
GET _cat/shards/logs-*,eventdata*?v&s=node&h=index
  ,shard,prirep,docs,node

GET logs-221998/_settings
```

```
index        shard prirep     docs node
logs-191998 0     p       9697882 instance-0000000008
logs-221998 0     p       8072360 instance-0000000008
logs-181998 0     p       2708746 instance-0000000009
logs-241998 0     p             0 instance-0000000009
eventdata   0     p      14072991 instance-0000000009
logs-211998 0     p      17647279 instance-0000000009
logs-201998 0     p      13053463 instance-0000000010
```

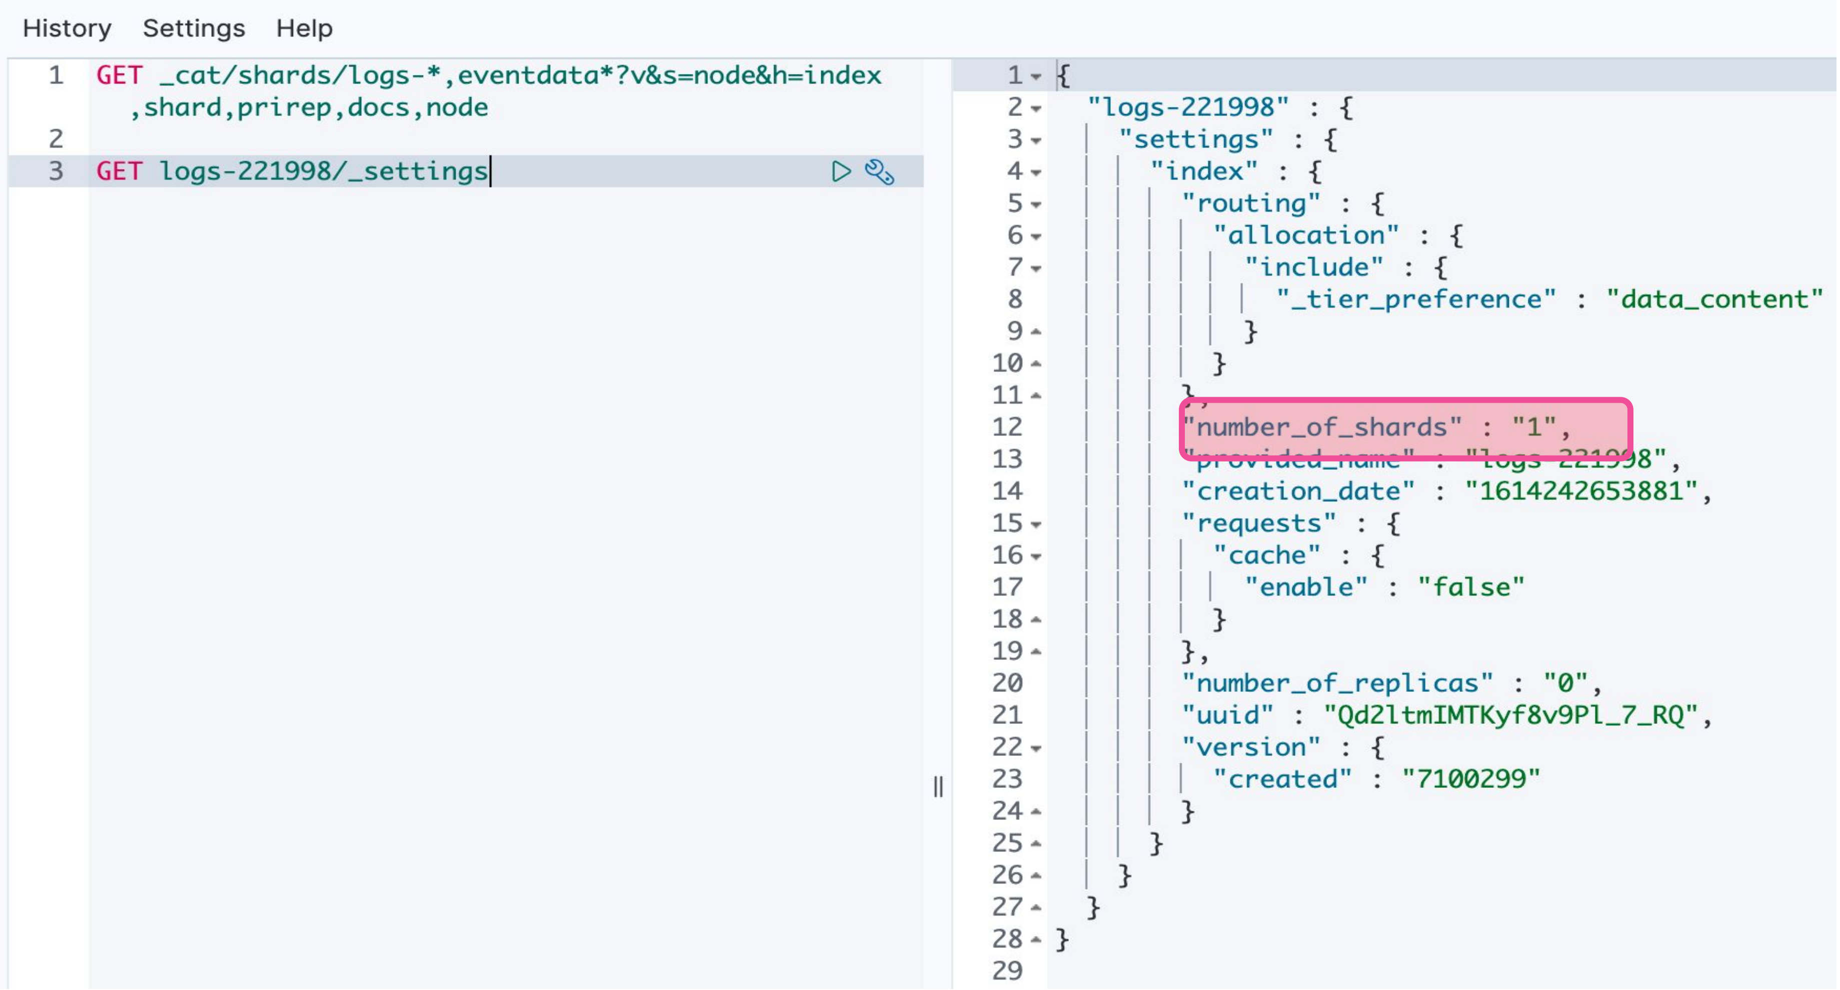

```
GET _cat/shards/logs-*,eventdata*?v&s=node&h=index
  ,shard,prirep,docs,node

GET logs-221998/_settings
```

```
{
  "logs-221998" : {
    "settings" : {
      "index" : {
        "routing" : {
          "allocation" : {
            "include" : {
              "_tier_preference" : "data_content"
            }
          }
        },
        "number_of_shards" : "1",
        "provided_name" : "logs-221998",
        "creation_date" : "1614242653881",
        "requests" : {
          "cache" : {
            "enable" : "false"
          }
        },
        "number_of_replicas" : "0",
        "uuid" : "Qd2ltmIMTKyf8v9Pl_7_RQ",
        "version" : {
          "created" : "7100299"
        }
      }
    }
  }
}
```

Log Analysis

→ Elasticsearch Logging

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/logging.html

https://www.elastic.co/guide/en/elasticsearch/reference/7.11/configuring-filebeat.html

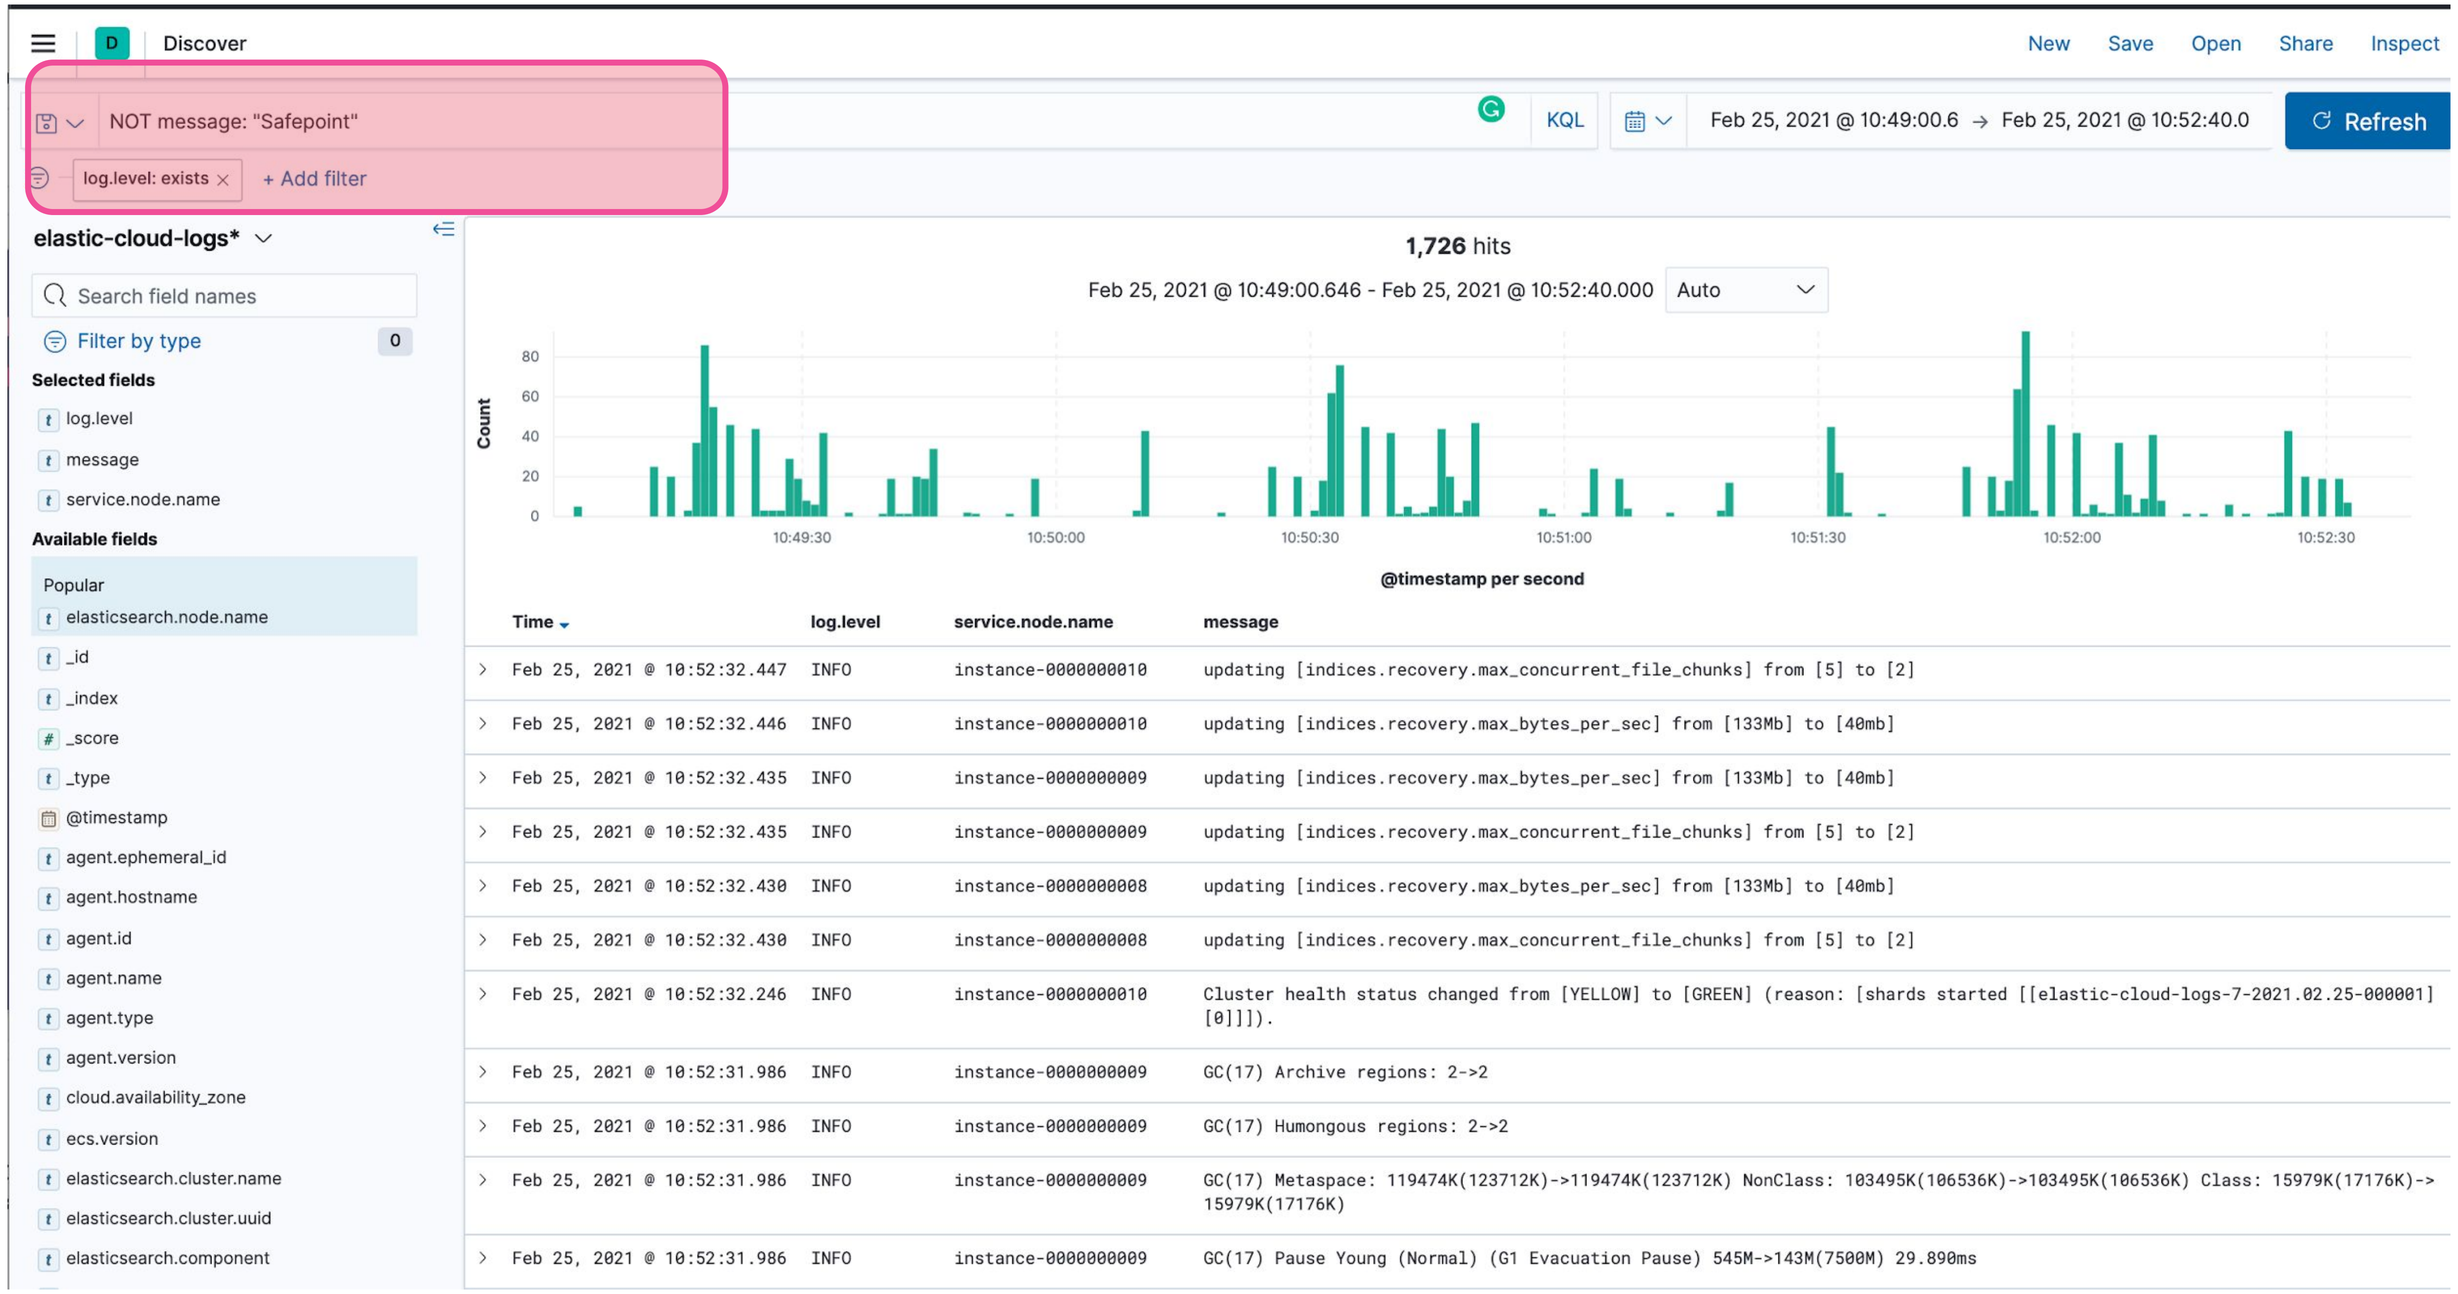

NOT message: "Safepoint"

log.level: exists

Feb 25, 2021 @ 10:49:00.6 → Feb 25, 2021 @ 10:52:40.0

1,726 hits

Feb 25, 2021 @ 10:49:00.646 - Feb 25, 2021 @ 10:52:40.000

Time	log.level	service.node.name	message
Feb 25, 2021 @ 10:52:32.447	INFO	instance-0000000010	updating [indices.recovery.max_concurrent_file_chunks] from [5] to [2]
Feb 25, 2021 @ 10:52:32.446	INFO	instance-0000000010	updating [indices.recovery.max_bytes_per_sec] from [133Mb] to [40mb]
Feb 25, 2021 @ 10:52:32.435	INFO	instance-0000000009	updating [indices.recovery.max_bytes_per_sec] from [133Mb] to [40mb]
Feb 25, 2021 @ 10:52:32.435	INFO	instance-0000000009	updating [indices.recovery.max_concurrent_file_chunks] from [5] to [2]
Feb 25, 2021 @ 10:52:32.430	INFO	instance-0000000008	updating [indices.recovery.max_bytes_per_sec] from [133Mb] to [40mb]
Feb 25, 2021 @ 10:52:32.430	INFO	instance-0000000008	updating [indices.recovery.max_concurrent_file_chunks] from [5] to [2]
Feb 25, 2021 @ 10:52:32.246	INFO	instance-0000000010	Cluster health status changed from [YELLOW] to [GREEN] (reason: [shards started [[elastic-cloud-logs-7-2021.02.25-000001][0]]]).
Feb 25, 2021 @ 10:52:31.986	INFO	instance-0000000009	GC(17) Archive regions: 2->2
Feb 25, 2021 @ 10:52:31.986	INFO	instance-0000000009	GC(17) Humongous regions: 2->2
Feb 25, 2021 @ 10:52:31.986	INFO	instance-0000000009	GC(17) Metaspace: 119474K(123712K)->119474K(123712K) NonClass: 103495K(106536K)->103495K(106536K) Class: 15979K(17176K)->15979K(17176K)
Feb 25, 2021 @ 10:52:31.986	INFO	instance-0000000009	GC(17) Pause Young (Normal) (G1 Evacuation Pause) 545M->143M(7500M) 29.890ms

Common Resources Shared

Sizing - how many shards per node and what size

- ➔ https://www.elastic.co/guide/en/elasticsearch/reference/7.11/size-your-shards.html
- ➔ https://www.elastic.co/blog/how-many-shards-should-i-have-in-my-elasticsearch-cluster
- ➔ https://www.elastic.co/guide/en/cloud/current/ec-reference-hardware.html
- ➔ https://benchmarks.elastic.co/
- ➔ https://esrally.readthedocs.io/

Storage

- ➔ https://www.elastic.co/blog/how-to-design-your-elasticsearch-data-storage-architecture-for-scale
- ➔ https://www.elastic.co/guide/en/elasticsearch/reference/7.11/tune-for-disk-usage.html

JVM Heap - do not go over ~30Gb heap

- ➔ https://www.elastic.co/blog/a-heap-of-trouble

Hot/Warm/Cold architectures for time series data

- ➔ https://www.elastic.co/blog/optimizing-costs-elastic-cloud-hot-warm-index-lifecycle-management

elastic

Common Resources Shared

Tuning for search - slow searches

➔ https://www.elastic.co/blog/advanced-tuning-finding-and-fixing-slow-elasticsearch-queries
➔ https://www.elastic.co/guide/en/elasticsearch/reference/7.11/tune-for-search-speed.html

Tuning for ingest - use bulk!

➔ https://www.elastic.co/guide/en/elasticsearch/reference/7.11/tune-for-indexing-speed.html
➔ https://www.elastic.co/guide/en/elasticsearch/reference/7.11/docs-bulk.html

Upgrading the Stack - be prepared, test and snapshots!

➔ https://www.elastic.co/webinars/expert-tips-for-upgrading-the-elk-stack
➔ https://www.elastic.co/guide/en/elastic-stack/7.11/upgrading-elastic-stack.html

Secure the Stack

➔ https://www.elastic.co/blog/configuring-ssl-tls-and-https-to-secure-elasticsearch-kibana-beats-and-logstash

Optimize Mappings

➔ https://www.elastic.co/blog/strings-are-dead-long-live-strings

elastic

Wrapping Up

Triage incidents

- ➔ How critical is it?
- ➔ Do we need urgent care or is there a workaround to stabilize?

Have tools ready

- ➔ REST APIs / Support diagnostics
- ➔ Monitoring & Alerts
- ➔ Log Analysis / Kibana Discover
- ➔ Search Elastic discuss, Stackoverflow, Elastic GitHub repos, etc..

Lessons learned

- ➔ Follow best practices
- ➔ Prevent future incidents - proactively investigate unexpected logs, etc.

elastic

Q & A

Thank You

elastic