How a single API endpoint saved 3000 CPU

MAERSK

Maersk Observability Platform

MAERSK

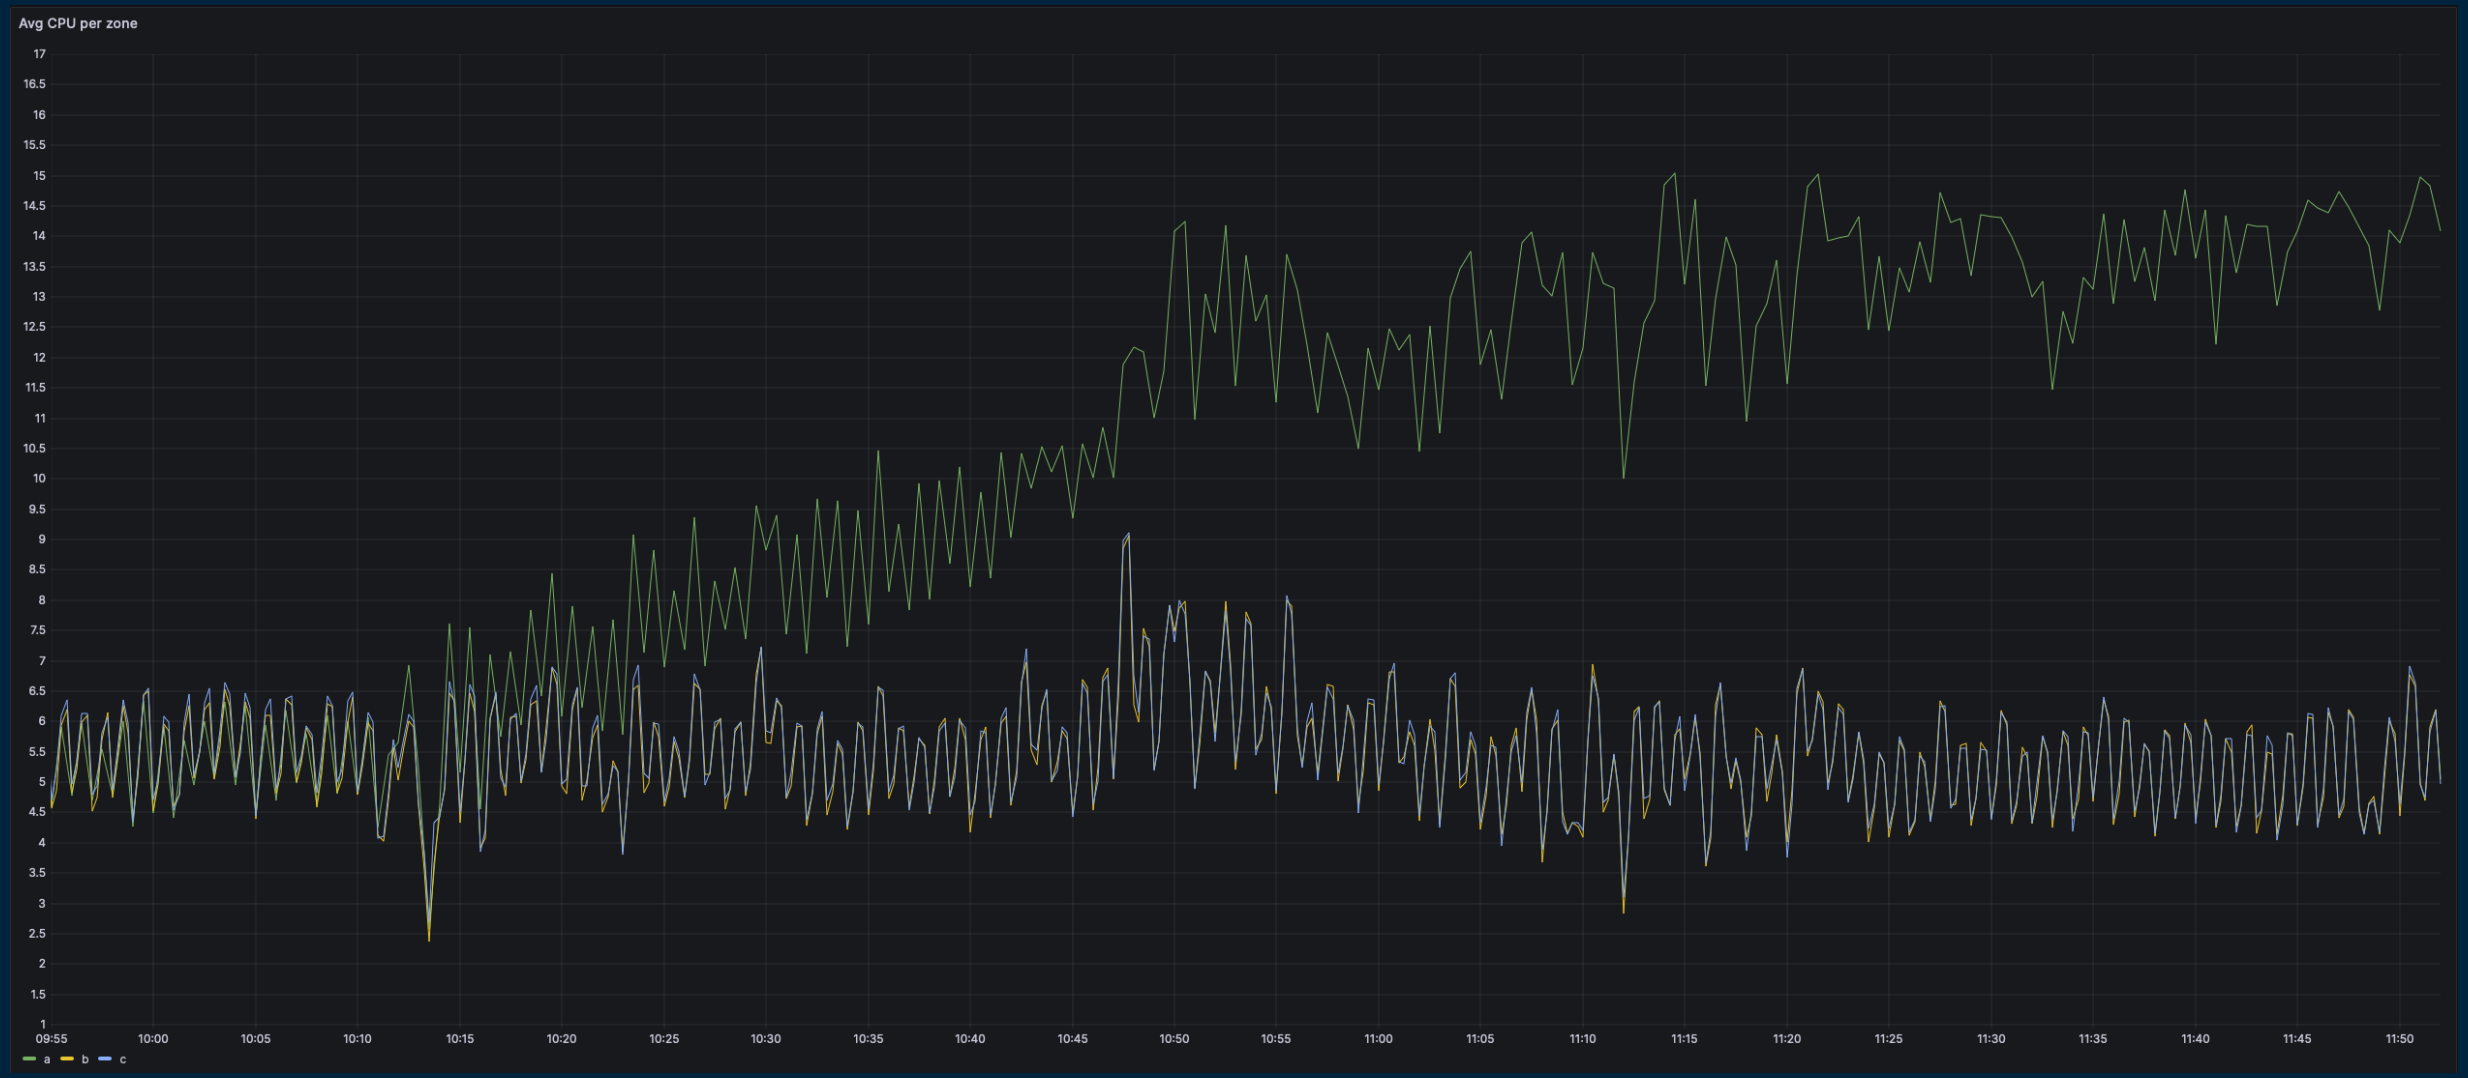
Avg CPU per zone
17
16.5
16
15.5
15
14.5
14
13.5
13
12.5
12
11.5
11
10.5
10
9.5
9
8.5
8
7.5
7
6.5
6
5.5
5
4.5
4
3.5
3
2.5
2
1.5
1
09:55
10:00
10:05
10:10
10:15
10:20
10:25
10:30
10:35
10:40
10:45
10:50
10:55
11:00
11:05
11:10
11:15
11:20
11:25
11:30
11:35
11:40
11:45
11:50
a
b
c

MAERSK

What is Mimir? What is an ingester?

Open-source time-series database for metrics

The write path

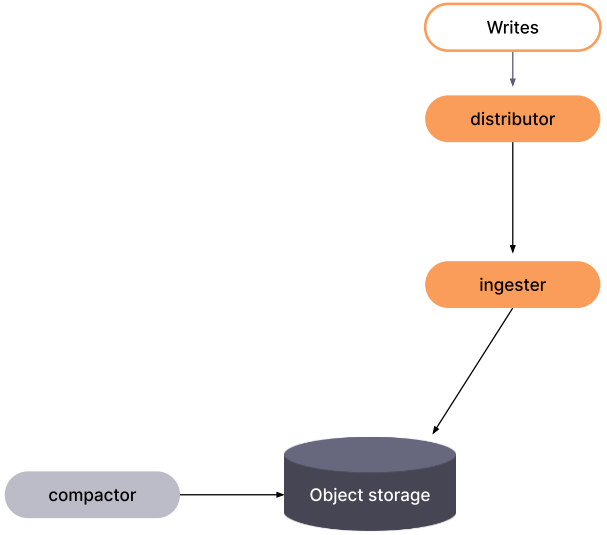

The read path

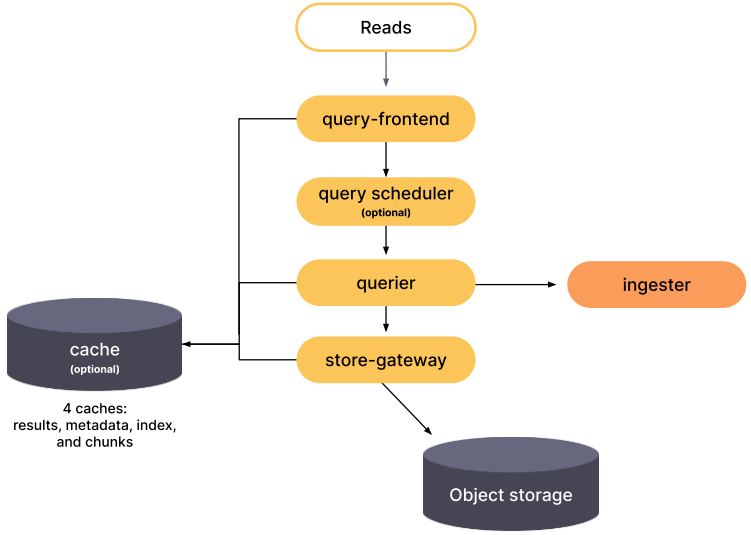

MAERSK

Why are CPU spikes an issue?

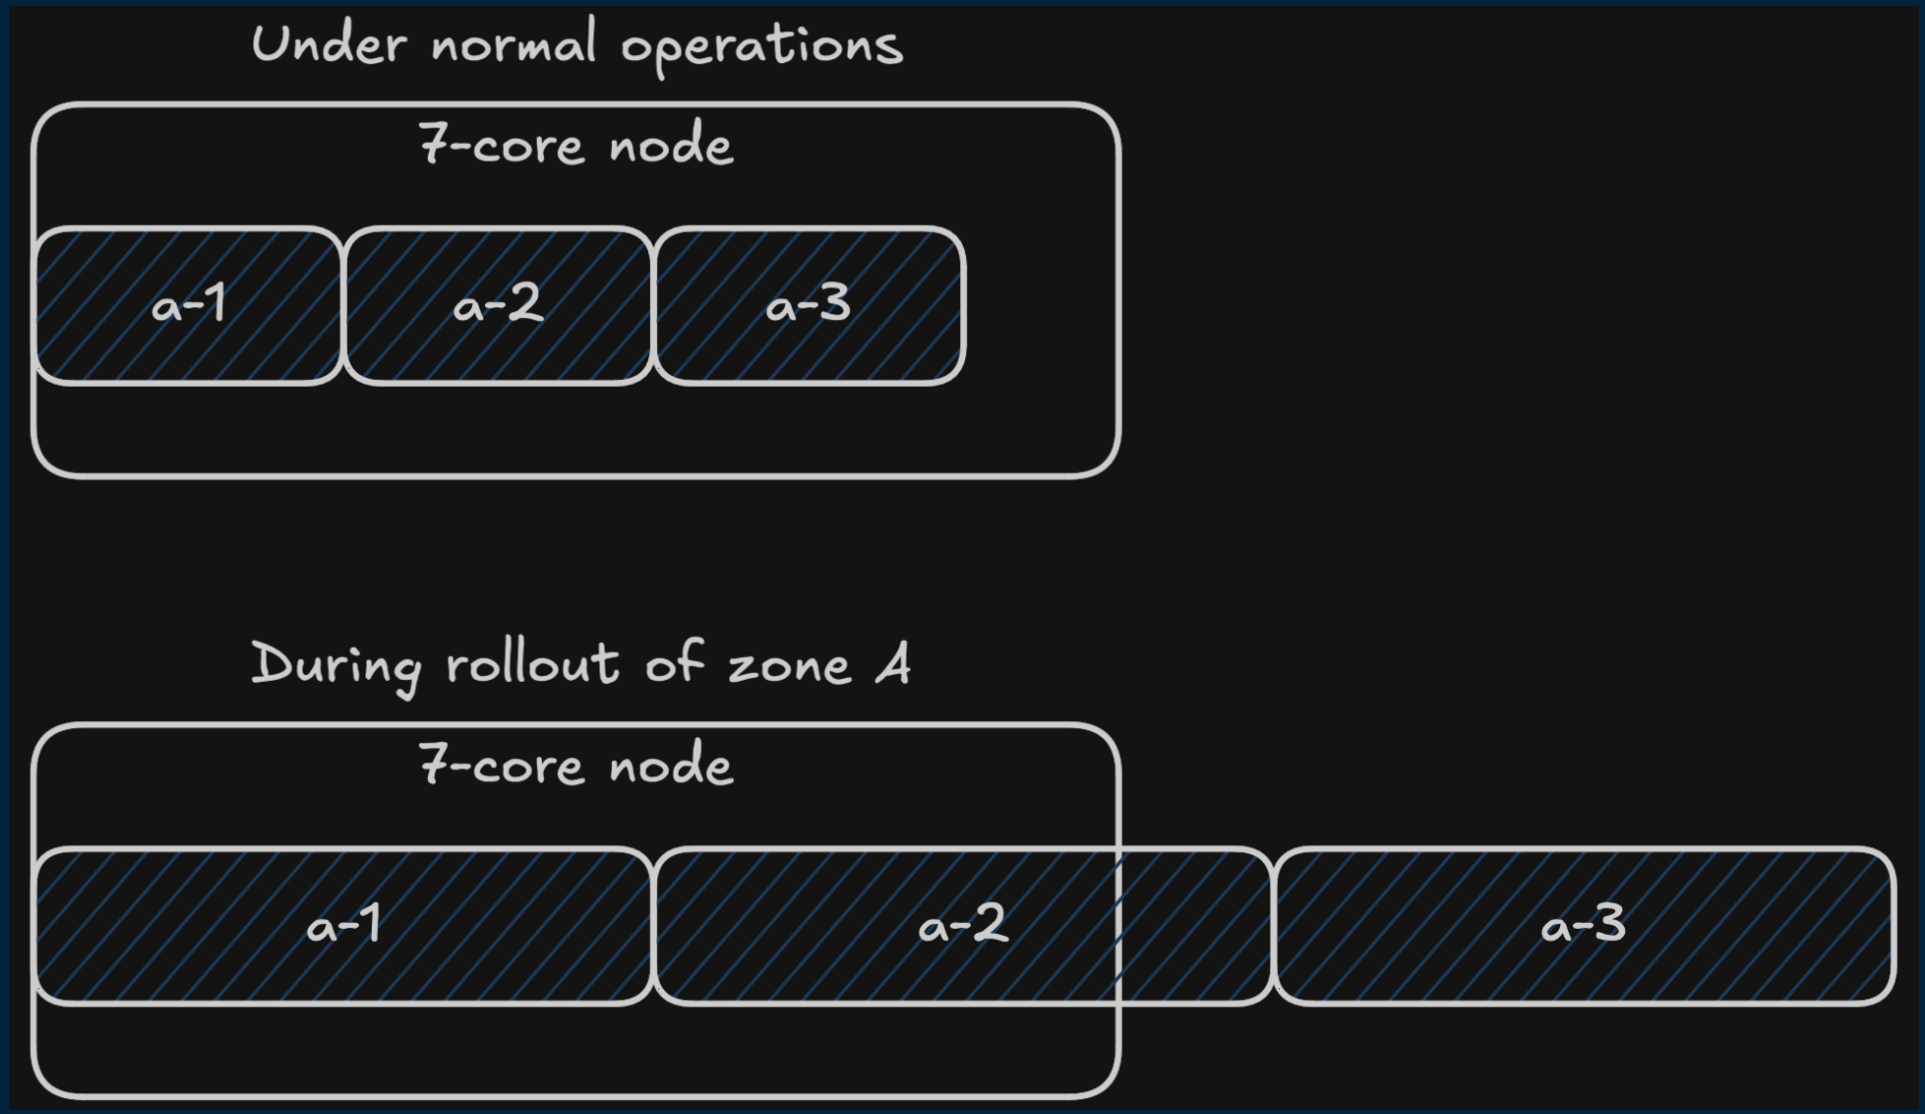

MAERSK

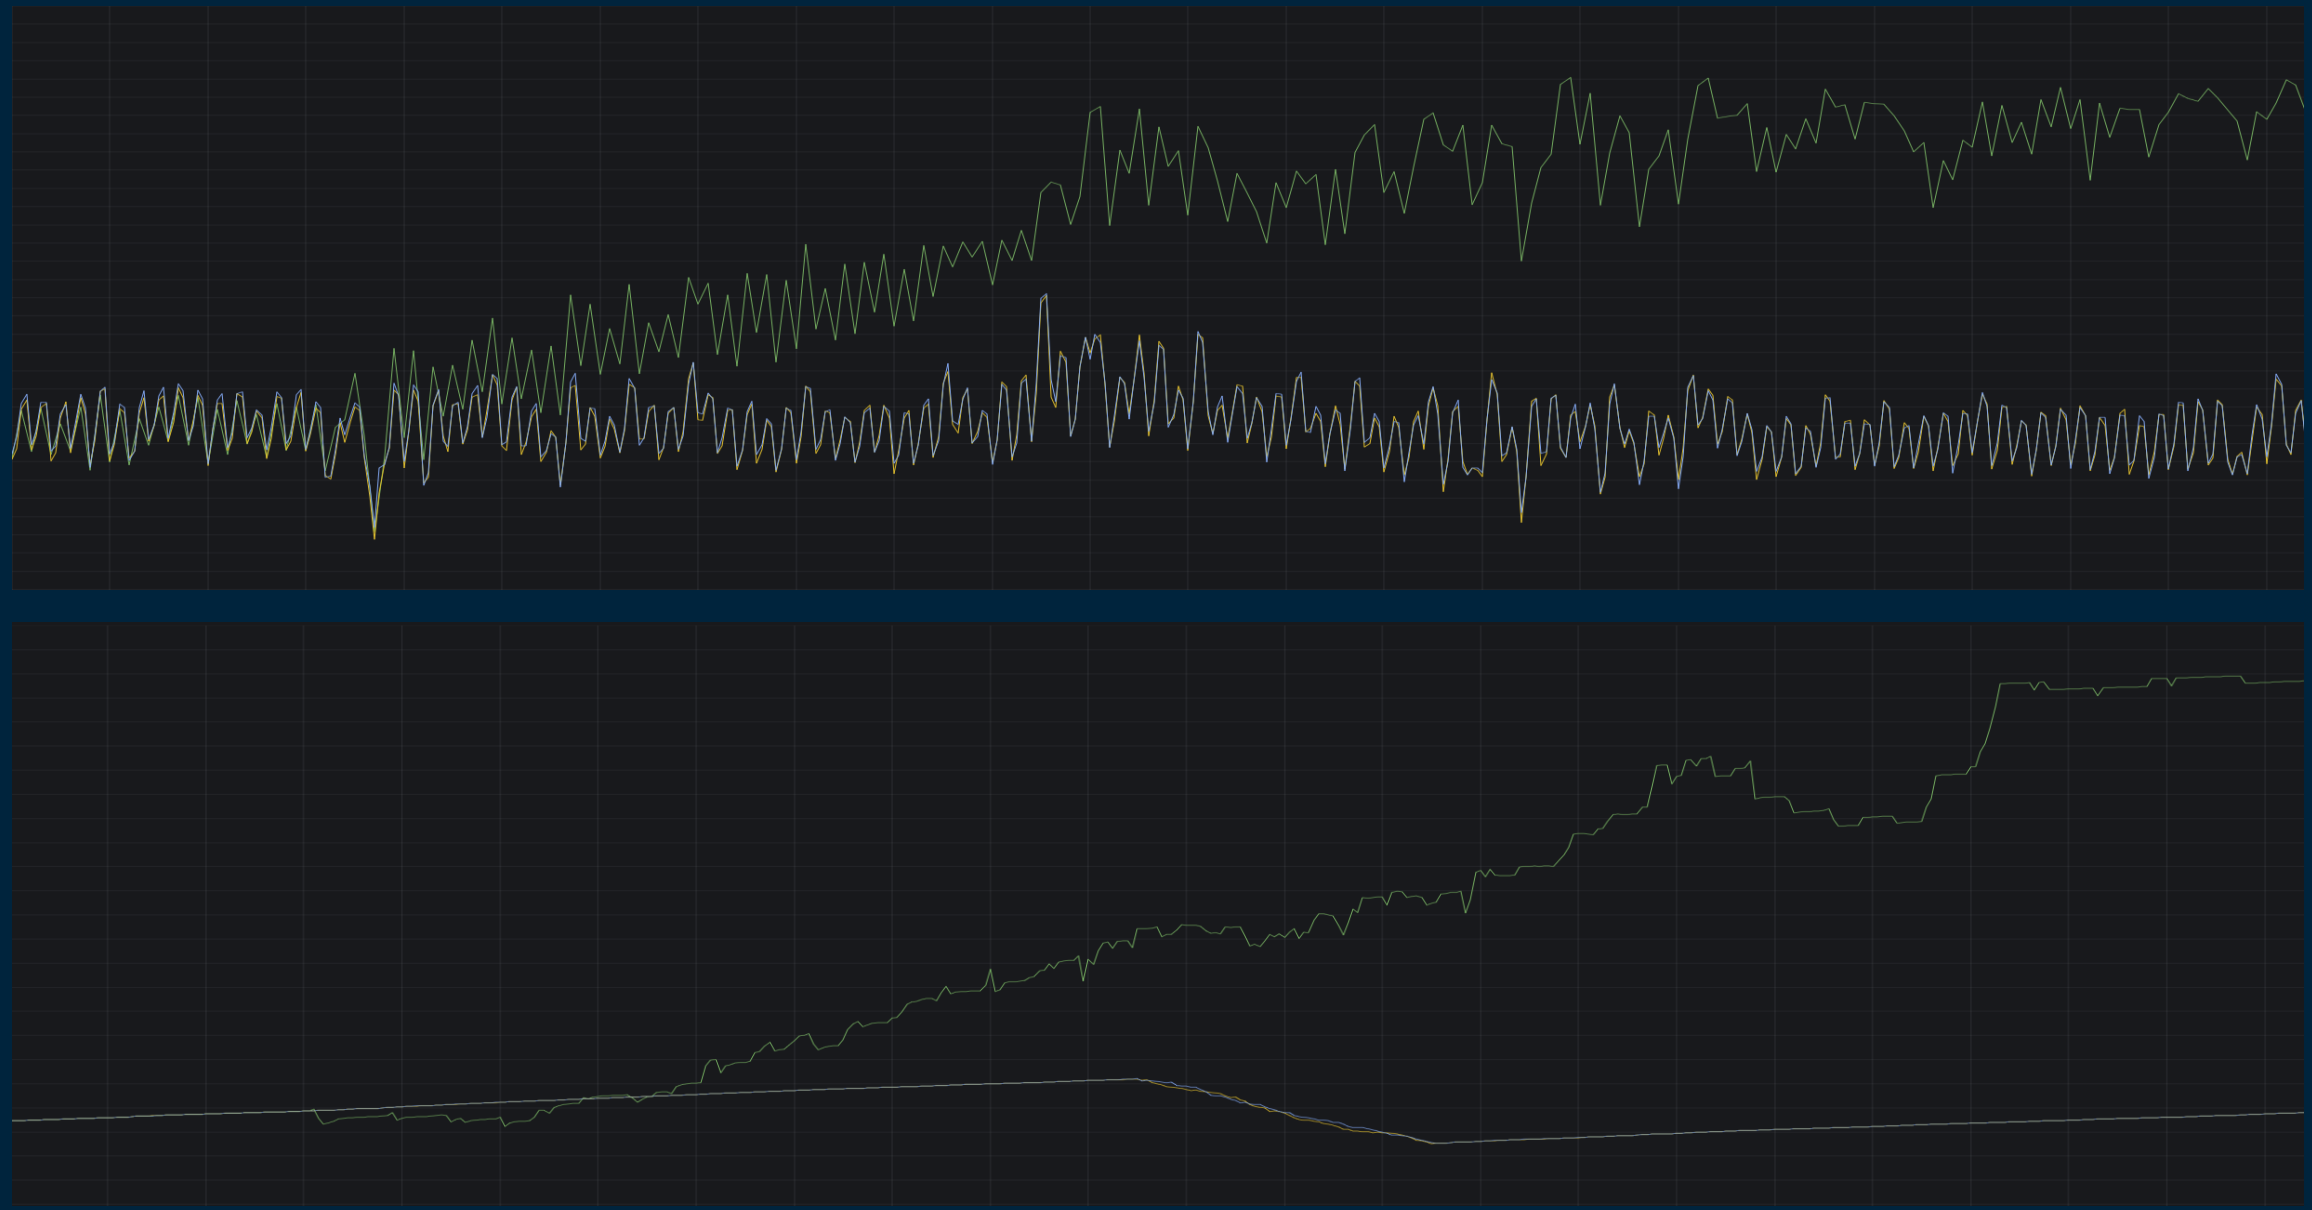

MAERSK

What is a metric series?

response_time{}

response_time{status_code=”200”}
response_time{status_code=”400”}
response_time{status_code=”500”}

response_time{status_code=”200”, path=”/read”}
response_time{status_code=”400”, path=”/read”}
response_time{status_code=”500”, path=”/read”}
response_time{status_code=”200”, path=”/write”}
response_time{status_code=”400”, path=”/ write”}
response_time{status_code=”500”, path=”/ write”}

response_time{status_code=”200”, path=”/read”, continent=”eu”}
response_time{status_code=”400”, path=”/read”, continent=”eu”}
response_time{status_code=”500”, path=”/read”, continent=”eu”}
response_time{status_code=”200”, path=”/write”, continent=”eu”}
response_time{status_code=”400”, path=”/ write”, continent=”eu”}
response_time{status_code=”500”, path=”/ write”, continent=”eu”}
response_time{status_code=”200”, path=”/read”, continent=”na”}
response_time{status_code=”400”, path=”/read”, continent=”na”}
response_time{status_code=”500”, path=”/read”, continent=”na”}
response_time{status_code=”200”, path=”/write”, continent=”na”}
response_time{status_code=”400”, path=”/ write”, continent=”na”}
response_time{status_code=”500”, path=”/ write”, continent=”na”}
response_time{status_code=”200”, path=”/read”, continent=”oc”}
response_time{status_code=”400”, path=”/read”, continent=”oc”}
response_time{status_code=”500”, path=”/read”, continent=”oc”}
response_time{status_code=”200”, path=”/write”, continent=”oc”}
response_time{status_code=”400”, path=”/ write”, continent=”oc”}
response_time{status_code=”500”, path=”/ write”, continent=”oc”}

Each unique set of label values is one metric series

MAERSK

What makes up an ingester's CPU usage?

(metric_series_count * reads_per_second) + writes_per_second + misc = CPU usage

	Few reads per second	**Many reads per second**
Low metric series count	Low CPU usage	Medium CPU usage
High metric series count	Medium CPU usage	High CPU usage

This is why we care about metric cardinality!

MAERSK

Why do metric series proliferate during rollouts?

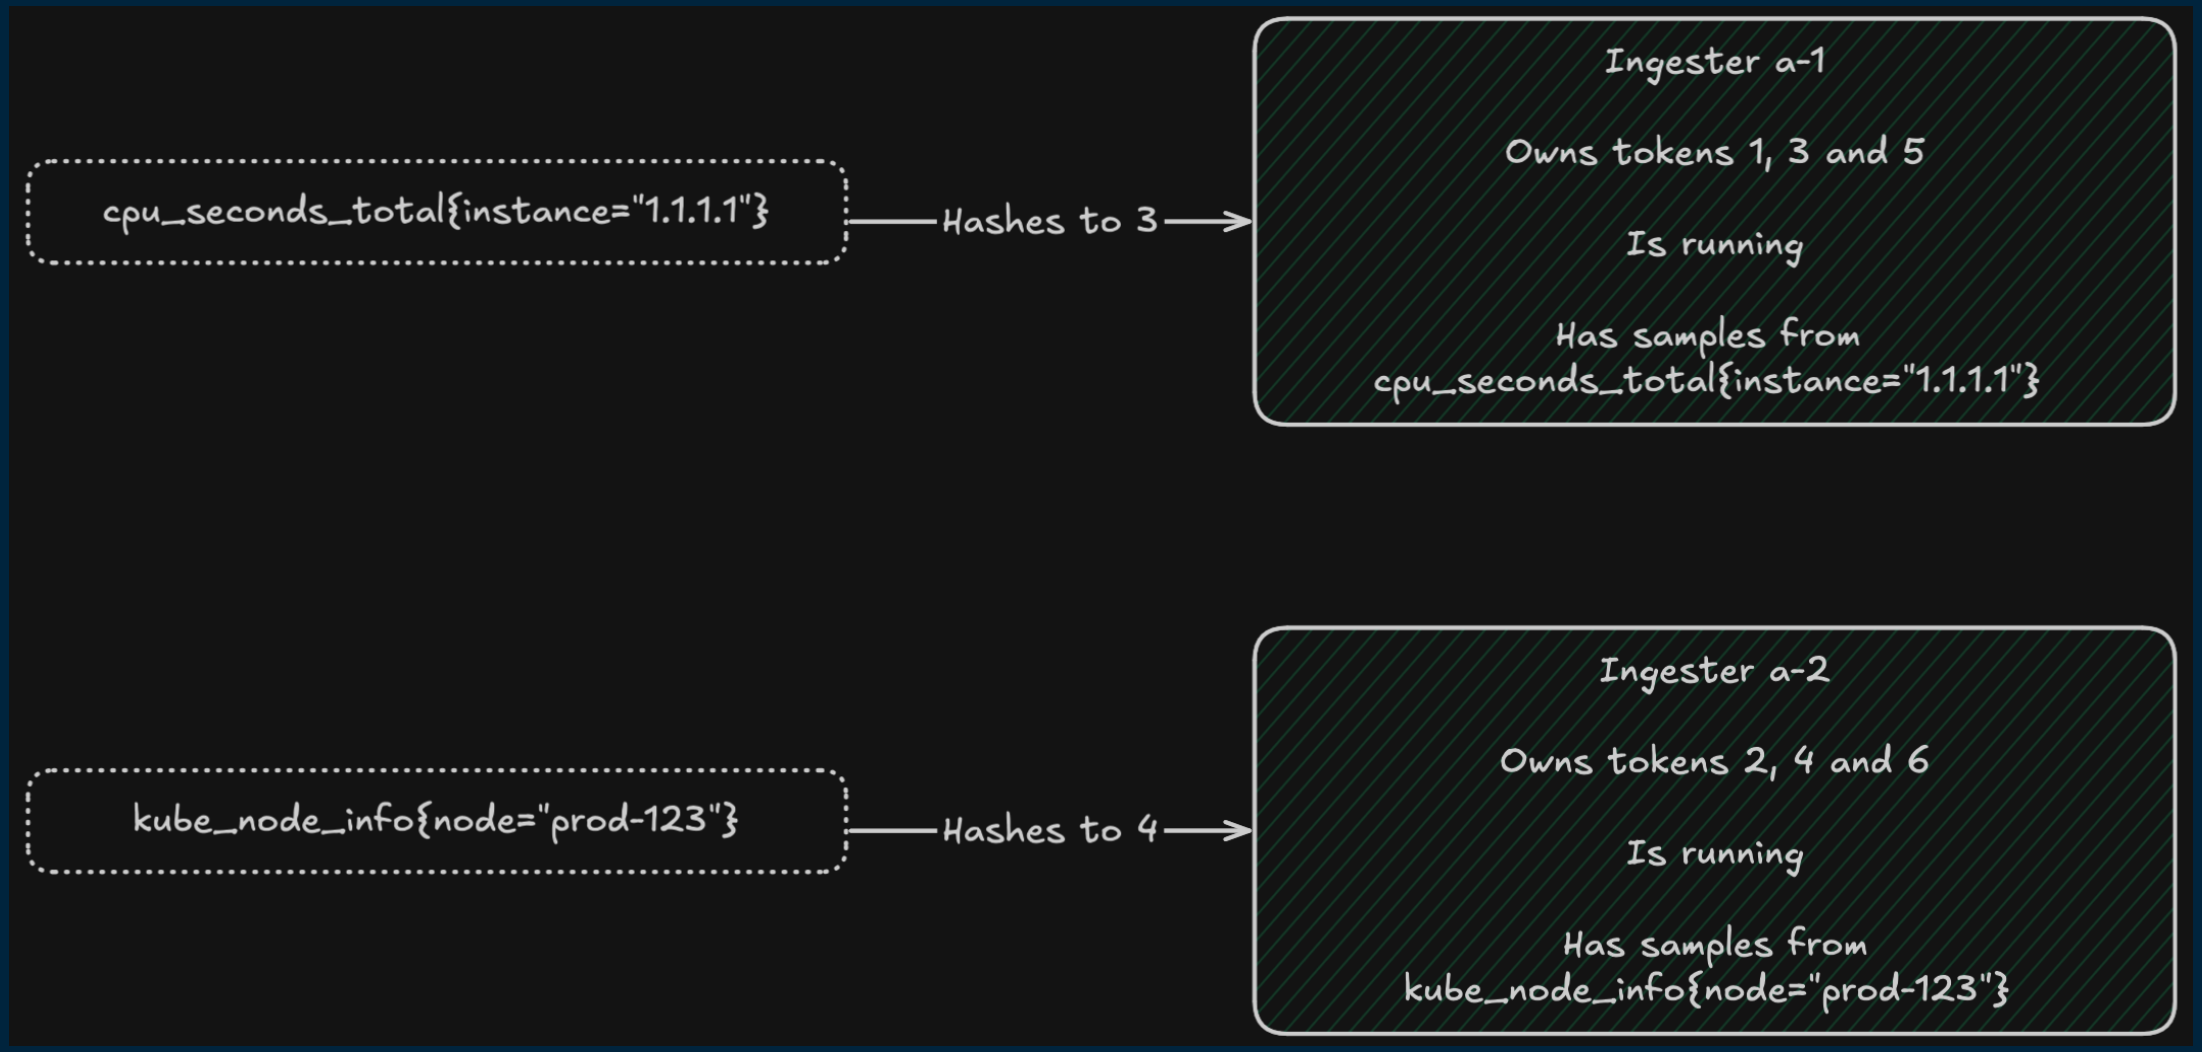

MAERSK

Why do metric series proliferate during rollouts?

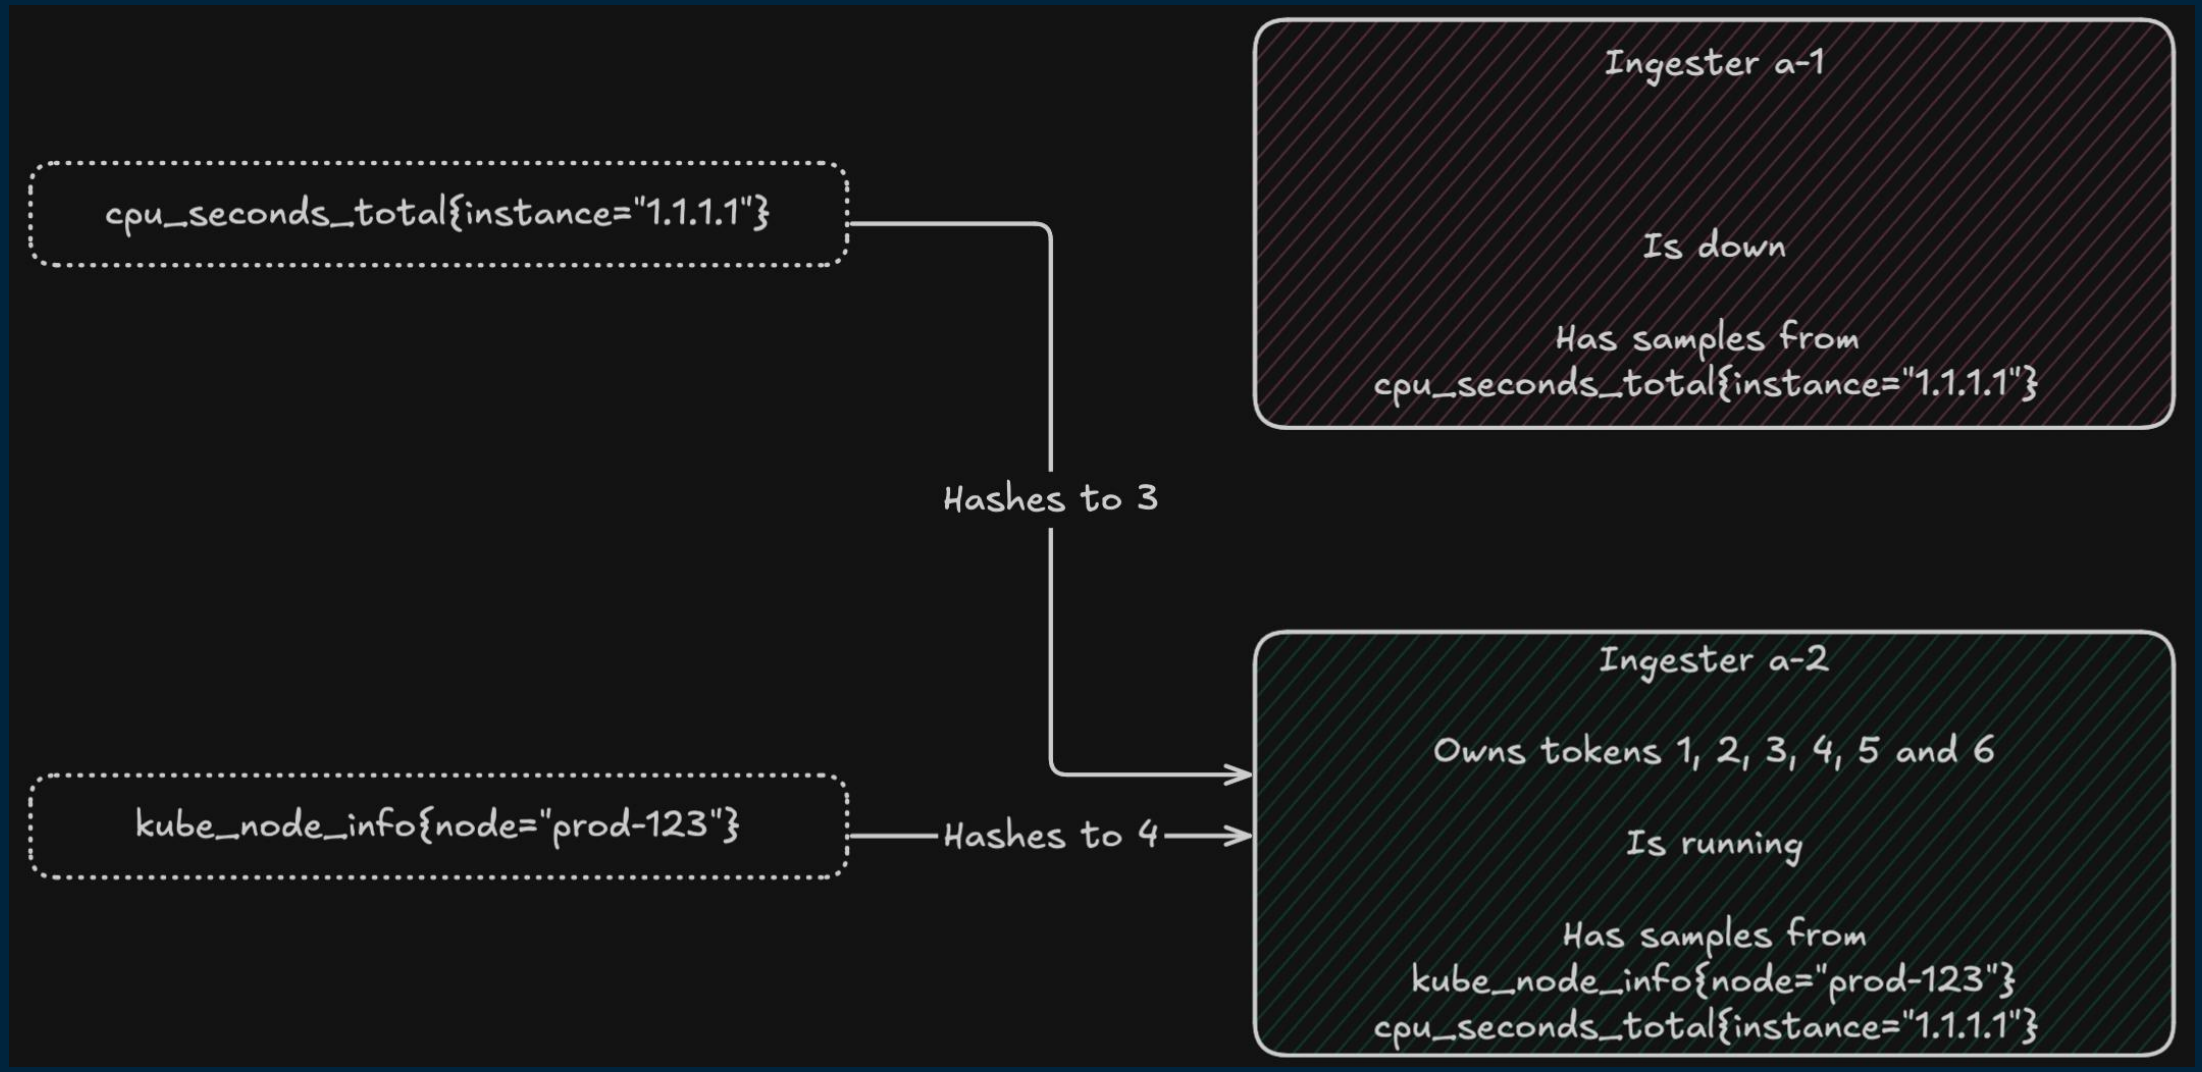

MAERSK

Why do metric series proliferate during rollouts?

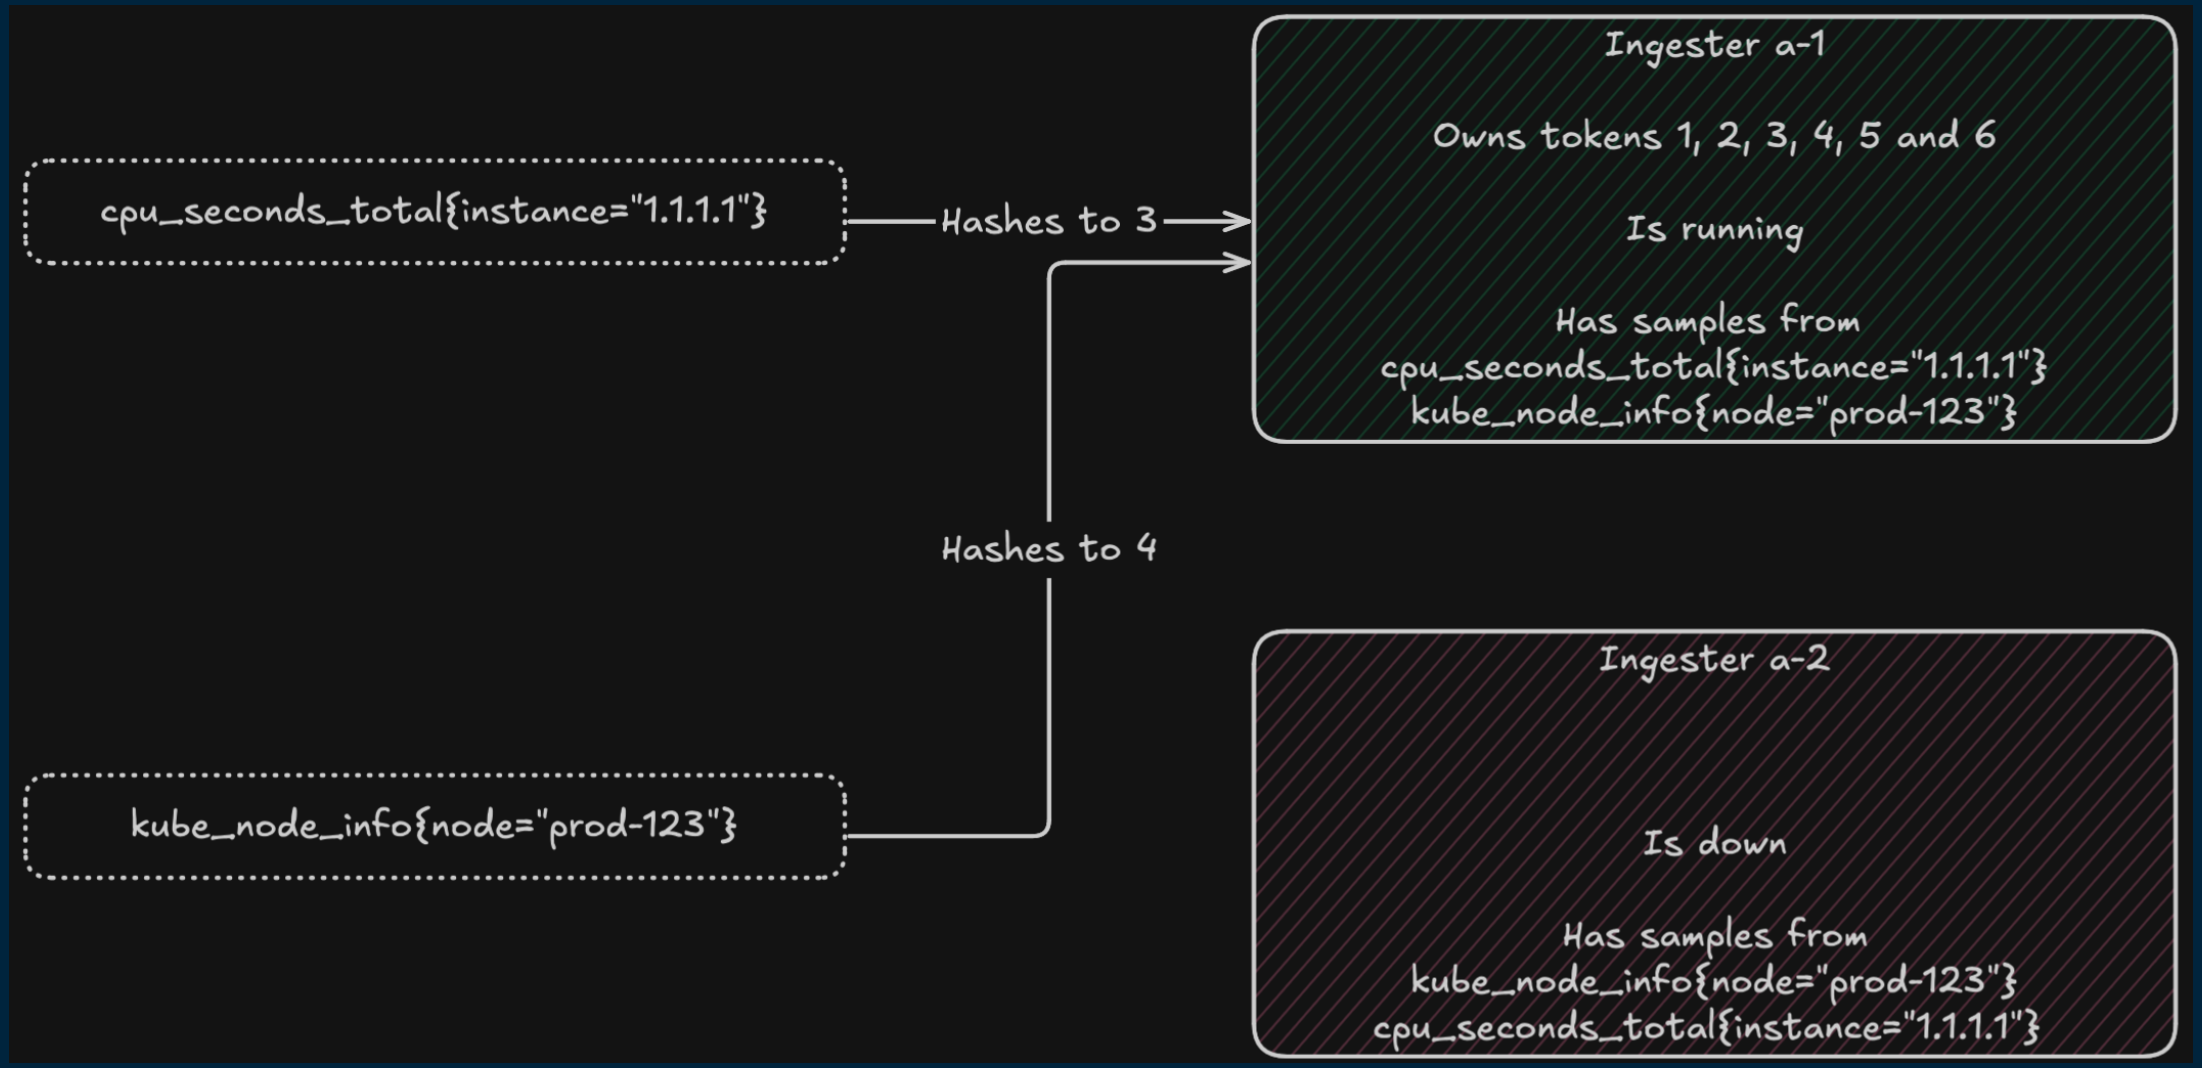

MAERSK

Why do metric series proliferate during rollouts?

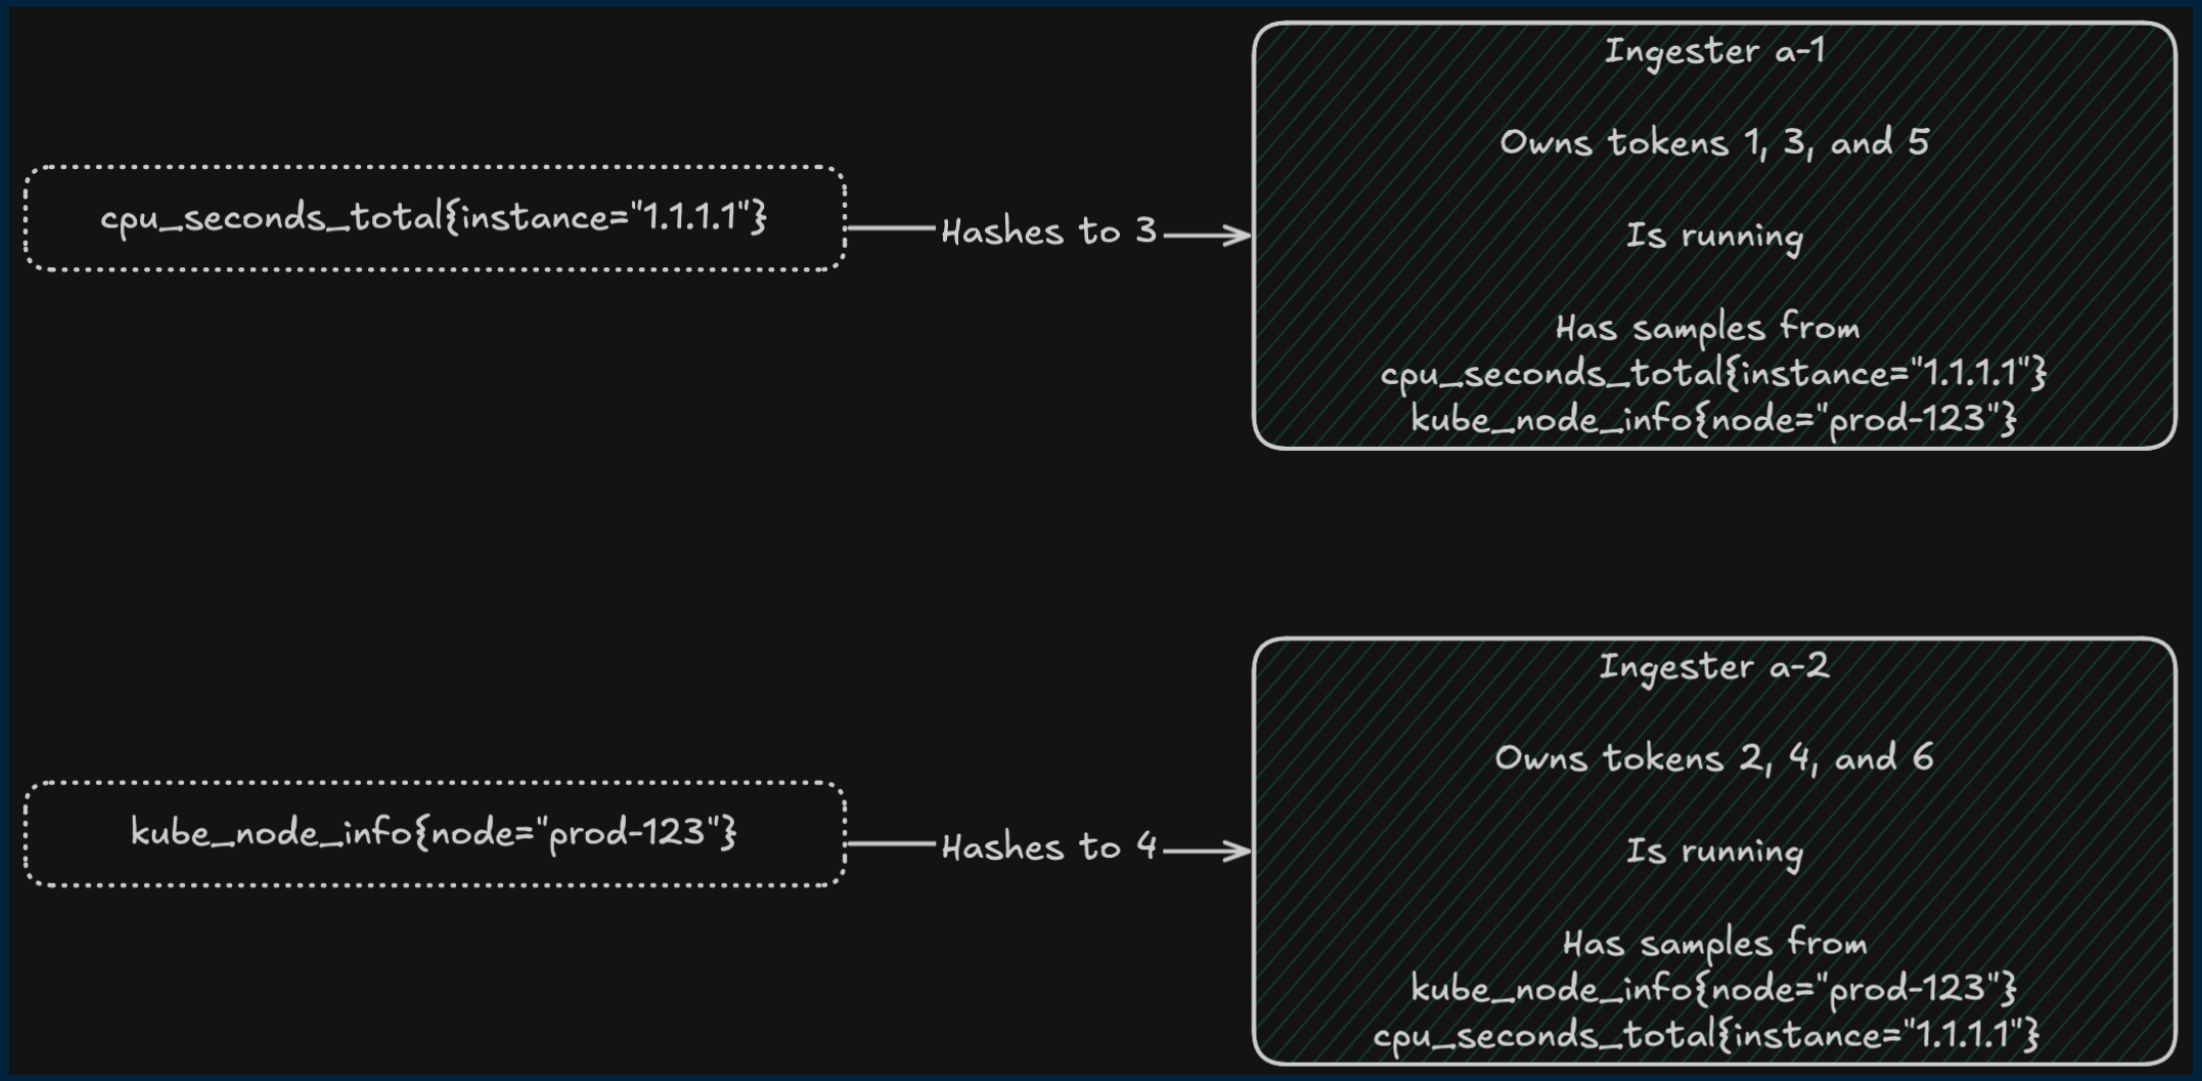

MAERSK

Let’s recap

- Read latency increases during rollouts *because*
- Nodes don’t have enough CPU time to serve demand *because*
- Ingester CPU usage surges *because*
- The amount of metric series held by each ingester increases by ~100% *because*
- Series ownership is handed back and forth between ingesters *because*
- Ingesters are continuously terminated and leave the ring

MAERSK

Potential solutions

Don’t do rollouts

MAERSK

Potential solutions

Don’t leave the ring on shutdown

Pros:

- Metric series no longer proliferate during rollouts.
- Easy to implement.

Cons:

- Multi-zone evictions cause downtime.

Zone A:

- Ingester a-1
- Ingester a-2
- Ingester a-3
- Ingester a-4

Zone B:

- Ingester b-1
- Ingester b-2
- Ingester b-3
- Ingester b-4

Zone C:

- Ingester c-1
- Ingester c-2
- Ingester c-3
- Ingester c-4

It takes only a single terminated ingester for the entire zone to be considered unhealthy!

MAERSK

Potential solutions

Conditionally leave the ring on shutdown

Pros:

- Metric series no longer proliferate during rollouts.
- Multi-zone evictions don't cause downtime

Cons:

- Requires buy-in from Mimir team and upstream contribution.

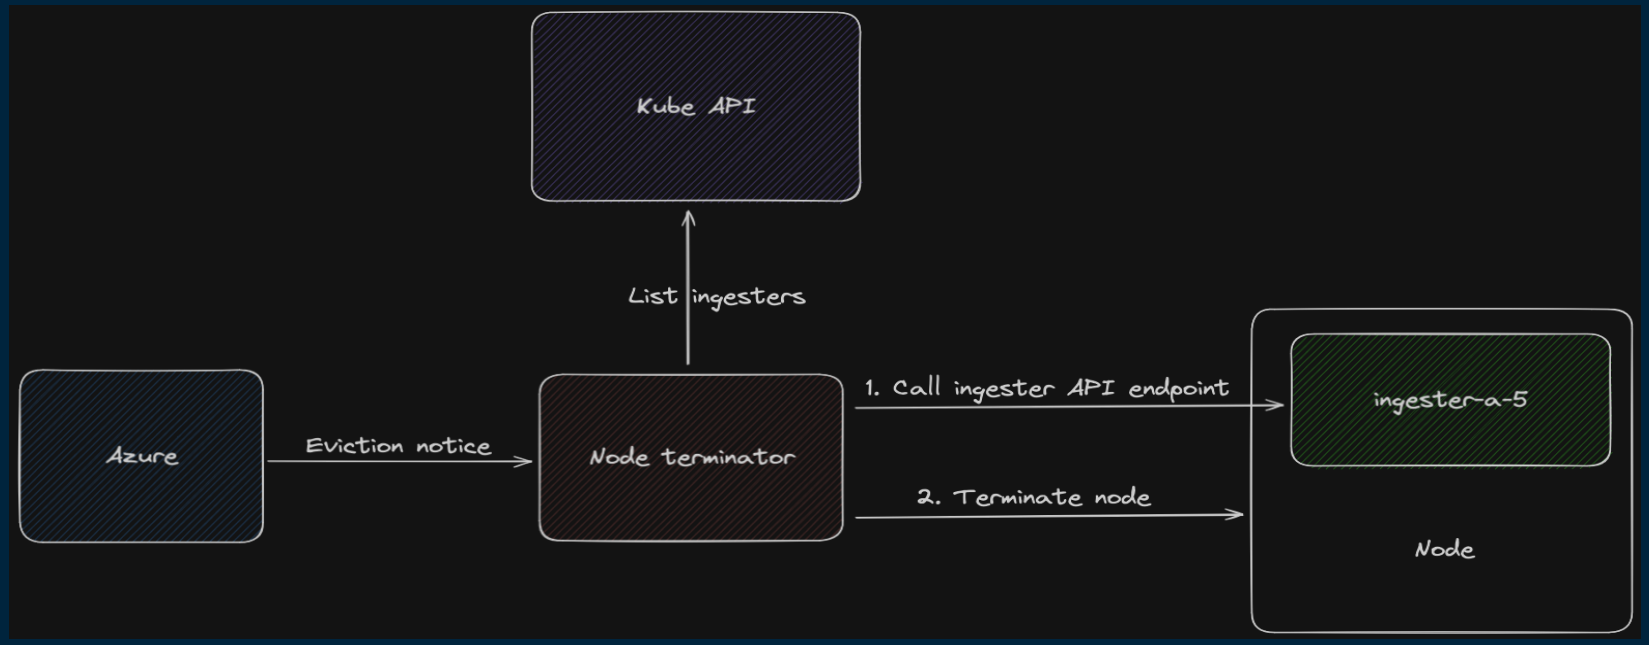

MAERSK

Gains 💪

CPU savings are equivalent to ~$13700 per month

MAERSK

Q&A

- Ingester CPU utilisation balloons during rollouts because metric series proliferate
- Ingesters cannot always stay in the ring because of spot node evictions
- The introduction of a new API endpoint allows ingesters to conditionally leave the ring

MAERSK