Taking your cluster to the Doctor: Keeping Elasticsearch Healthy

—

By Jurgens du Toit / jrgns / eagerelk.com

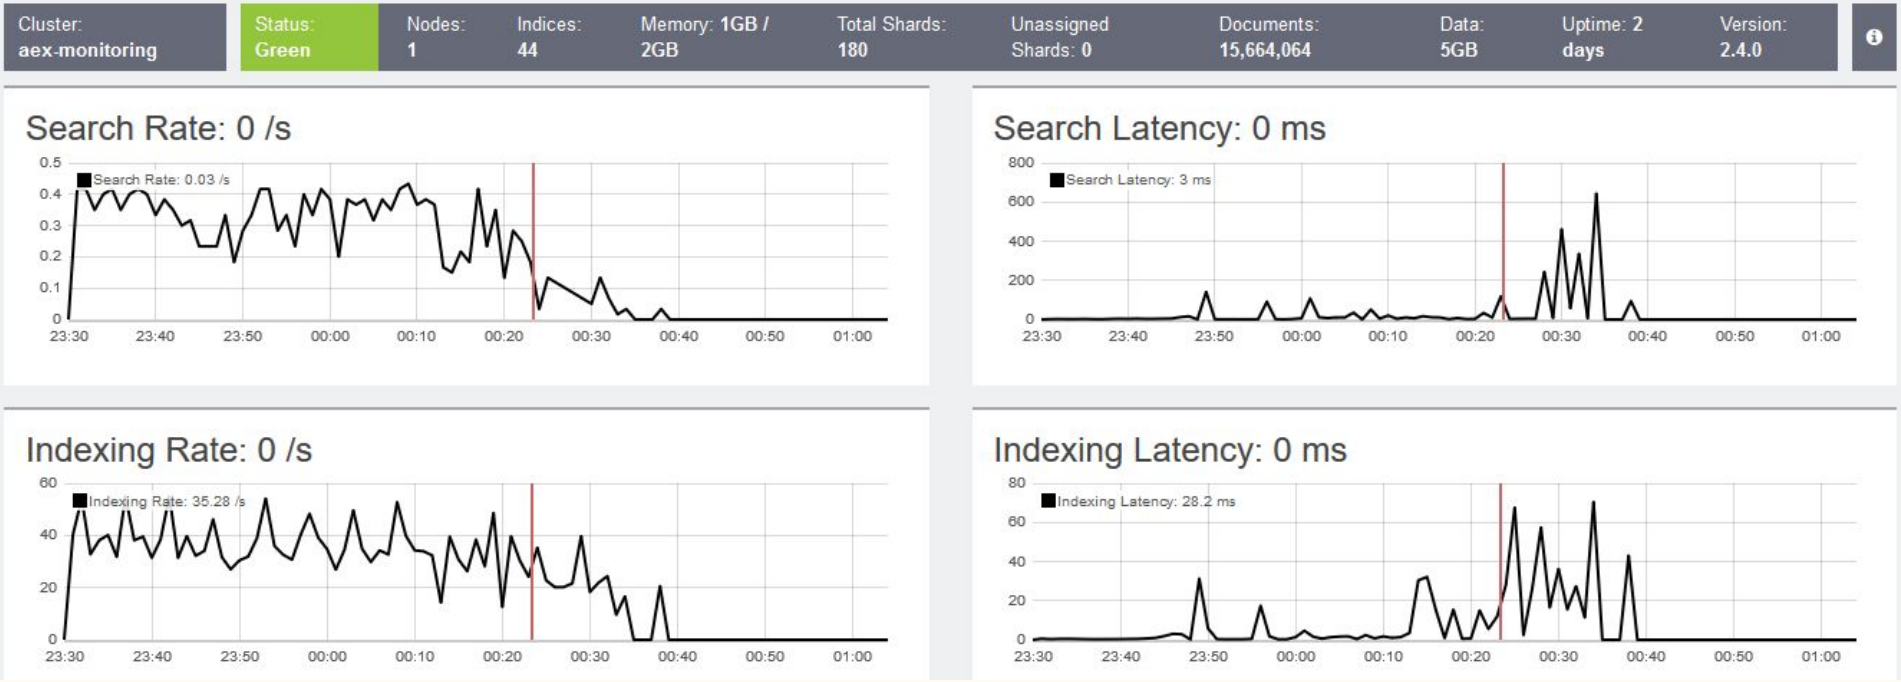

A Failing Cluster

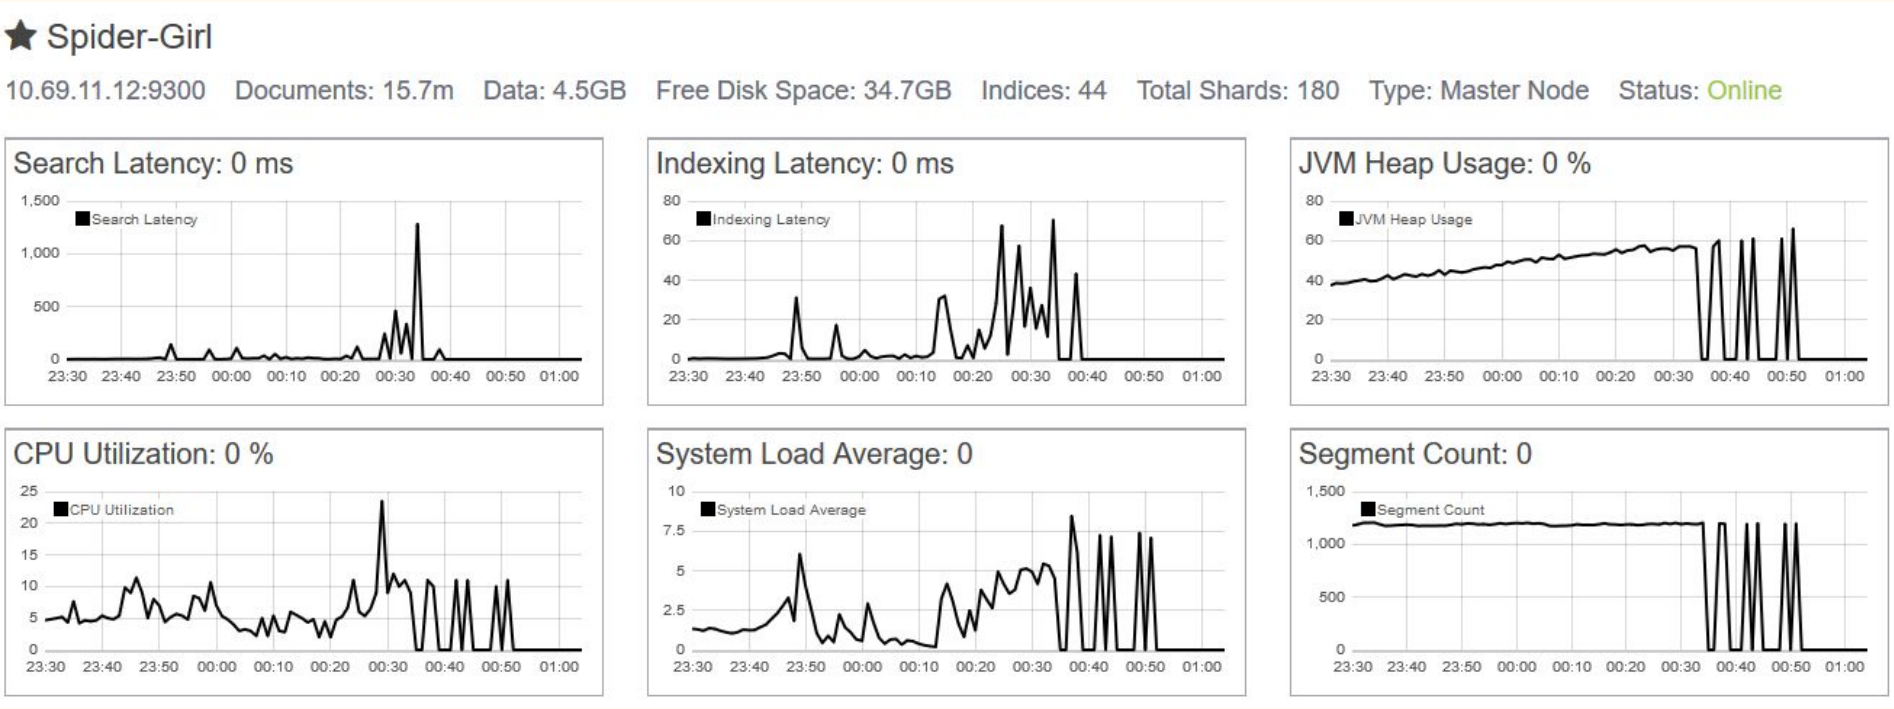

A Failing Cluster

Cluster

- Cluster Health
  - Green is Good
  - Yellow is OK
  - Fear the REDS
- Number of Nodes
- Node System Health
  - CPU
  - Memory
  - Disk I/O

Cluster fixes

- Check number of replicas
- Check stuck or transferring shards

Memory

- Too small heap size
- Too large heap size
- Too little system memory
- Swap is on

Memory fixes

- ES_HEAP_SIZE = 50% of SM
- Turn swap off
-

Disk Space

- Less than 20% left

Disk Spaces fixes

- Cluster data on separate FS
- Clean out old data
- Hot / Warm / Cold Setup

Segments & Shards

- Too many segments
- Shards larger than 50GB
- Unassigned / relocating shards

Segment & Shard Fixes

- Force merge older indices
- Properly profile your data
- Properly profile your server
- Choose number of shards based on profiling

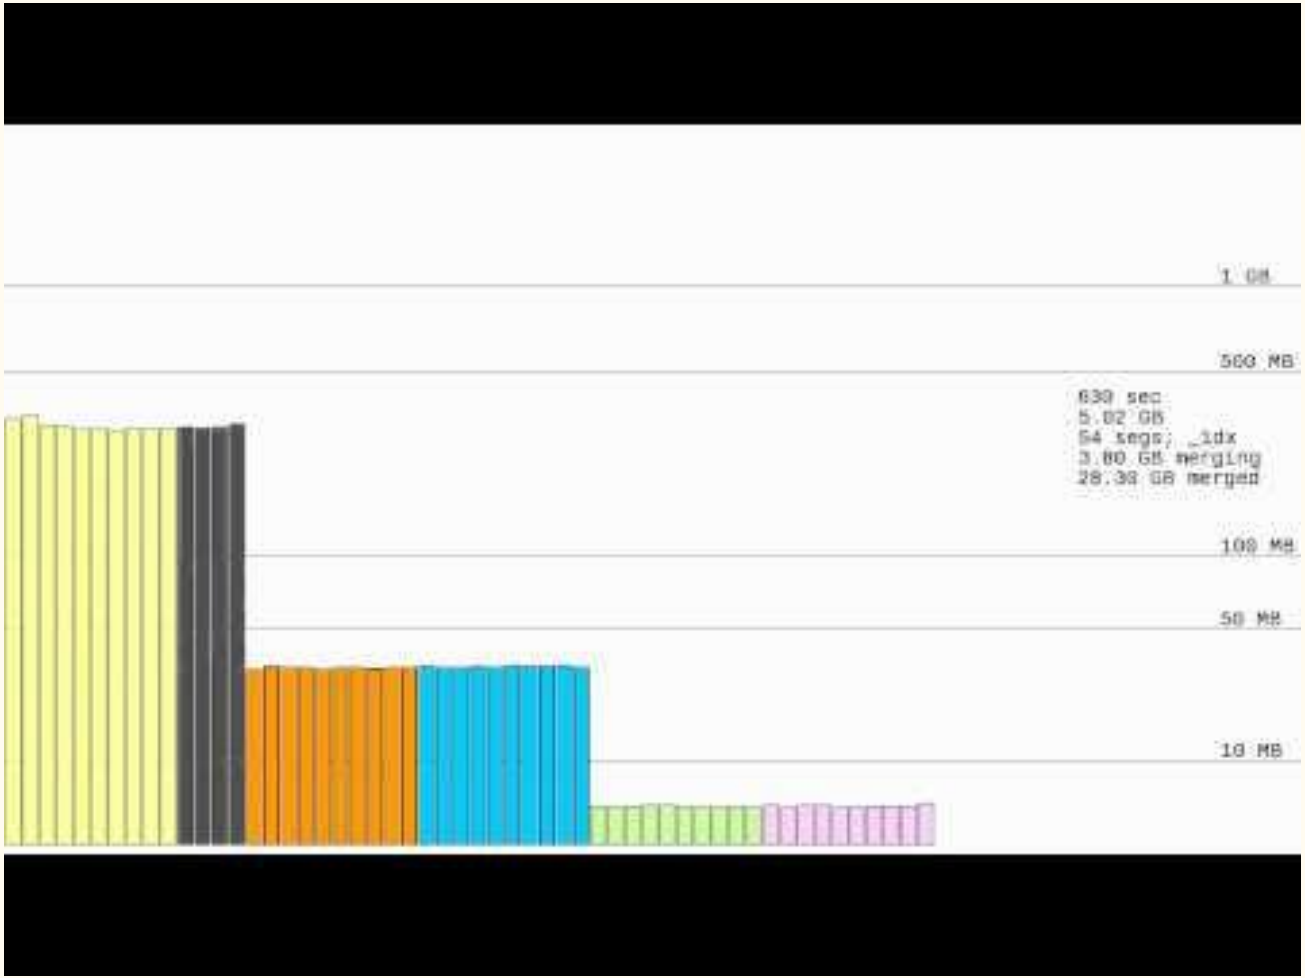
1 GB
500 MB
630 sec
5.02 GB
54 segs; _1dx
3.80 GB merging
28.38 GB merged
100 MB
50 MB
10 MB

JVM

- It’s all about the memory
- Duration of GCs
- Number of GCs
- Separate pool sizes

JVM Fixes

- Get your memory and shard settings right
- OR, spend 15 years learning about how the JVM manages memory

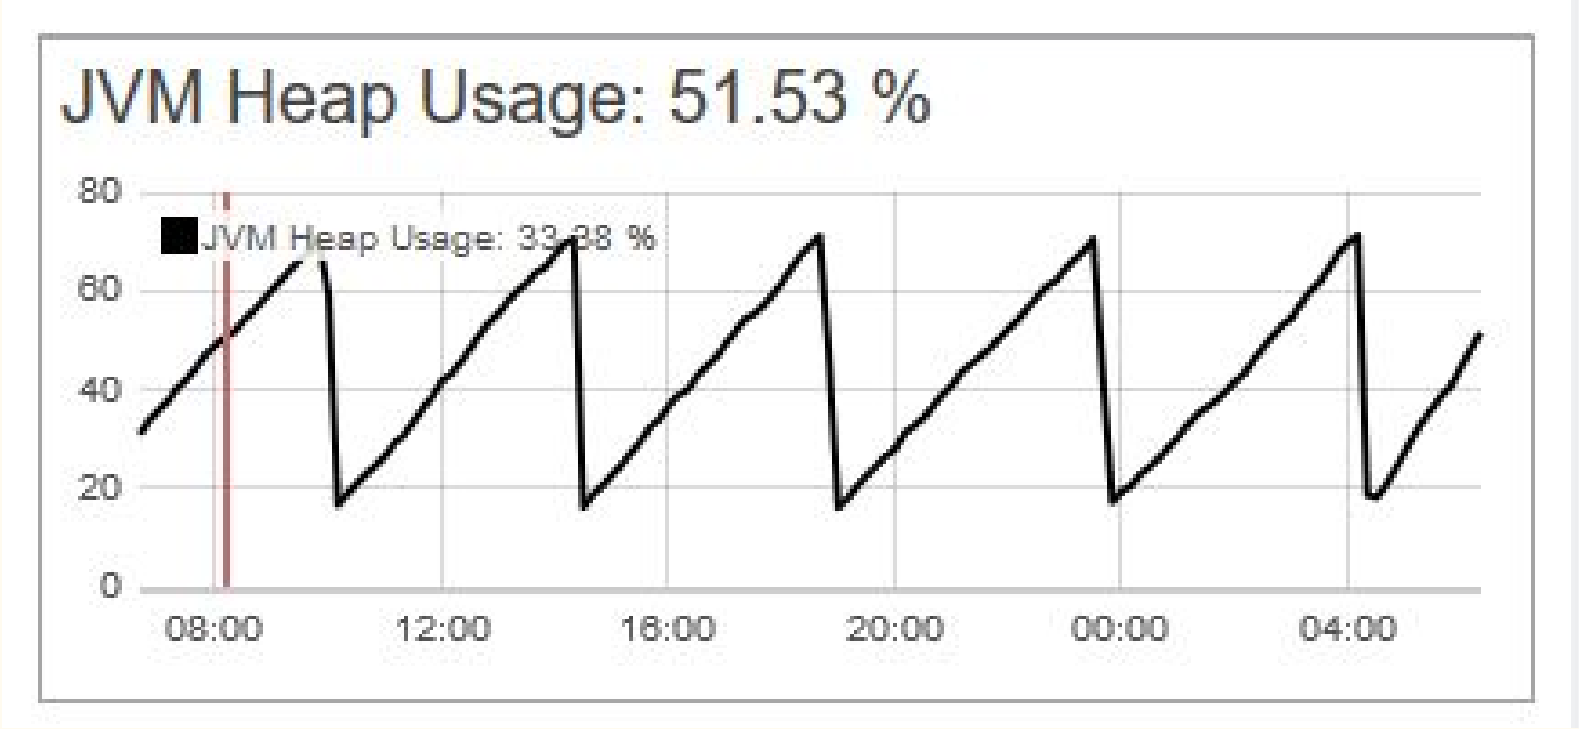
JVM Heap Usage: 51.53 %
JVM Heap Usage: 33.38 %
80
60
40
20
0
08:00
12:00
16:00
20:00
00:00
04:00

Good GC

Index Rate

- Affected by
  - Memory
  - Disk
  - Segments
- Declining indexing rate predicts trouble
- Watch out for declined requests

Index Rate Fixes

- Get your memory settings right
- Quicker disks
- Larger index thread pool

Search Rate

- Affected by
  - Memory
  - Disk
  - Nodes & Shards
- Declining search rate predicts trouble
- Watch out for declined requests

Search Rate Fixes

- Get your memory settings right
- Get your number of shards and nodes right
- Larger search thread pool

Circuit Breakers

- Indicates issues with
  - Fielddata
  - Requests
  - Overall

Why is it important

- Can indicate configuration issues
- Can also point to application and usage issues

Resources

- https://www.datadoghq.com/blog/monitor-elasticsearch-performance-metrics/
- https://www.oreilly.com/ideas/10-elasticsearch-metrics-to-watch
- https://www.elastic.co/guide/en/marvel/current/installing-marvel.html
- https://www.opsdash.com/blog/elasticsearch-monitoring.html

Tools

- DataDog
  https://www.datadoghq.com
- Marvel
  https://www.elastic.co/products/marvel
- OpsDash
  https://www.opsdash.com/
- ProxES
  http://proxes.io / http://github.com/EagerELK/proxes

Questions?

@jrgns

jrgns@jrgns.net

eagerelk.com