

Scaling Aerospike

Lessons for Data-Intensive Application Developers



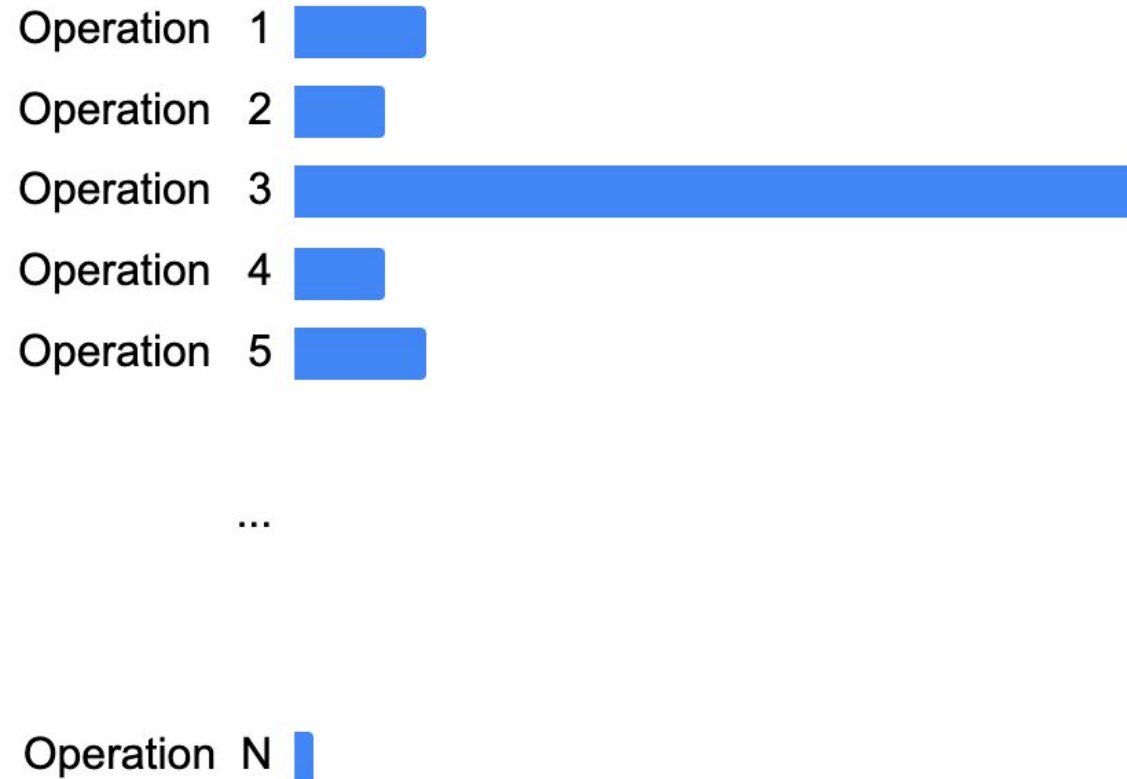
Latency Distribution of Sub-Operations of a Task



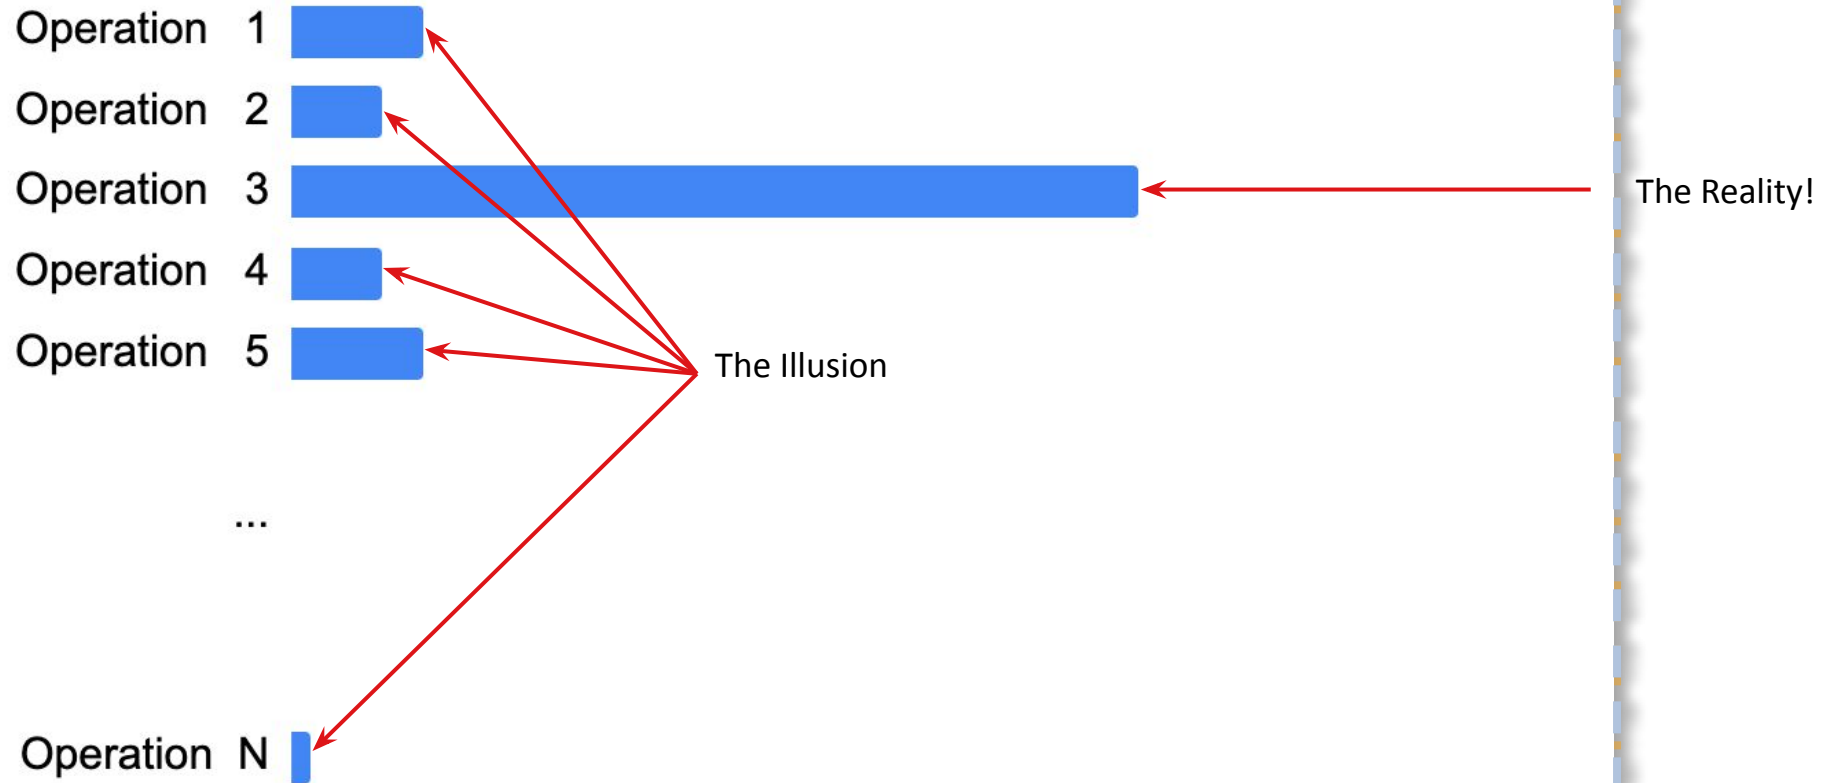
Latency Distribution of Sub-Operations of a Task

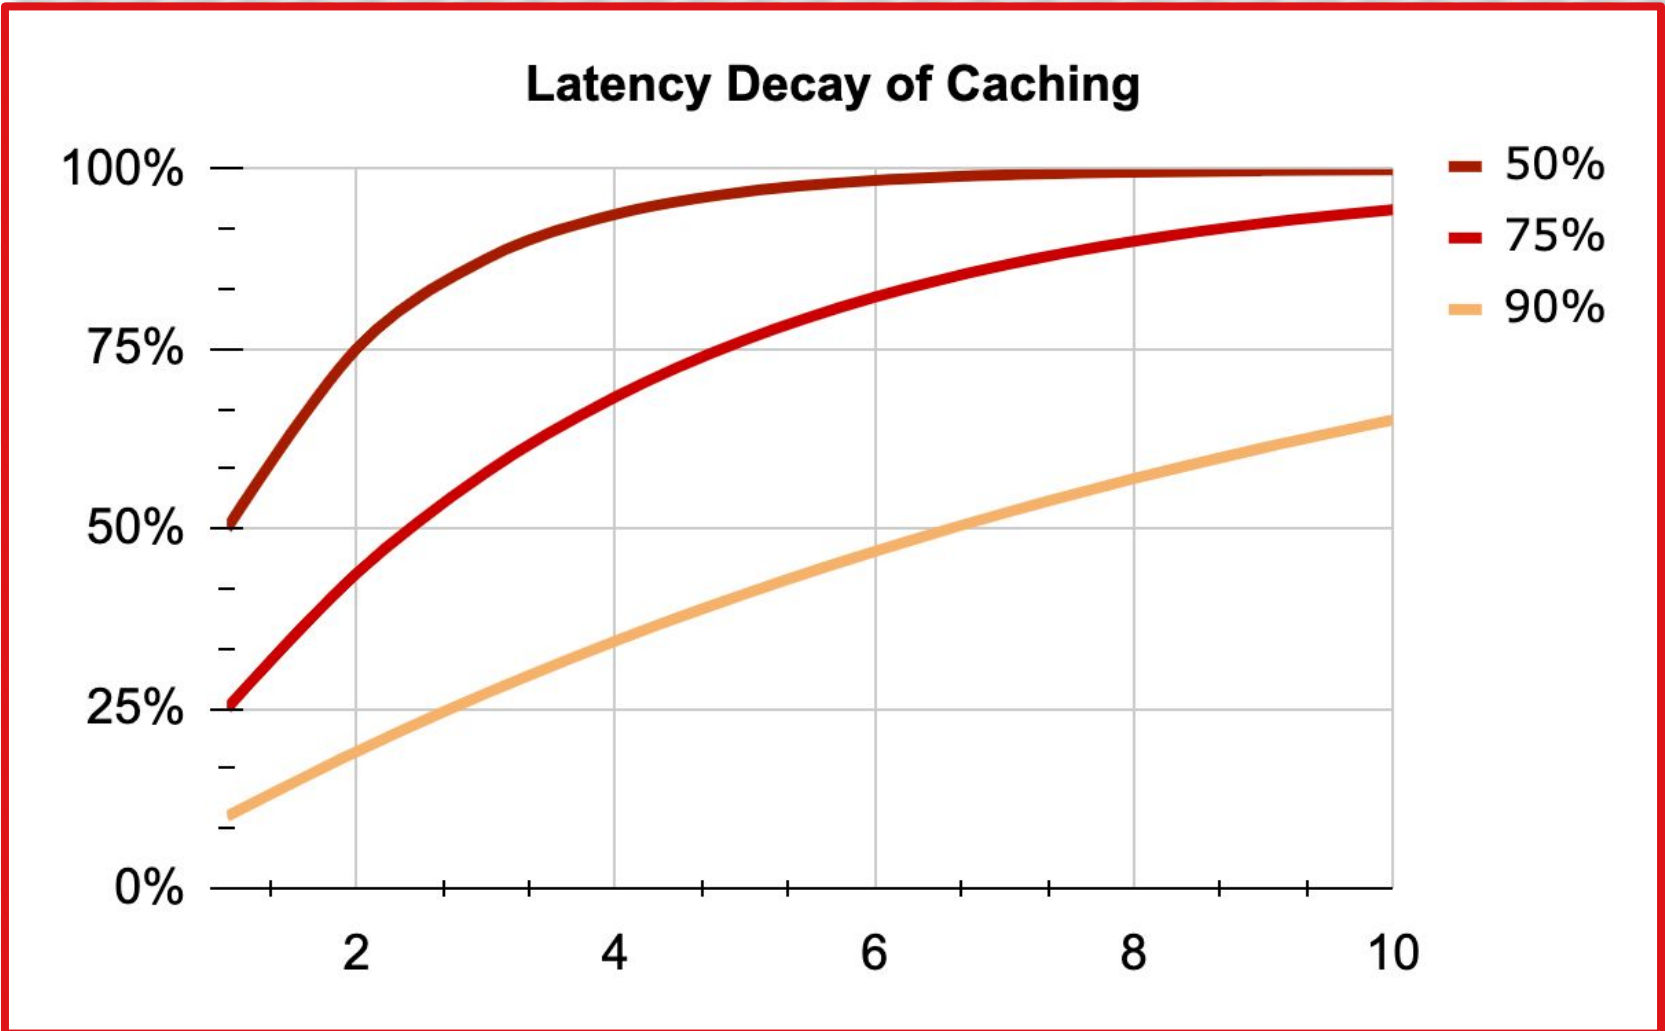


Latency Distribution of Sub-Operations of a Task with a Cache



Latency Distribution of Sub-Operations of a Task with a Cache





Aerospike NoSQL Database



Performant

- Serves data from disk with speed similar to a Cache.
- Handles high throughput.
- Scales From GB to PB.
- Predictable performance.



Highly Available

- Resilient.
- Self-healing Capability.
- Multi-DC Replication.
- Rack/AZ Awareness.



Efficient

- Low cost.
- Use resource efficiently.
- High Availability with RF=2.
- Reduce CO2 emissions!



Interoperable

- Key/value, document, graph.
- Kubernetes Operator.
- AWS Graviton compatible.
- Kafka, Spark, Pulsar, Trino, Elasticsearch integration.

Software Design

- Software design is a practice in prioritisation of competing concerns.
- When prioritising, some trade-offs have to be made.
- Trade-offs cause imperfections!
- Two solutions may solve the same problem, but they have different imperfection.



Example: High-Availability Trade-Off/Imperfection

When building a Highly Available distributed system, would you design it to be:

1. Efficient during failures, but only effective during normal times.
2. Efficient during normal times, but only effective during failures.



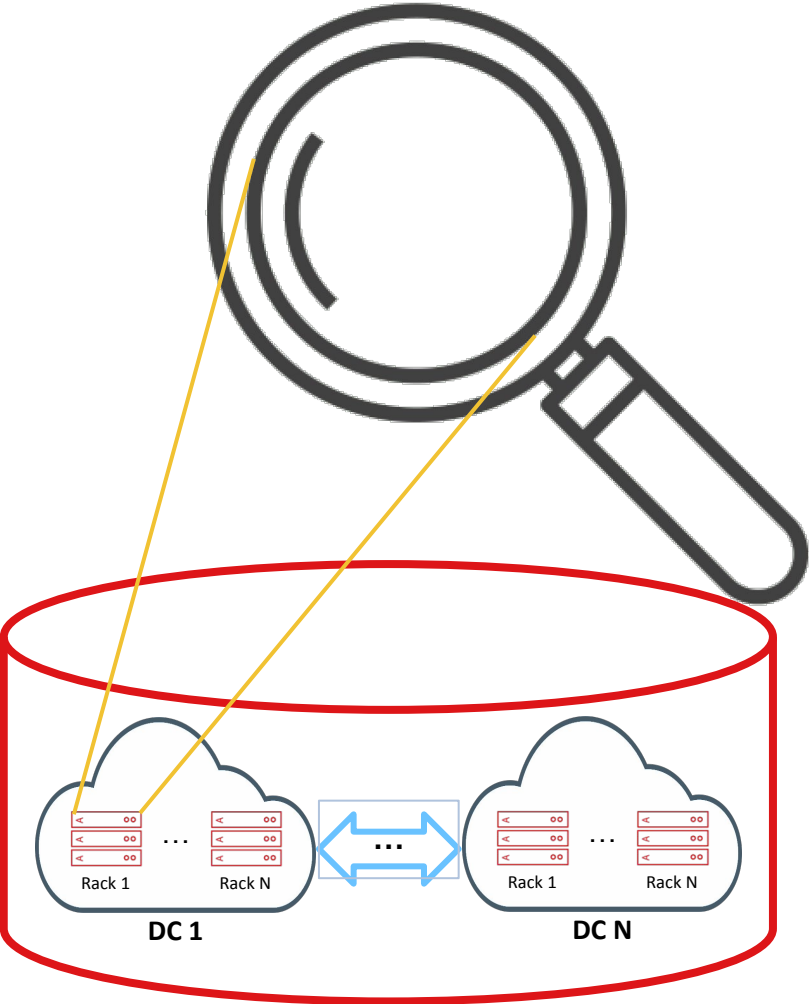
“It Depends!” 🙄

Let's have a look at AWS SLAs:

- EC2 instance level uptime during a month: 99.5% (Uptime: 716:20, Downtime: 3:36)
- EC2 region level uptime during a month: 99.99% (Uptime: 719:56, Downtime: 0:04)



Demystifying Aerospike Architecture



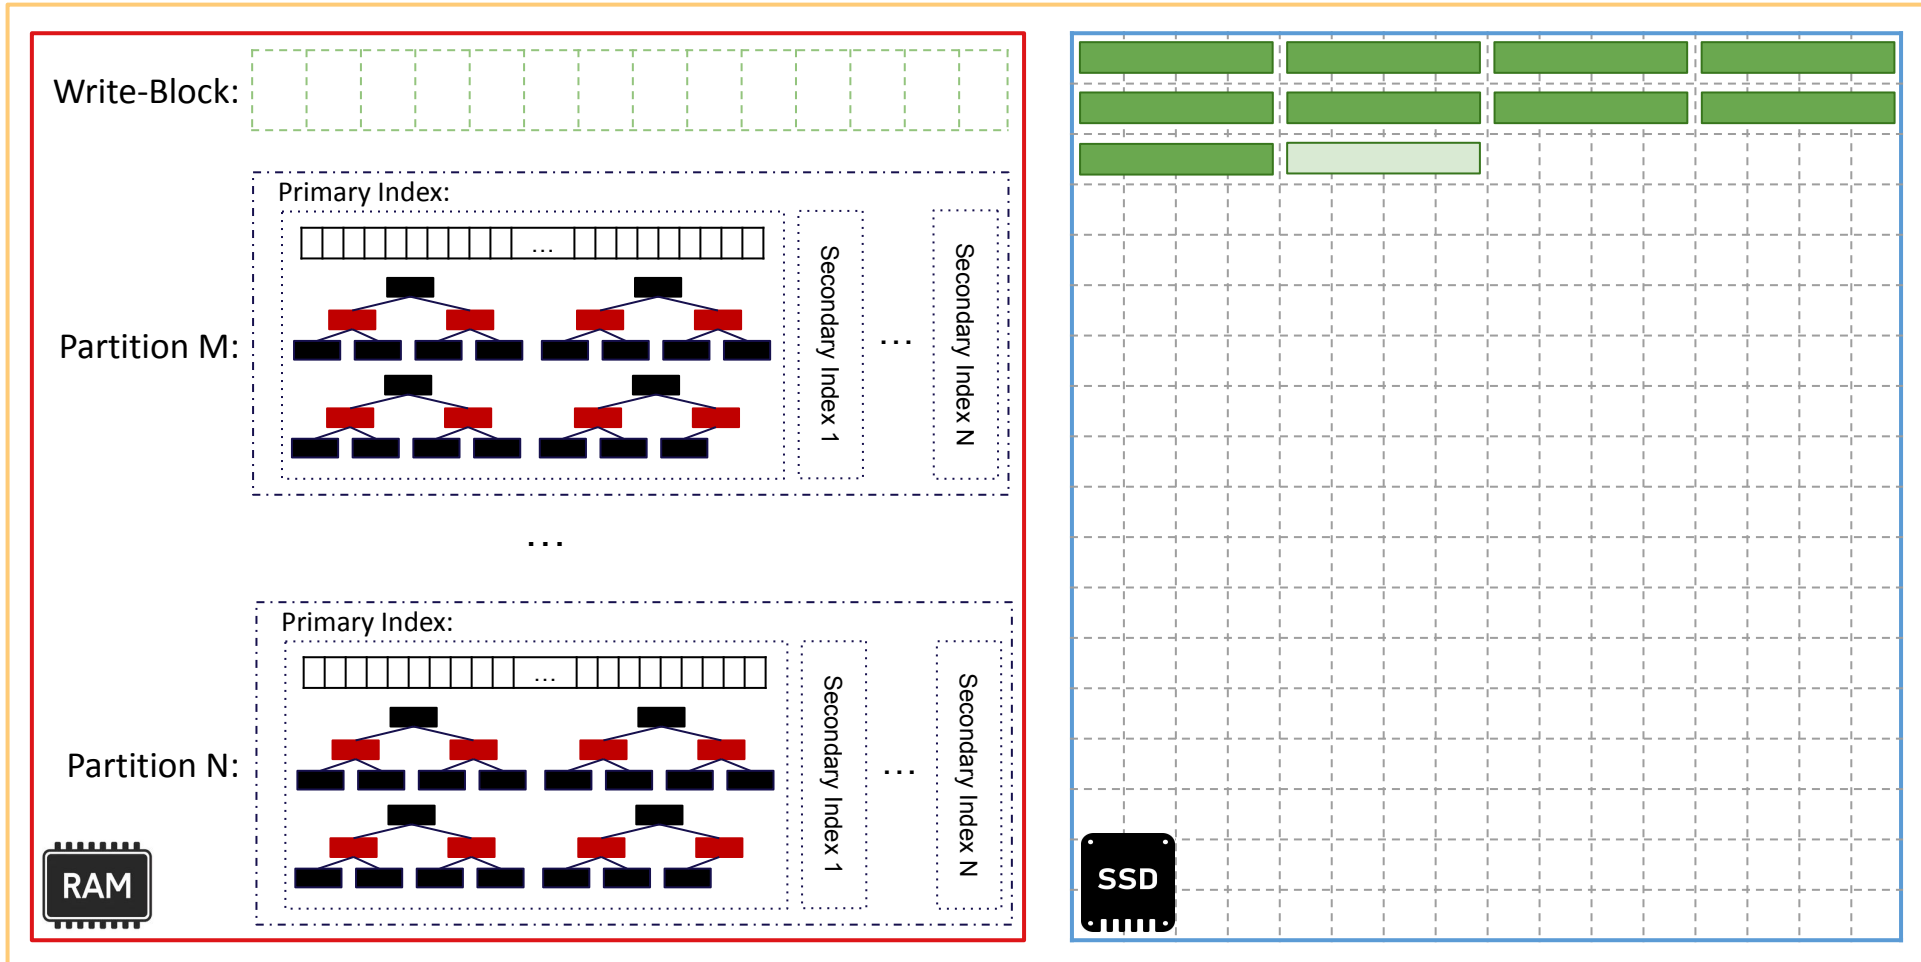
Data Distribution

Nodes: A, B, C

4096 Partitions

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| A | B | A | C | C | B | A | C | B | A | B | C | A | B | C | A | C | B | C | A | C | B | B | A | ... | A | B | A | C | C | B | A | C | B | A | B | C | A | B | C | A | C | B | C | A | C | B | B | A | | | | | | | | | | | |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 | 40 |

Node Architecture



Write-Block Magic

- A Write-Block gets flushed to the disk:
 - i. Overwritten In-place: every second.
 - ii. Immutably: When it doesn't have enough space for a new record.
- Write-Block guarantees both durability and fast access on disk. databases write the data on the disk twice:
 - i. Durability (eg. OpLog, WriteAheadLog, CommitLog, Journal)
 - ii. Fast access (Some kind of B-Tree or B+Tree)
- Some other benefits of Write-Block:
 - i. Because Write-Block size is fixed, recovering and reusing space is simple.
 - ii. Disk would never need to be defragged.

The Result



Example* of Performant and Efficient

| Name | ▲ Instance ID | Instance state ▼ | Instance type |
|--------------------|---------------------|------------------|---------------|
| aerolab4-Demo_1 | i-04a18f4d5c3849854 | ✔ Running | i4i.4xlarge |
| aerolab4-Demo_2 | i-0d8dd3bc198e5eea3 | ✔ Running | i4i.4xlarge |
| aerolab4-Demo_3 | i-06d83bfea16d28e76 | ✔ Running | i4i.4xlarge |
| aerolab4-Demo_4 | i-0a3ceaba775d77b00 | ✔ Running | i4i.4xlarge |
| aerolab4-Perseus_1 | i-097d76344a2fbf4e3 | ✔ Running | c5.9xlarge |

| Instance Name | vCPU Count | Memory (GB) | Storage (GB) |
|---------------|------------|-------------|--------------|
| i4i.4xlarge | 16 | 128 | 3750 |

| Node Count | RF | Capacity (TB) | Throughput |
|------------|----|---------------|------------|
| 4 | 2 | 3 | 600 K TPS |

| Region | AWS Annual Reserved Cost |
|---------|--------------------------|
| London | \$36,282 |
| Ireland | \$34,483 |
| Ohio | \$31,203 |

* This example can be presented to you as a live demo upon request.

Thank you :)



Developer Hub
—AEROSPIKE—

<https://developer.aerospike.com>



GitHub

<https://github.com/aerospike-examples>



Documentation
—AEROSPIKE—

<https://docs.aerospike.com>



—AEROSPIKE—
{ developers }