
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

ANT315

Optimizing Amazon OpenSearch Service domains for size and cost

Jon Handler (He/Him)

Senior Principal Solutions Architect,
Search Services
AWS

Aish Gunasekar

Solutions Architect
AWS

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

In this session

Amazon OpenSearch Service core concepts

Sizing/costing examples

Sizing: storage

Sizing: compute

Cost optimization: infrastructure

Cost optimization: workload

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Core Concepts

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



OpenSearch

Community-driven, open-source search and analytics suite. Consists of a search engine, OpenSearch; a visualization user interface, OpenSearch Dashboards; as well as a series of functionality-adding tools and plugins.



Amazon OpenSearch Service

Fully managed service that makes it easy to deploy, manage, and run OpenSearch cost effectively with industry-leading reliability, scalability, and security.

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

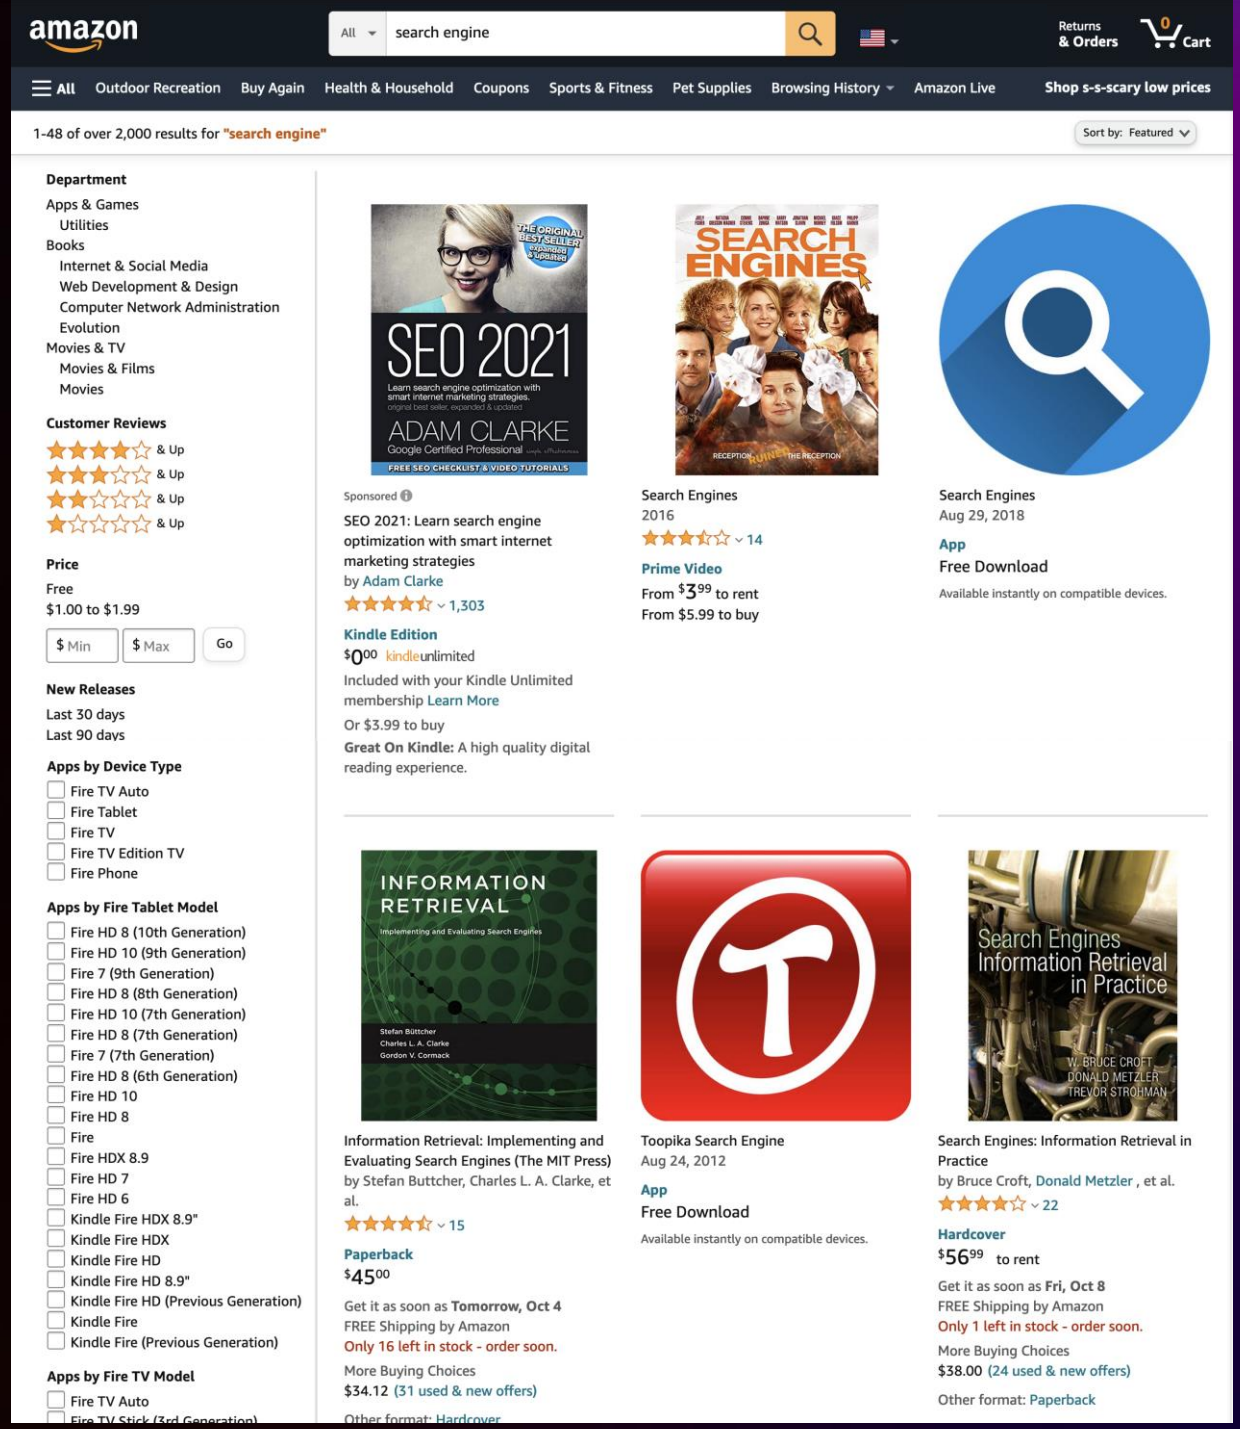
Searching

Search engines match terms to objects in a catalog

Match free text or structured data

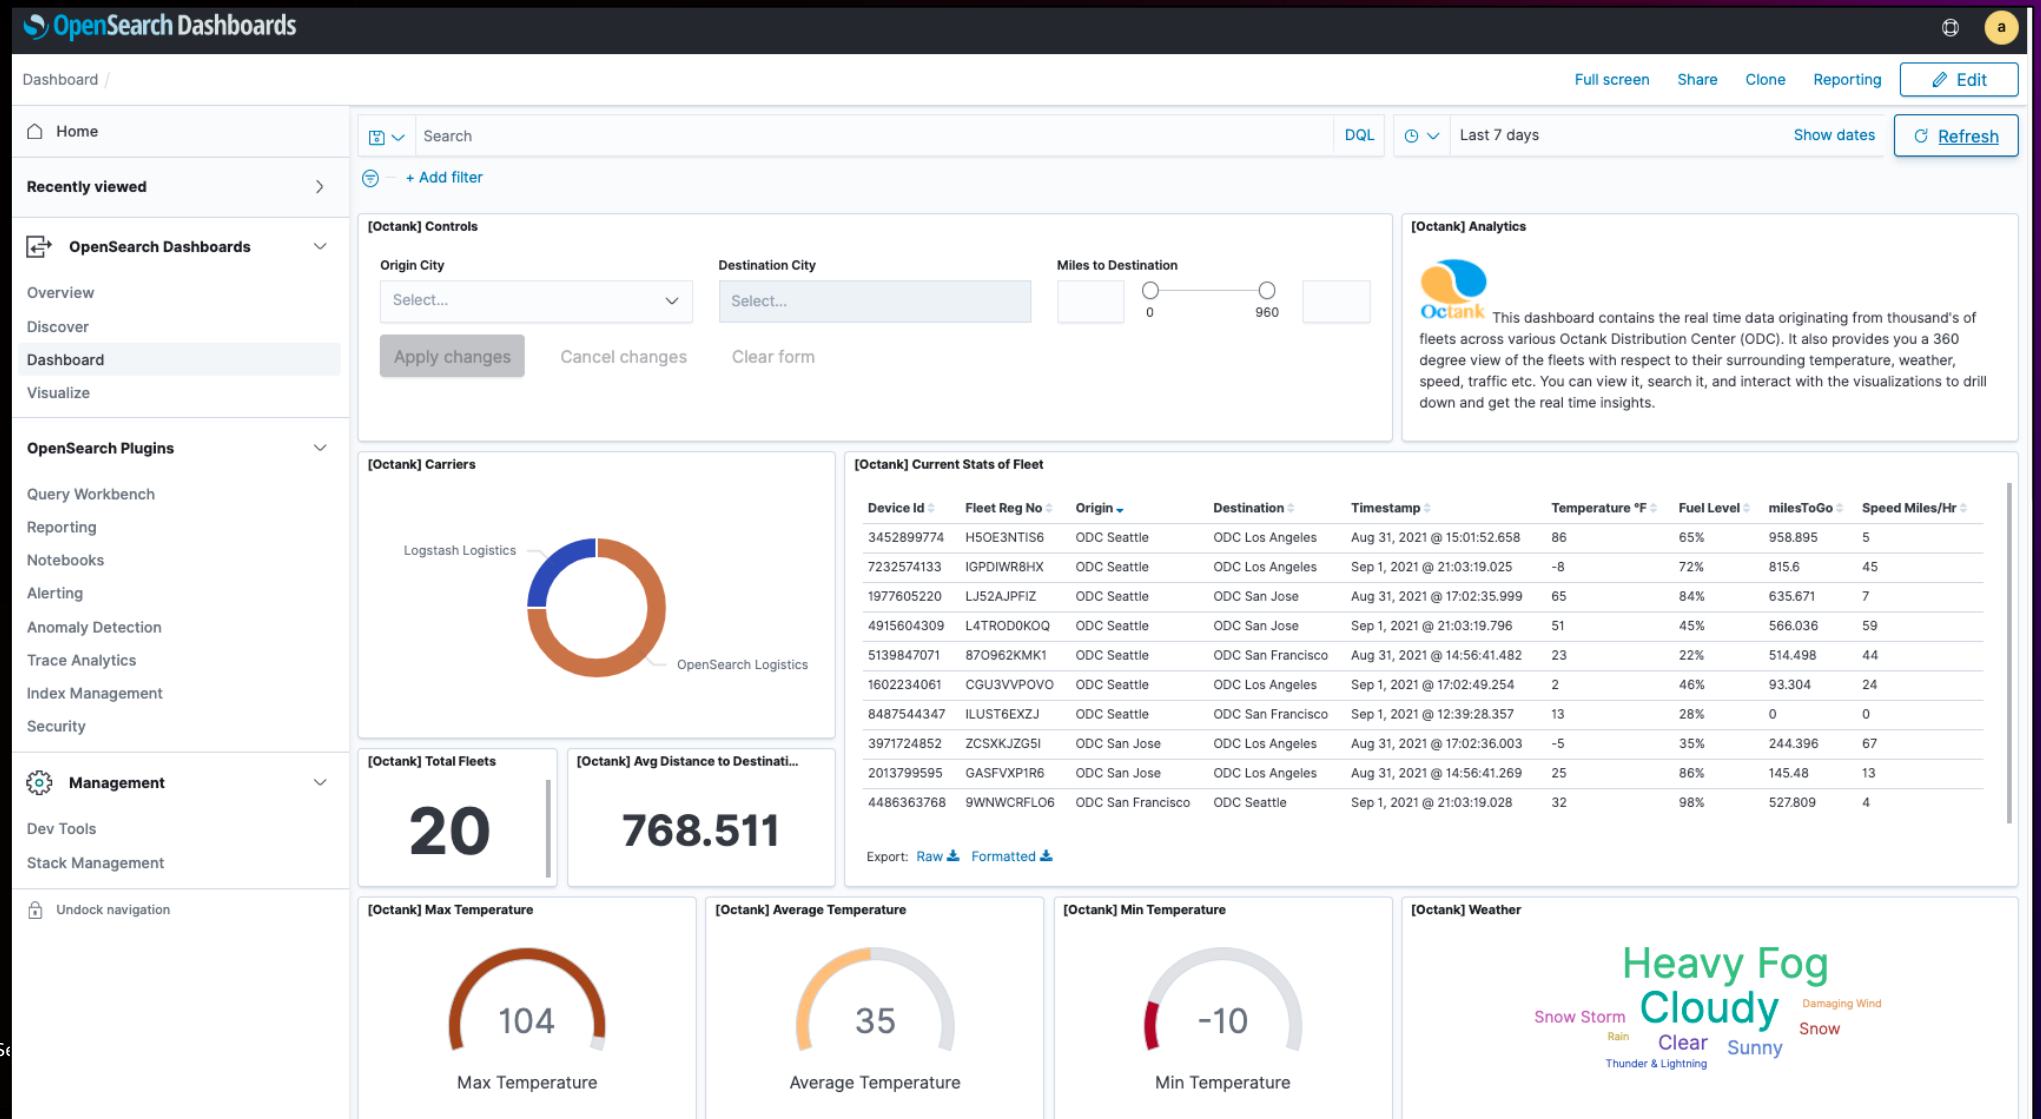
Relevance determines sort order

Facets enable users to drill in to the search results to narrow by attribute value



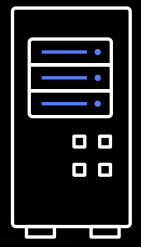
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Analyzing streaming data (logs)

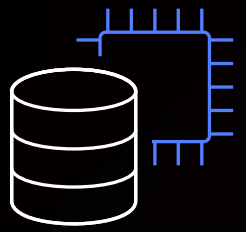


© 2022, Amazon Web Se

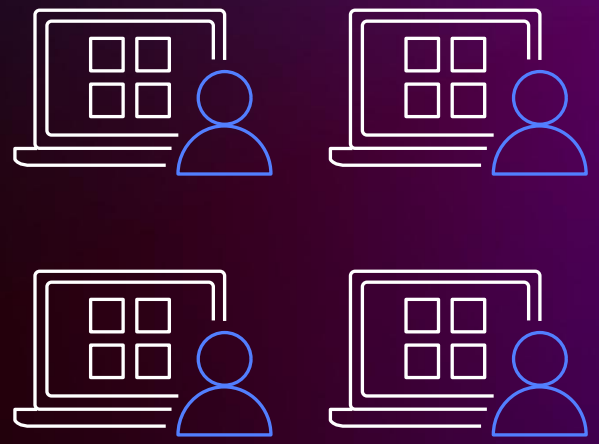
Interaction flow



Server, application, network, AWS, and other logs

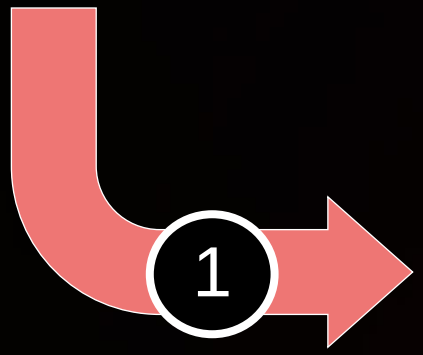


Application data



Application users, analysts, DevOps, and security

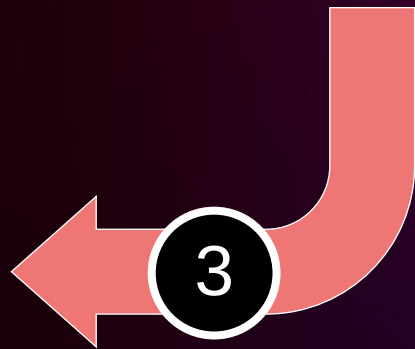
Indexed and instantly searchable



JSON documents



Amazon OpenSearch Service domain

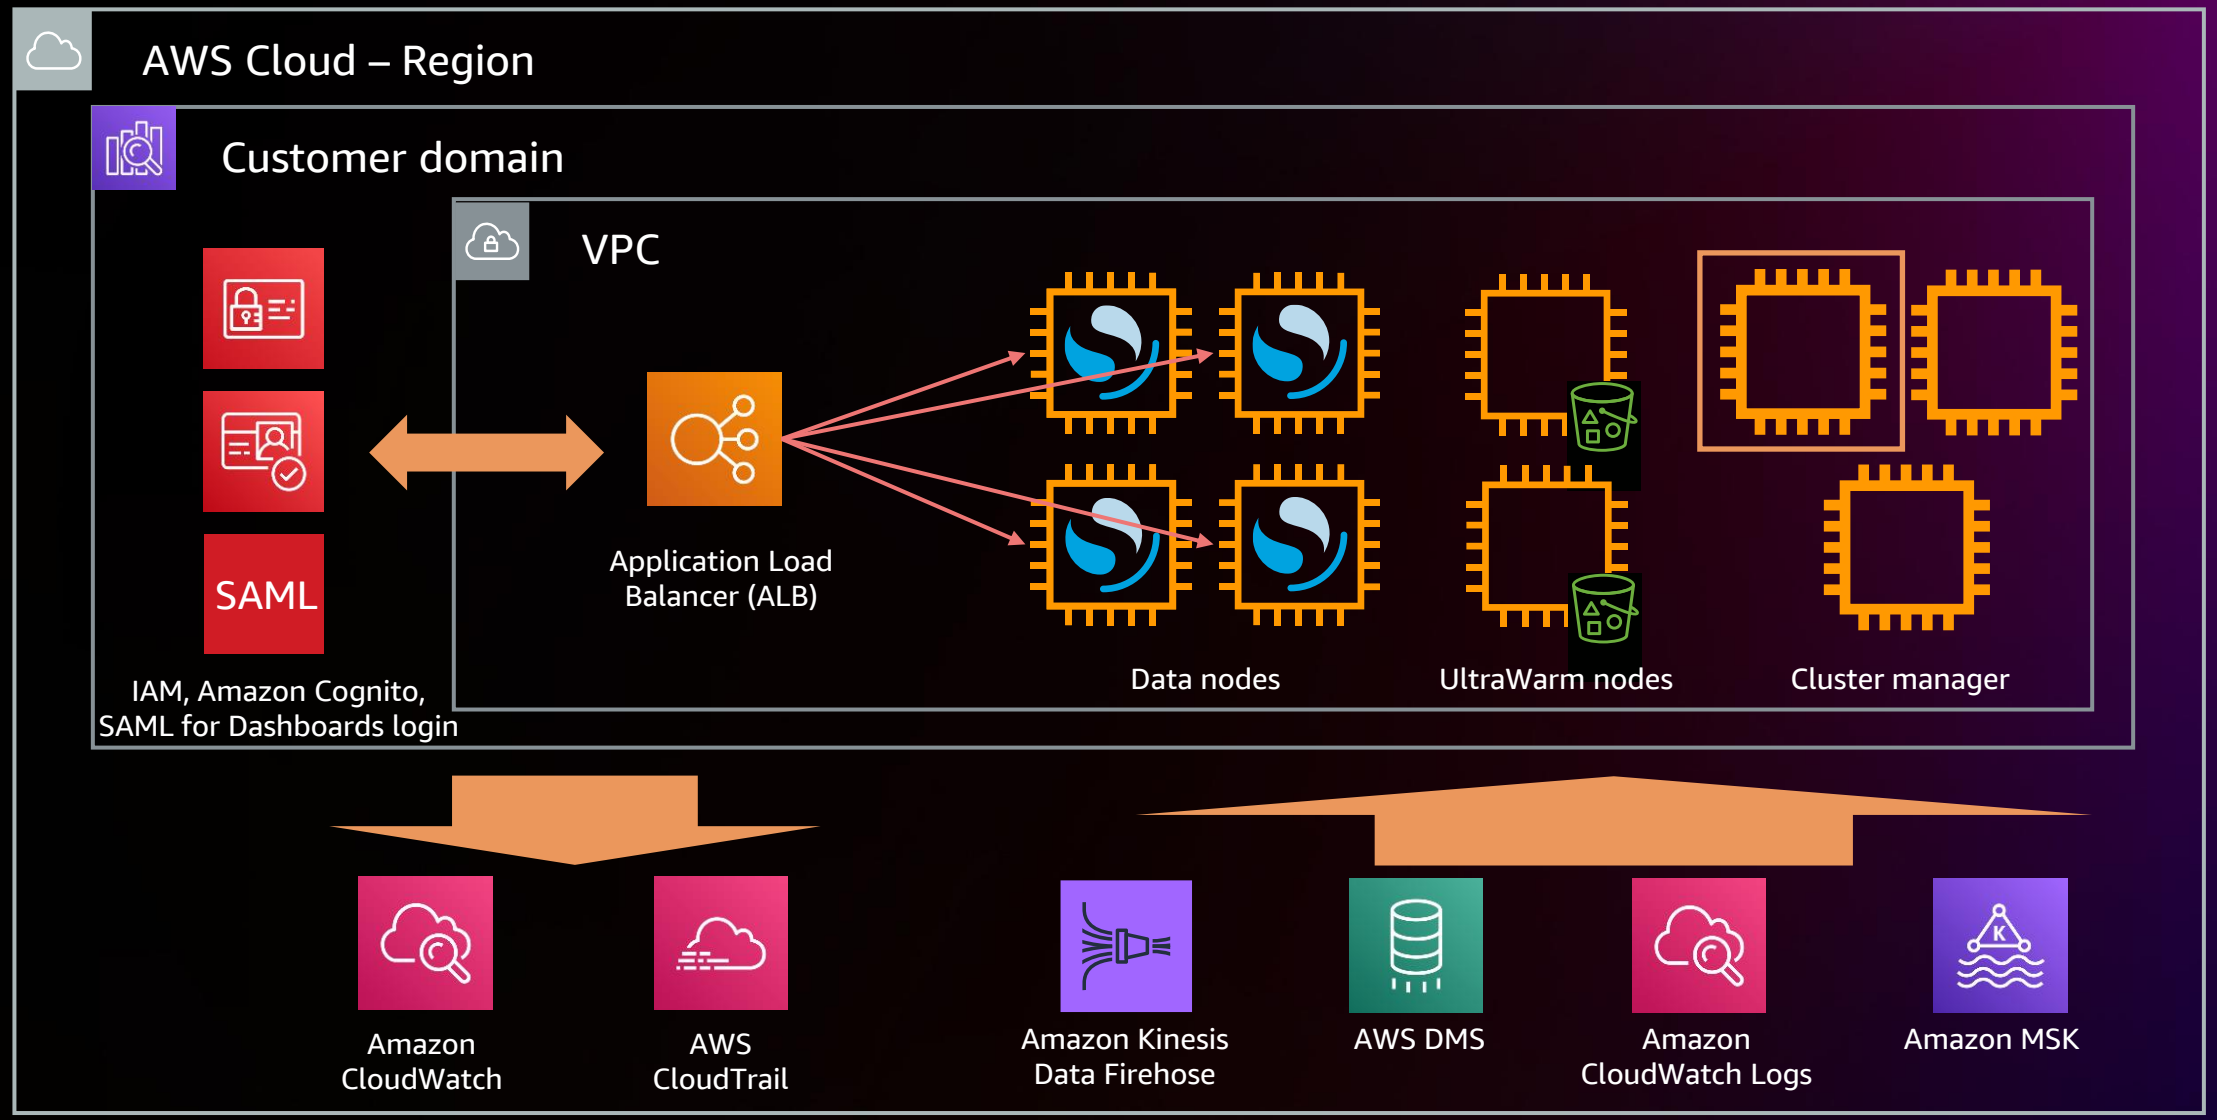


Rest APIs, clients, and OpenSearch Dashboards

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon ES deployment architecture



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data disposition in OpenSearch

Indices are the main container

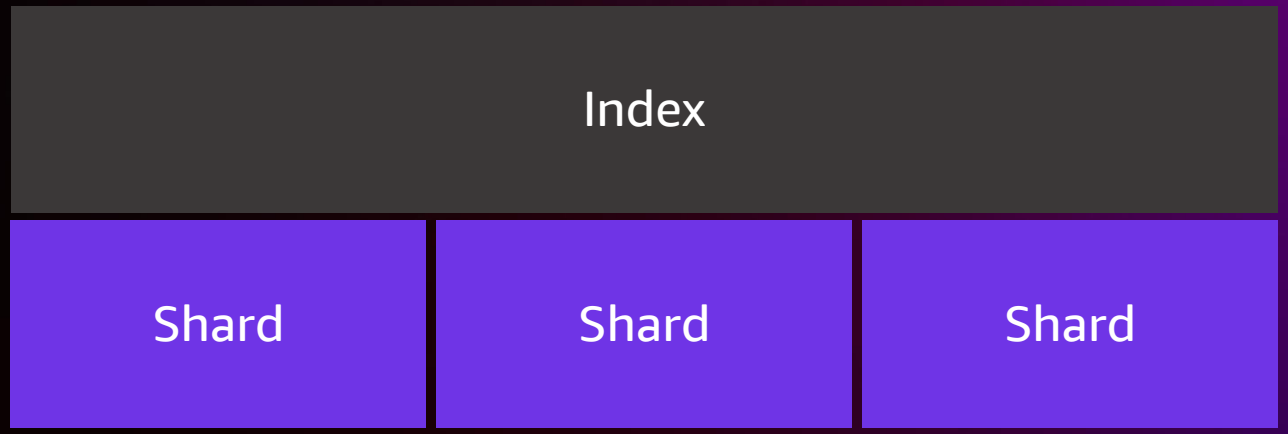


© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data disposition in OpenSearch

Indices are the main container

You create partitions by specifying a shard count – primary and replica



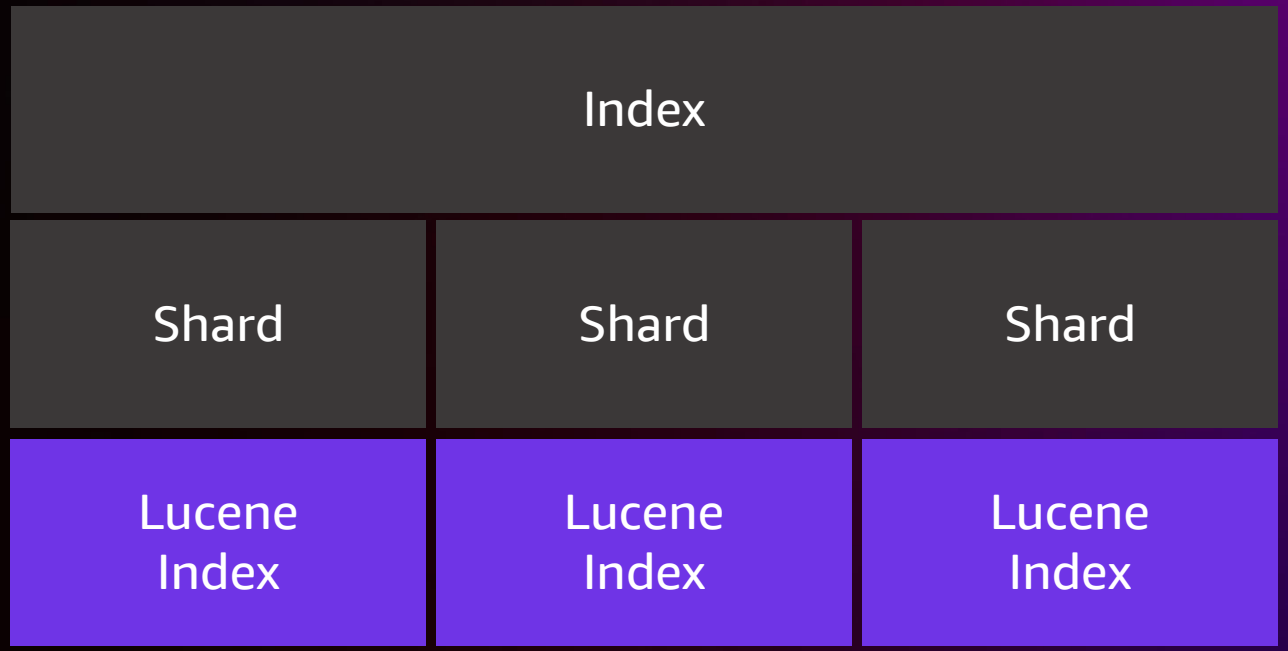
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data disposition in OpenSearch

Indices are the main container

You create partitions by specifying a shard count – primary and replica

Each shard is an instance of Apache Lucene



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

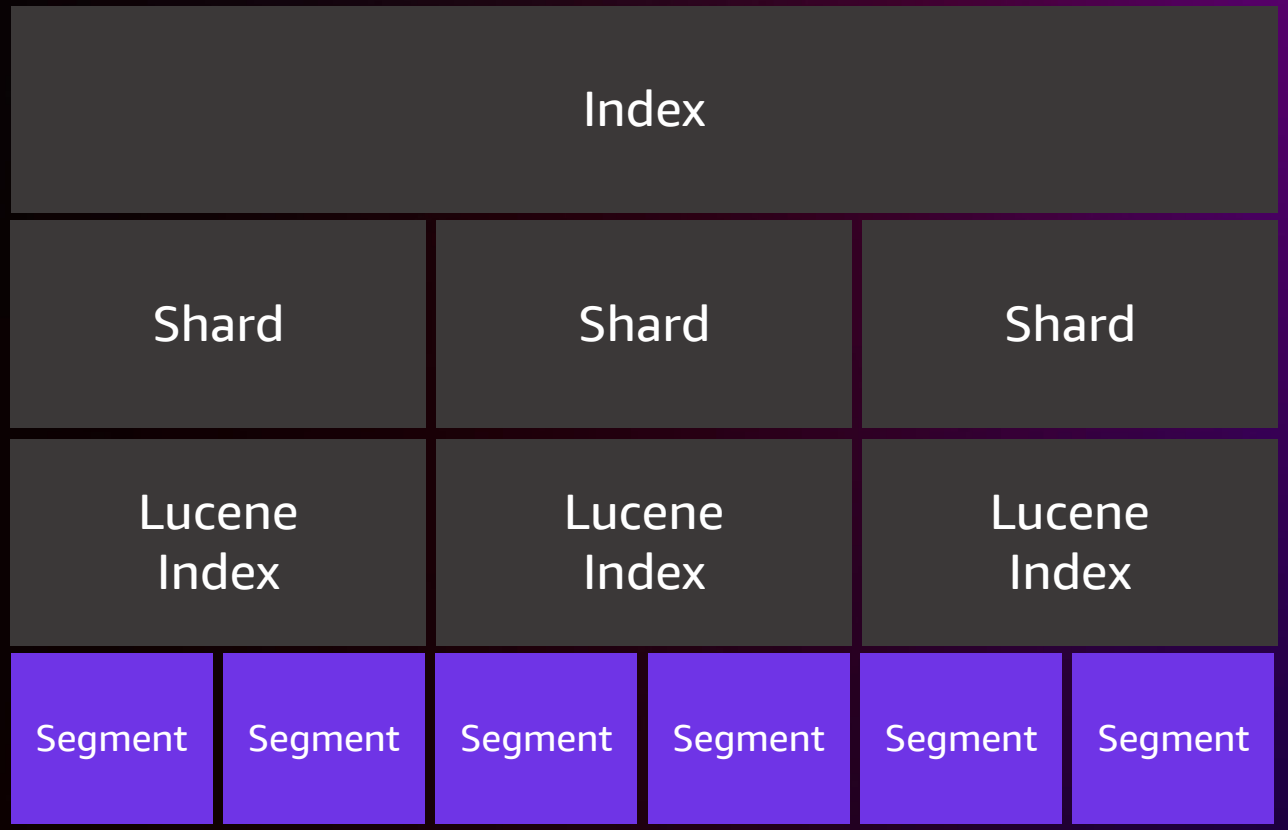
Data disposition in OpenSearch

Indices are the main container

You create partitions by specifying a shard count – primary and replica

Each shard is an instance of Apache Lucene

Lucene indices are comprised of segments



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data disposition in OpenSearch

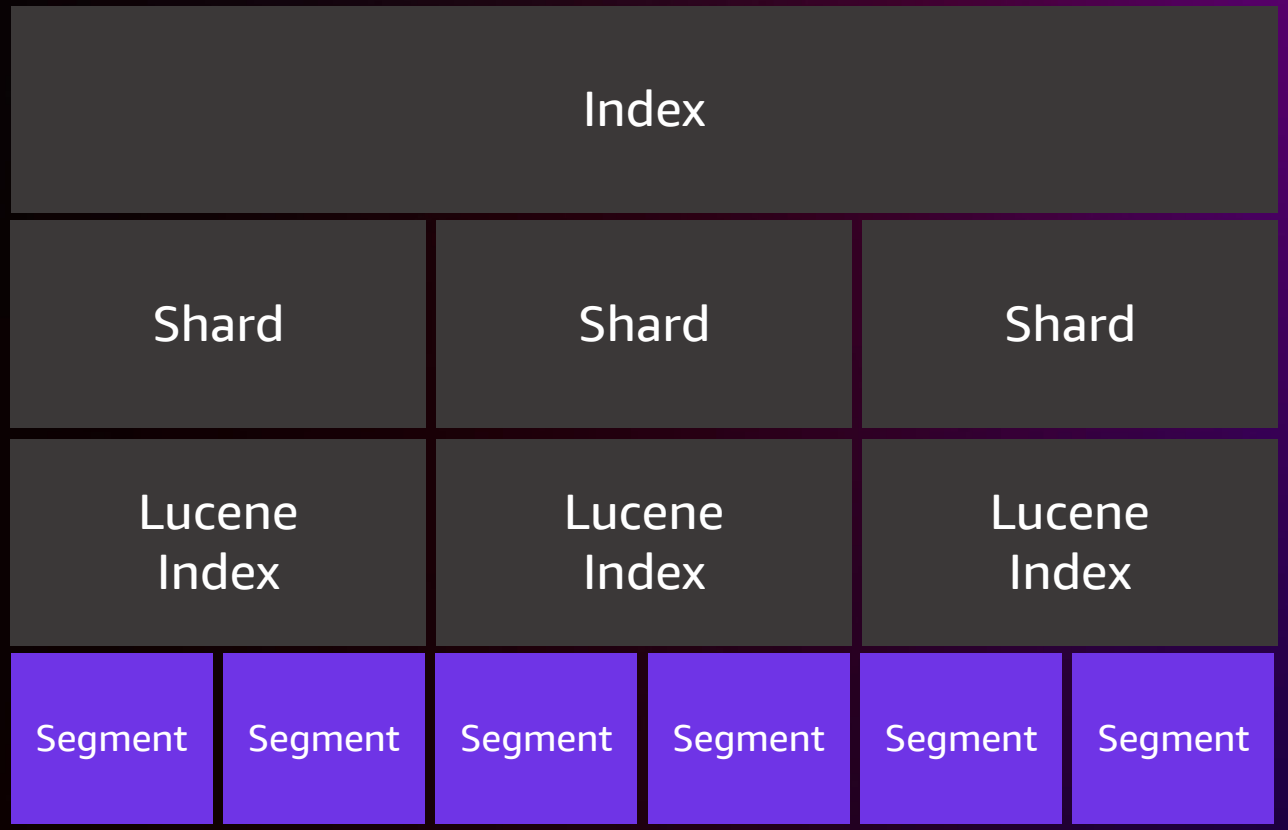
Indices are the main container

You create partitions by specifying a shard count – primary and replica

Each shard is an instance of Apache Lucene

Lucene indices are comprised of segments

OpenSearch distributes requests via the nodes to the shards



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data disposition in OpenSearch

Indices are the main container

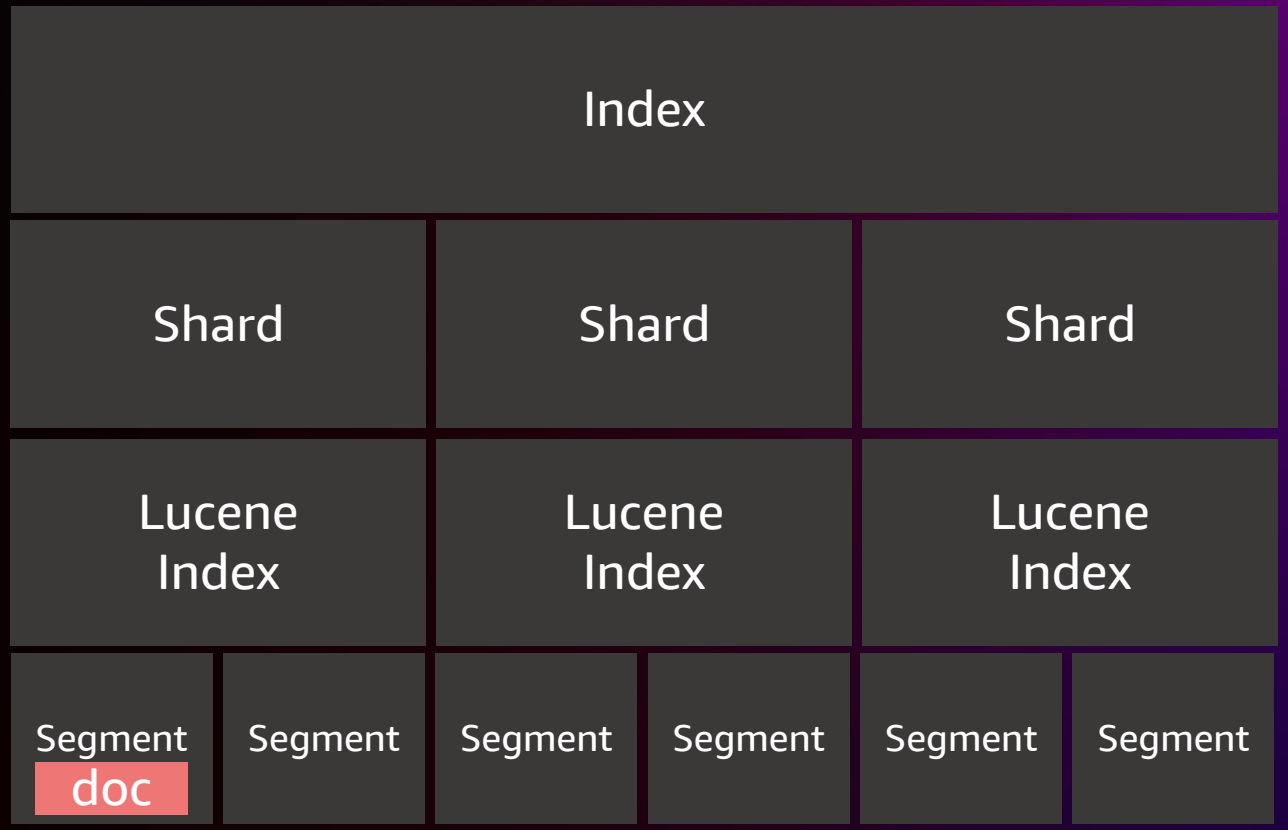
You create partitions by specifying a shard count – primary and replica

Each shard is an instance of Apache Lucene

Lucene indices are comprised of segments

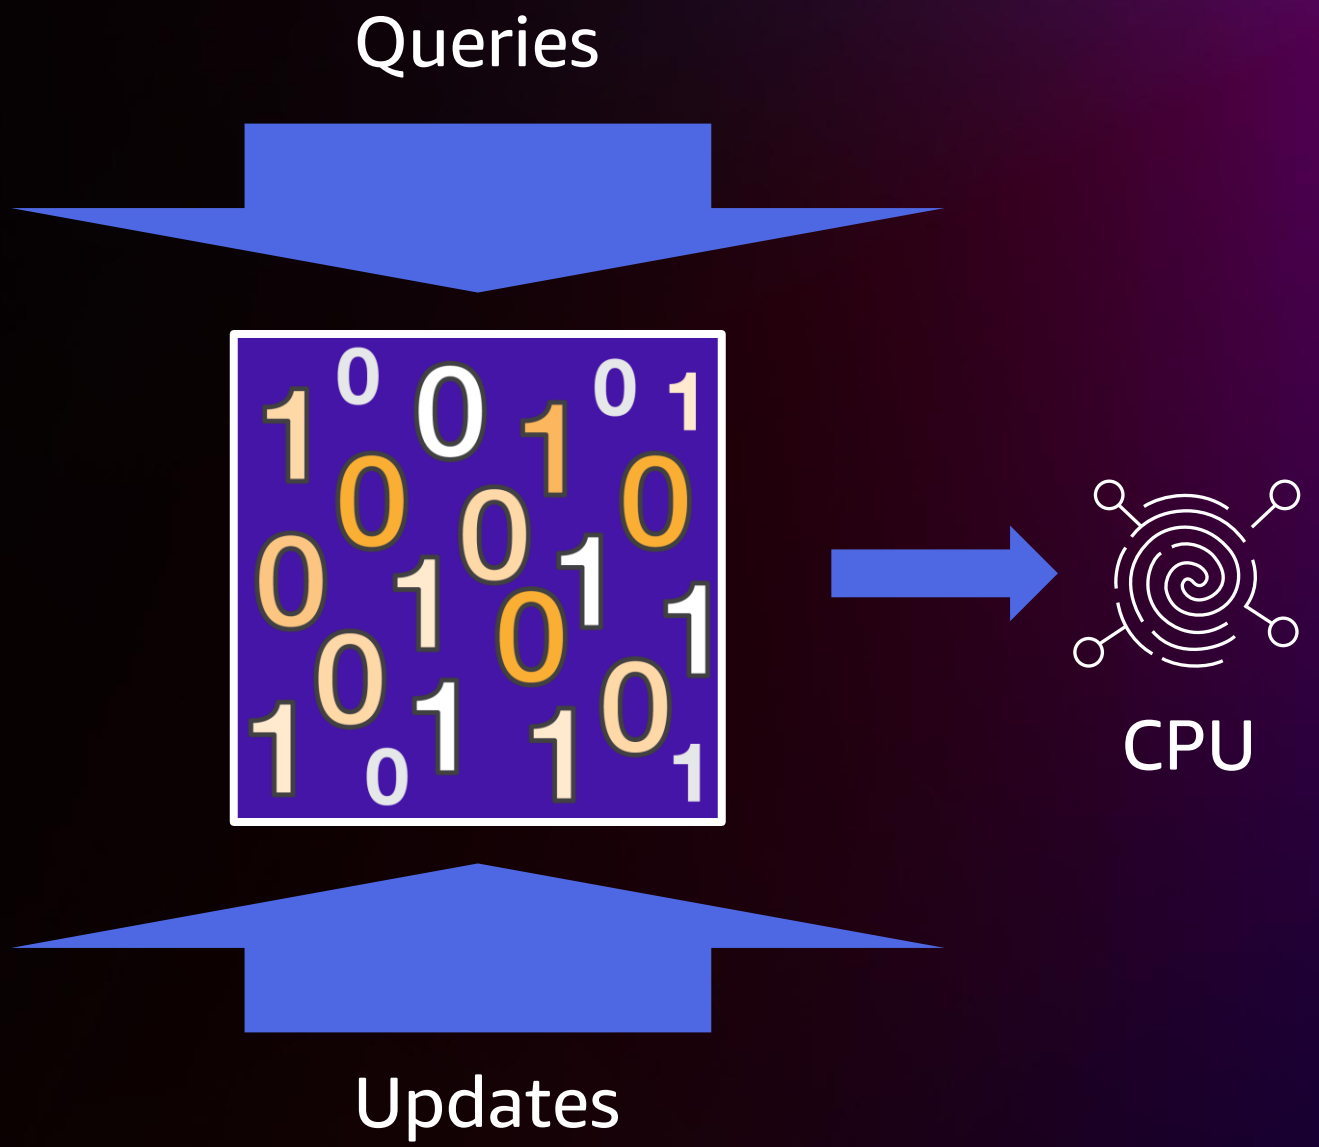
OpenSearch distributes requests via the nodes to the shards

Documents are JSON entities, all keys indexed



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Shards are the workers



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Sizing/Costing Examples

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Example sizing problems

Q: I have 3 TiB of logs that I'm ingesting daily. I want to keep them for 3 months.

- What infrastructure should I use?
- What will it cost?
- How do I optimize that cost?

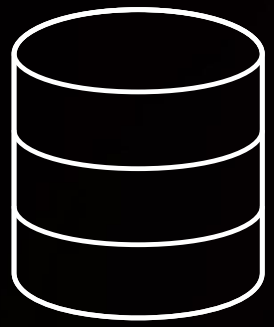
Q: I have a 5GiB product catalog with 1.5 million items. I have 4,000 concurrent site visitors, running 100 queries per second.

- What infrastructure should I use?
- What will it cost?
- How do I optimize that cost?

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Search: 5 GiB, 1.5M items, 4,000 concurrent visitors, 100 qps





Storage – 5 GB

Any instance type will suffice

Compute – 1 shard, 2+ replicas

Min CPUs? Test!

Recommendation: 3 R6gd.2xlarge hot data nodes, 3 C6g.large cluster managers

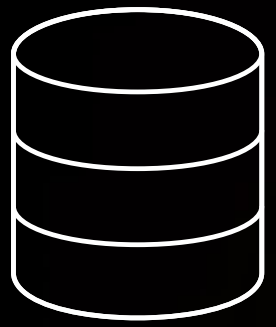
Monthly cost (on demand, us-east-1): $1,922.82 monthly

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Logs: 3 TiB/day for 90 days

All-hot



Storage – 760 TiB

Min 32 r6g.12xlarge
with 12 TiB GP3

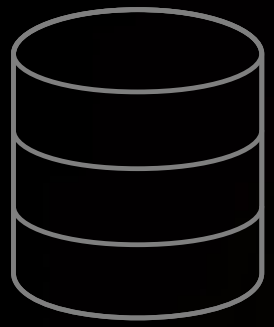
Recommendation: 63xR6g.4xlarge data nodes, 13x C6g.4xlarge cluster manager, 12 TiB EBS, 3-Zone replication

Cost (on demand, us-east-1): $155,329 monthly

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

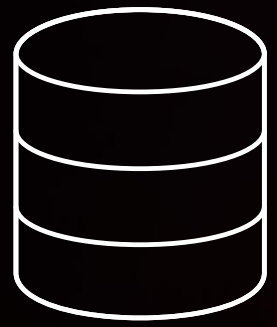
Logs: 3 TiB/day for 90 days

All-hot



Cost $155,329

14 Hot, 76 UltraWarm





Storage – 118/233 TiB

Recommendation: 12 R6g.4xlarge data nodes, 3 C6g.xlarge cluster manager, 12 UltraWarm1.large

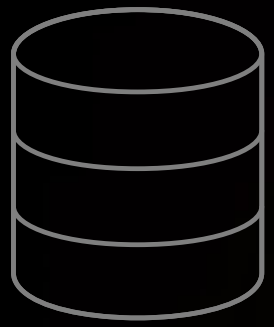
Cost (on demand, us-east-1): $55,733 monthly (65% savings vs. all hot)

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Logs: 3 TiB/day for 90 days

All-hot



Cost $160,480

14 Hot, 76 UltraWarm



Cost $55,733

65% saving

14 hot, 16 UW, 60 Cold



Storage – 118/49/184

76% saving

Recommendation: 12 R6g.4xlarge data nodes, 3 C6g.xlarge cluster manager, 3 UltraWarm1.large

Cost (on demand, us-east-1): $38,126 monthly (76% saving vs. all-hot)

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

To scale correctly, iterate

Set config

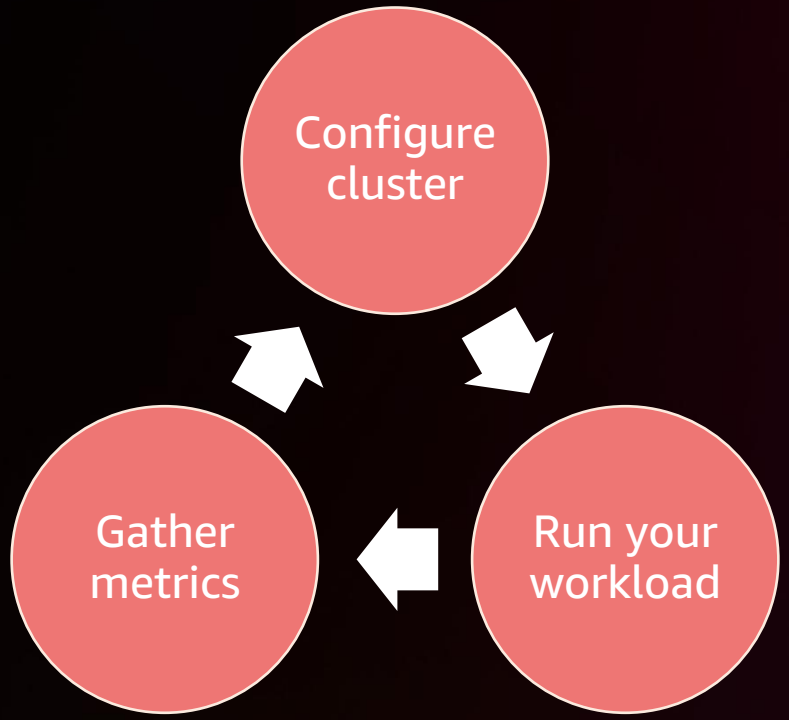
Instance type and storage

Monitor

CloudWatch, OpenSearch APIs for CPU, JVM, storage, queues

Adjust

Change instance types, storage amounts



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Key sizing parameters

Storage

Index data (primary and replica shards) is stored on disk or in Amazon S3

Instance count

Add instances for increased parallelism

Instance type

Deploy instances based on storage and compute needs

Shard count

Shards are the units of work and storage

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Capacity planning

Gather information

- Data per day
- Number of indexes
- Concurrency
- RPS (Docs, Qs)
- Index/query complexity
- Durability need
- HA need

Plan for storage

- Ephemeral/EBS
- Replication
- Retention
- UltraWarm
- Cold storage

Start with storage

Plan for compute

- Concurrency – number of readers/writers
- Active indexes/shards
- Total request counts and latencies

Adjust for compute
CPUs = 1.5 * active shards

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Sizing: Storage

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Storage needed

Storage needed = source * 1.1 * 2 * 7 * 1.15

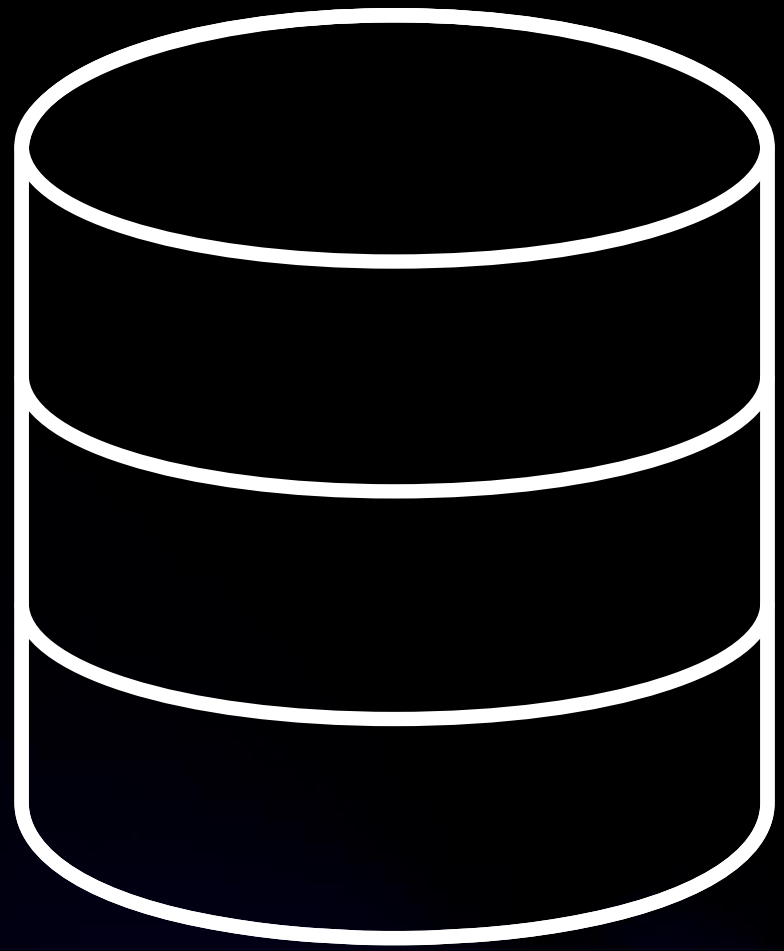
On disk, an index is ~1.1 x source

Each replica requires that much additional storage

For streaming data, multiply by retention (ignore for single indexes)

Add 15% for overhead

Deploy instances based on storage supplied



aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Storage needed – warm/cold

Storage needed = source * 1.1 * retention * 1.15

Additional stored replicas are not needed since the data is stored durably in S3

Amazon OpenSearch Service uses metadata stubs

```
GET _cat/indices?v
```

health	status	index	uuid	pri	rep	docs.count	docs.deleted	store.size	pri.store.size
green	open	.opendistro_security	6T8LBUqMQiqcdOdw2QSRQQ	1	1	9	3	115.6kb	57.8kb
green	open	domains-2022.09.12	5TUZNIiIRKeLN_8DZbuQCA	1	1	1743471	0	201.1mb	100.5mb
green	open	domains-2022.10.01	7rSAjDFxQBOsGy2l1QHgfA	1	1	1739062	0	199.4mb	99.7mb

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Scale shards by storage too

Set shard count so that primaries hold max 50 GB per shard (logs) or 30 GB per shard (search)



1 TB index
20-25 primary shards

Template or the _rollover API (or via ISM)

```
PUT _template/template1
{
  "index_patterns": ["logs-*"],
  "settings": {
    "number_of_shards": 1,
    "number_of_replicas": 1
  }
}
```

```
POST test000001/_rollover
{
  "conditions": {
    "max_size": "5gb",
    "max_age": "7d",
    "max_docs": 1000000
  }
}
```

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why are smaller shards better for search?

MATCH COUNT DRIVES QUERY PERFORMANCE, INDEXING PERFORMANCE IS MORE INDEPENDENT

Search

- Queries match many documents
- Per-match scoring and sorting are compute- and time-intensive
- Query rate is high

Logs

- Filtered by time range, usually small, fewer matches
- Usually aggregated, not scored or sorted
- Query rate is low

Smaller shards mean more parallelism

Larger shards work well

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Revisiting 3 TiB/day – 14 hot, 16 UW, 60 Cold

Sizing for storage only!

Hot storage needed = 3 TiB * 1.25 * 2 * 14 = 105 TiB

Storage limit R6g.4xlarge = 12 TiB; 105/12 = 9 r6g.4xlarge minimum

Warm storage needed = 3 TiB * 1.25 * 16 = 60 TiB

Storage limit ultrawarm1.large = 20 Tib; 60 / 20 = 3 ultrawarm1.large minimum

Cold storage needed = 3 TiB * 1.25 * 60 = 225 TiB

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Sizing: Compute

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Compute needed

vCPUs = 1.5 * active shards

assumes each shard occupies a vCPU 100%



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

“Active” index – an index processing requests

Streaming:
One index per day

- logs_11.28.2022
- logs_11.27.2022
- logs_11.26.2022
- logs_11.25.2022
- logs_11.24.2022
- logs_11.23.2022
- logs_11.22.2022

Search:
One index

- product_catalog

Whether you have a streaming use case or a search use case, you’ll most often have one index handling most traffic

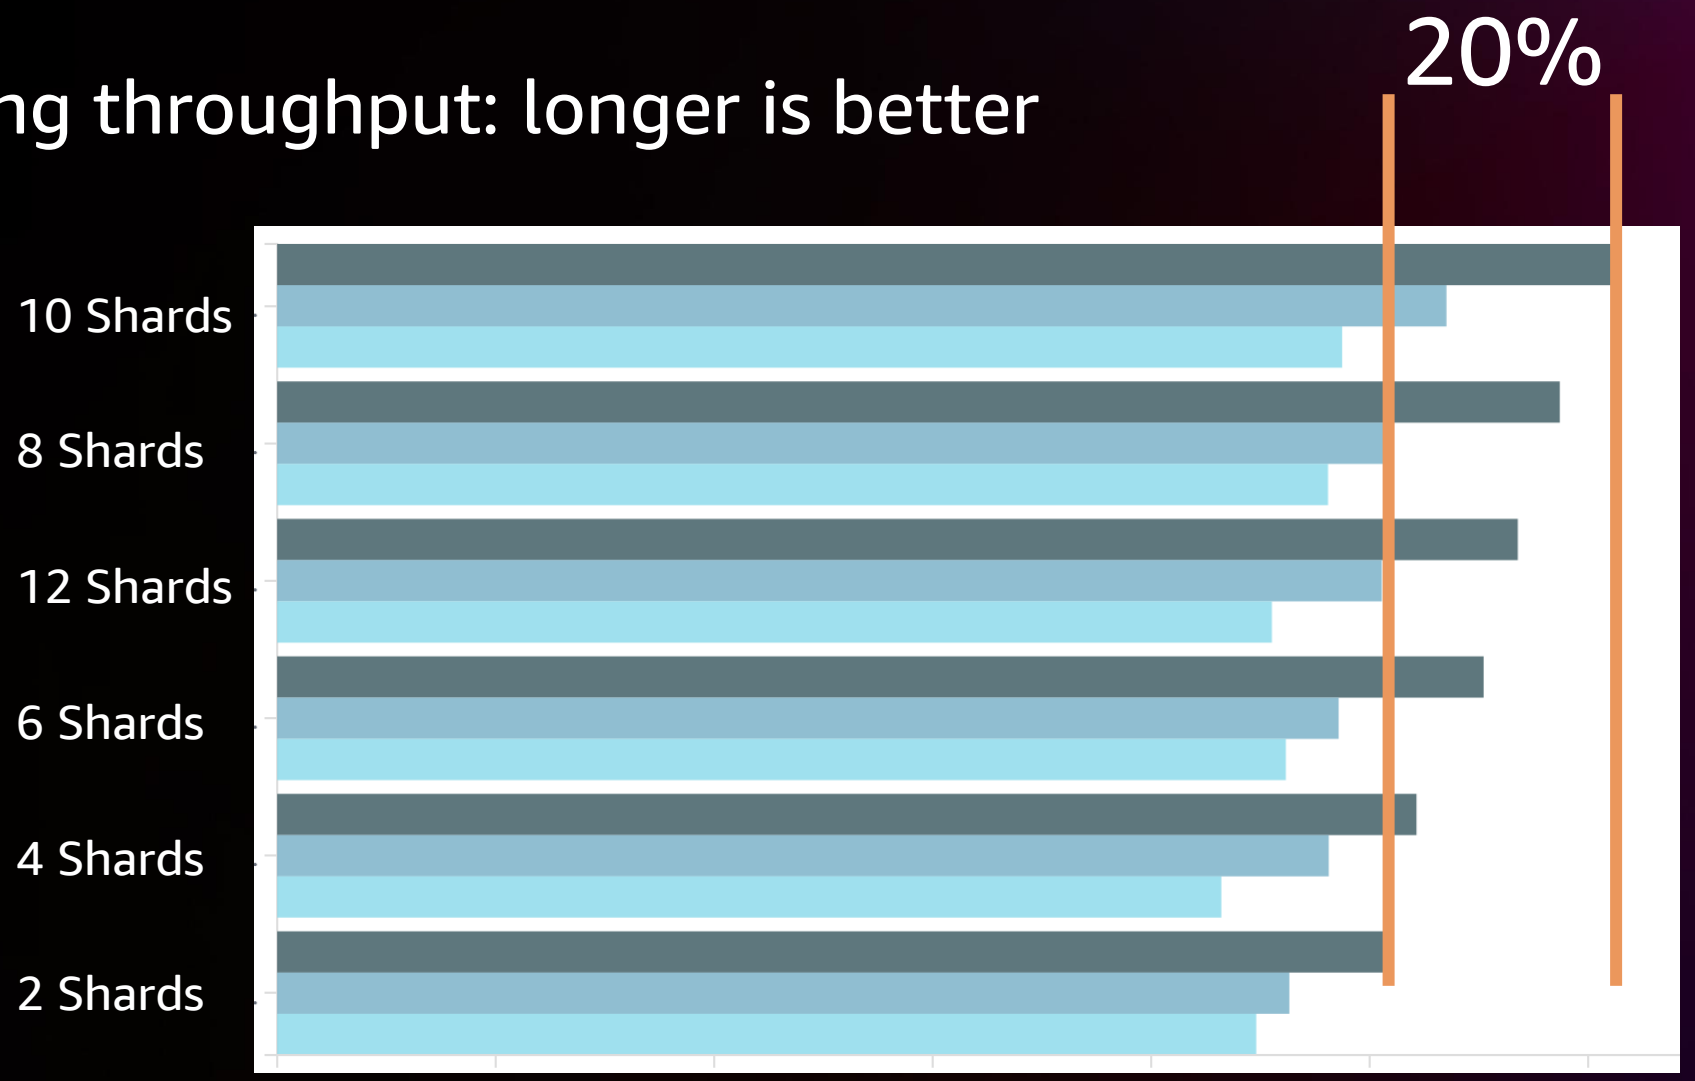
That’s the active index

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Find your “sweet spot” – Indexing

R6G.2XLARGE (8 VCPUS), LOGS DATA SET

Measuring throughput: longer is better



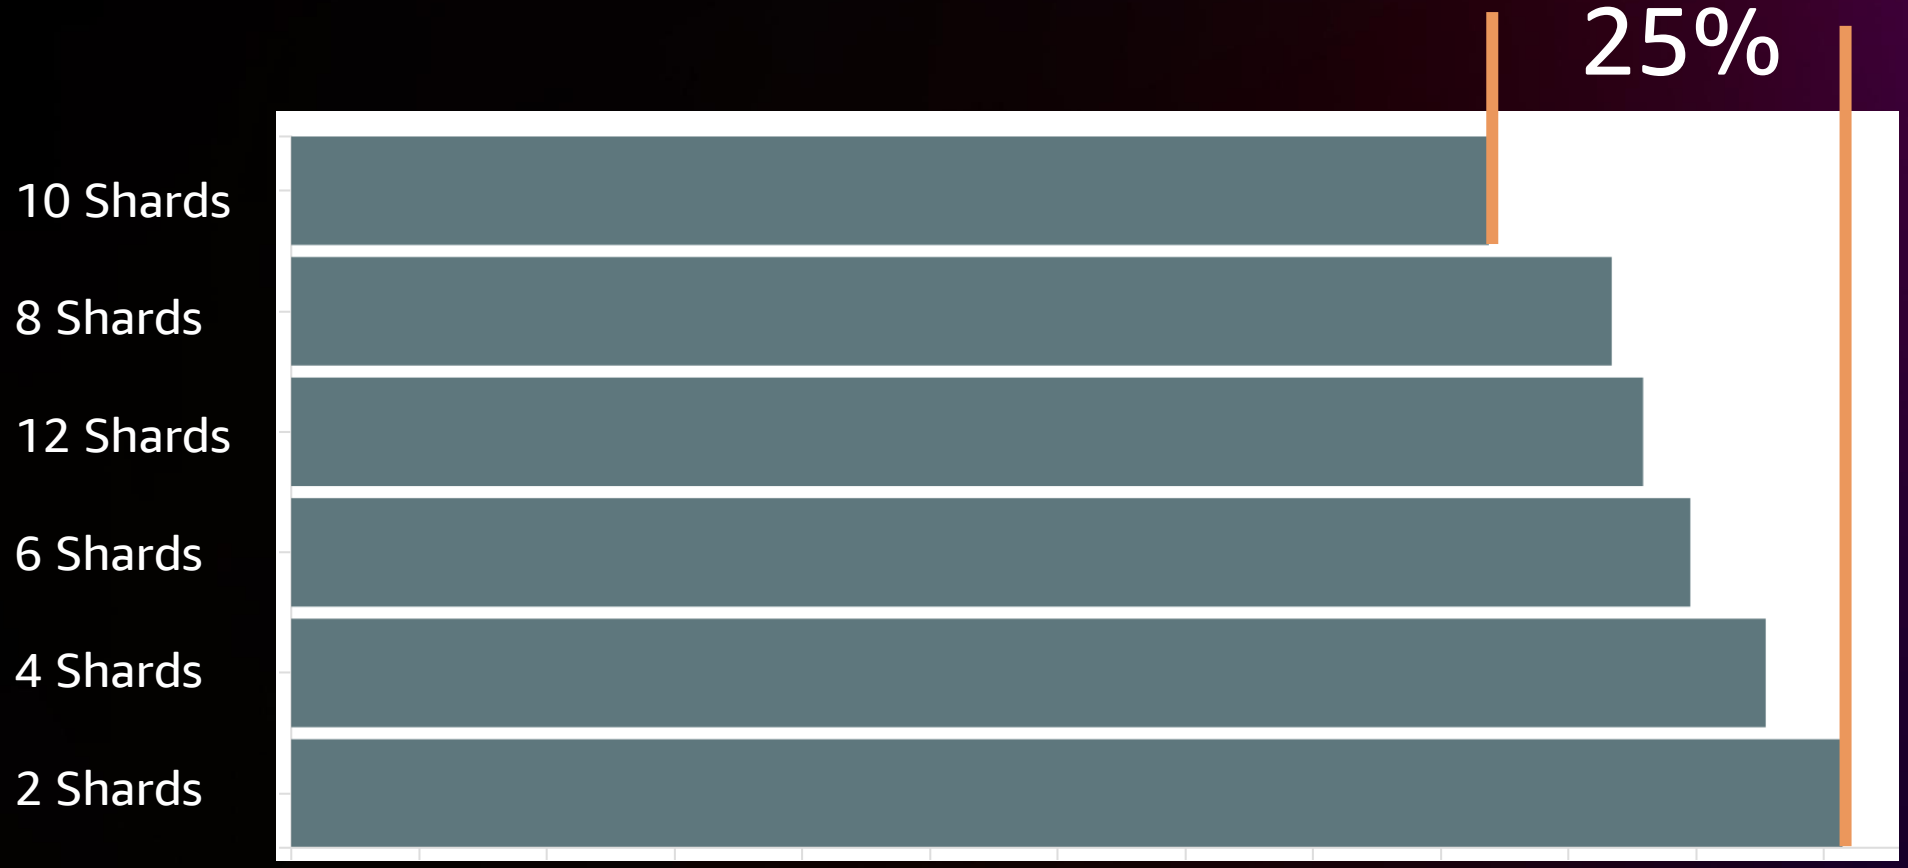
aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Find your “sweet spot” – Query

R6GD.2XLARGE (8VCPUS), SEARCH DATA SET

Measuring total test duration: shorter is better



aws

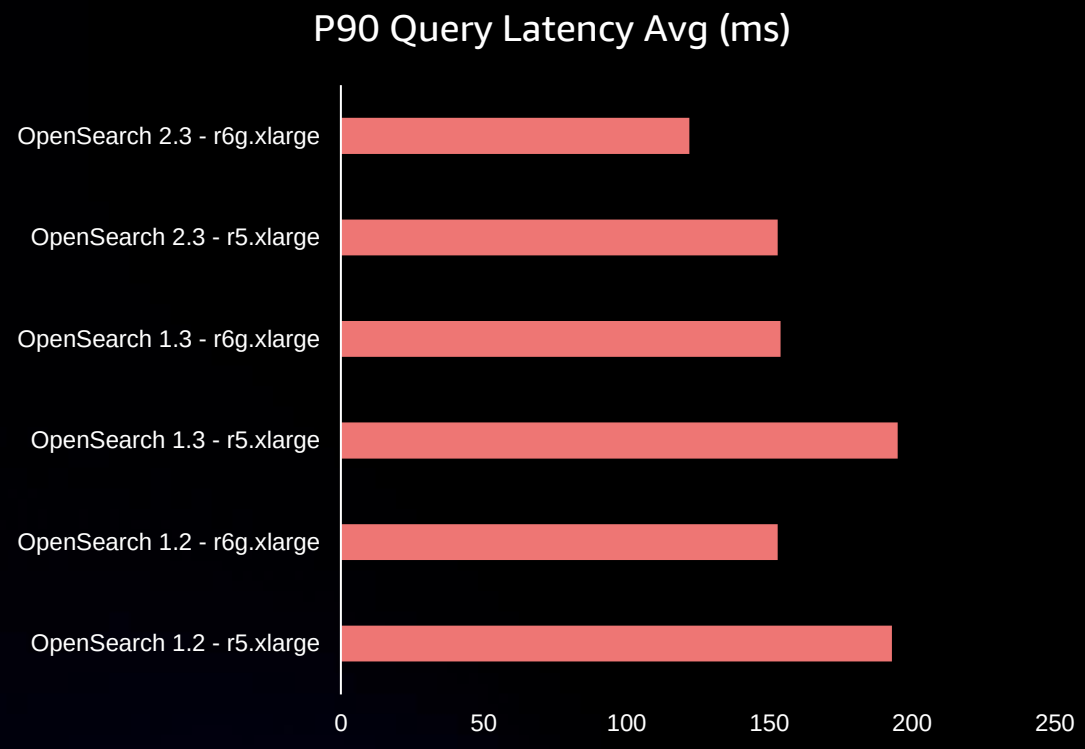
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Version and AWS Graviton2

Newer versions continue to offer performance improvements

Version 2.3 of OpenSearch

- 20% performance improvement over versions 1.3 and 1.2



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

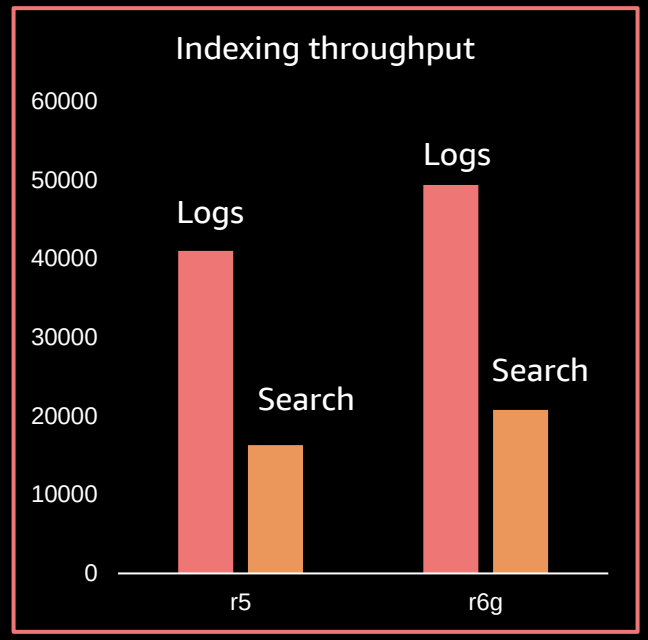
Use AWS Graviton2 instances

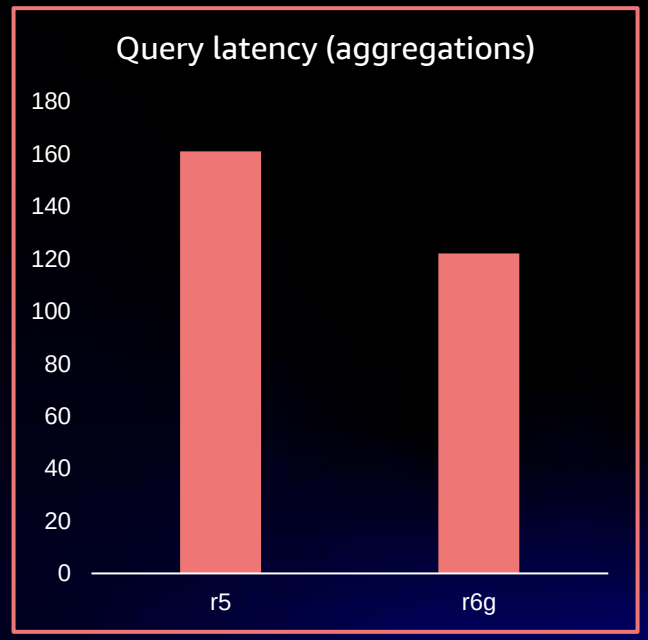
10% cost reduction (r6g vs. r5)

+20% indexing throughput (logs)

+27% indexing throughput (search)

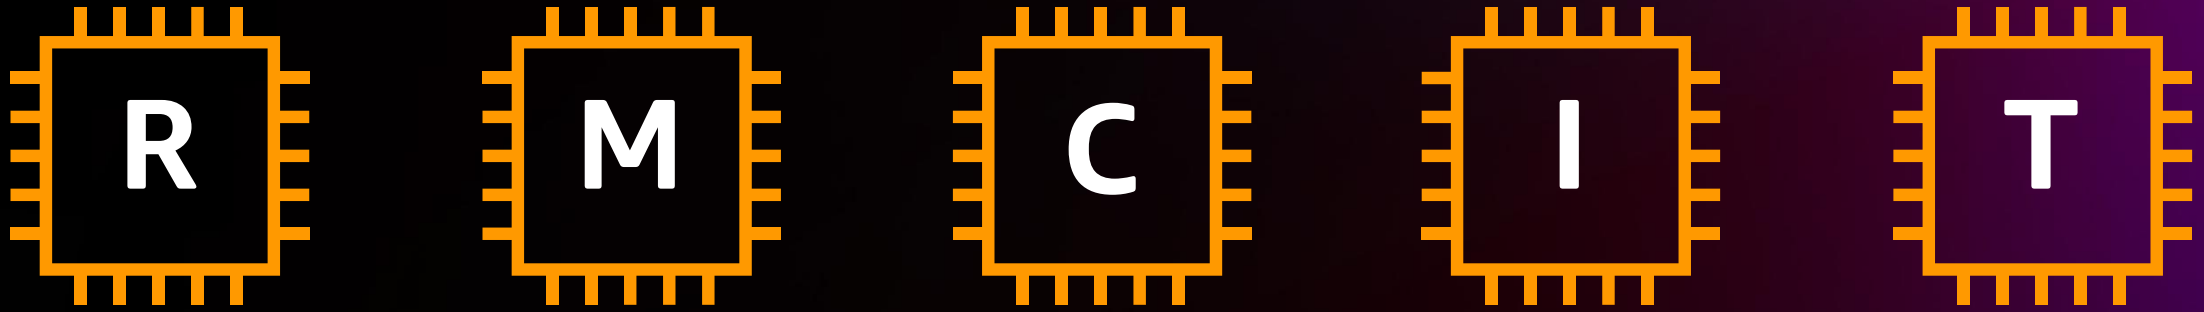
-24% query latency

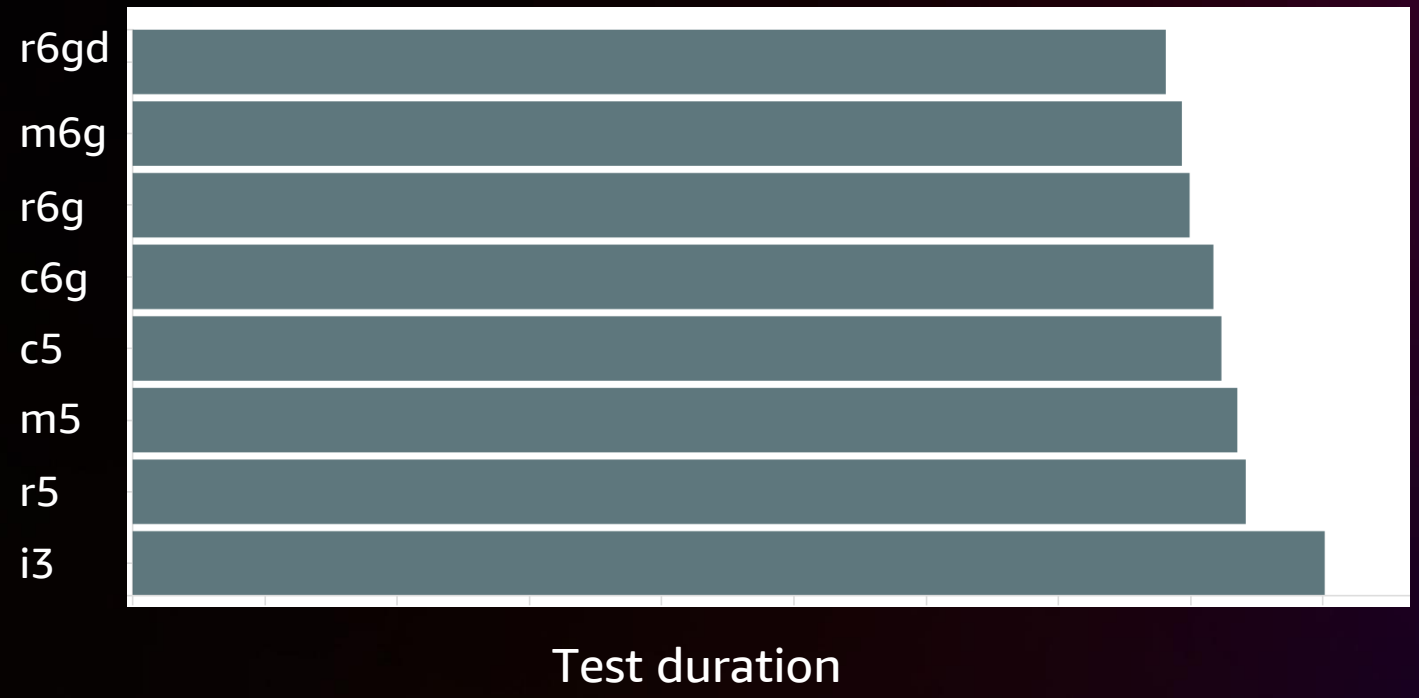




© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Instance classes



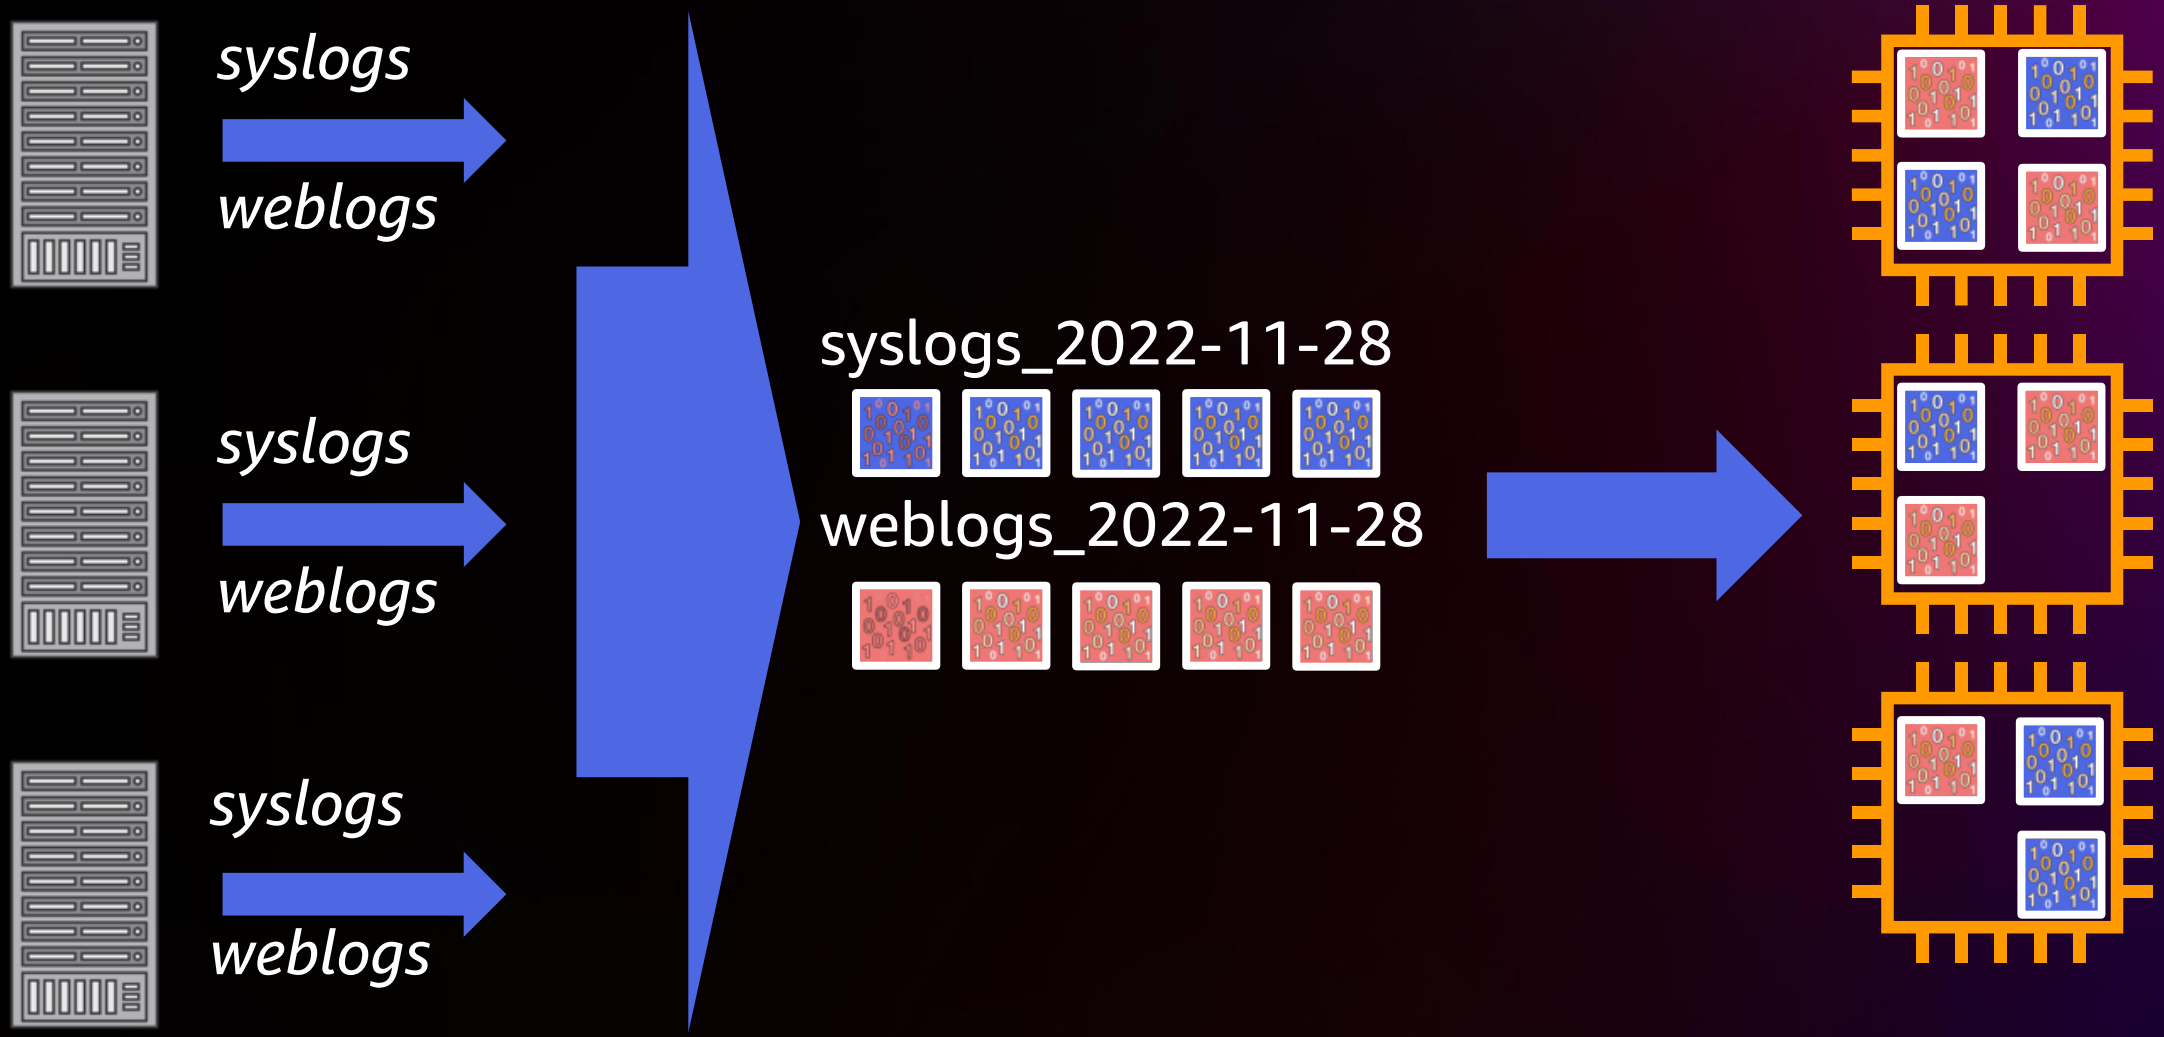


aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

For log analytics, sizing for storage usually provides enough compute. But, the number of active data streams can push your resource needs higher.

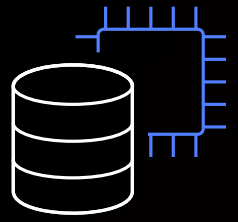
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Multiple data streams overlap on data nodes



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Per-node resources



Storage

On disk, an index with 1 replica needs ~2.2x the source size; use this as an initial scale point



CPU

Each shard receiving a request uses a CPU to process that request; use 1.5 vCPU per active shard



RAM

The Java heap is 50% of total RAM; off-heap is file system cache; scale horizontally for heap pressure



Queues

OpenSearch queues are for surges; monitor queue depth, scale horizontally

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Java heap

Java heap is 50% of RAM up to ~31.5 GB

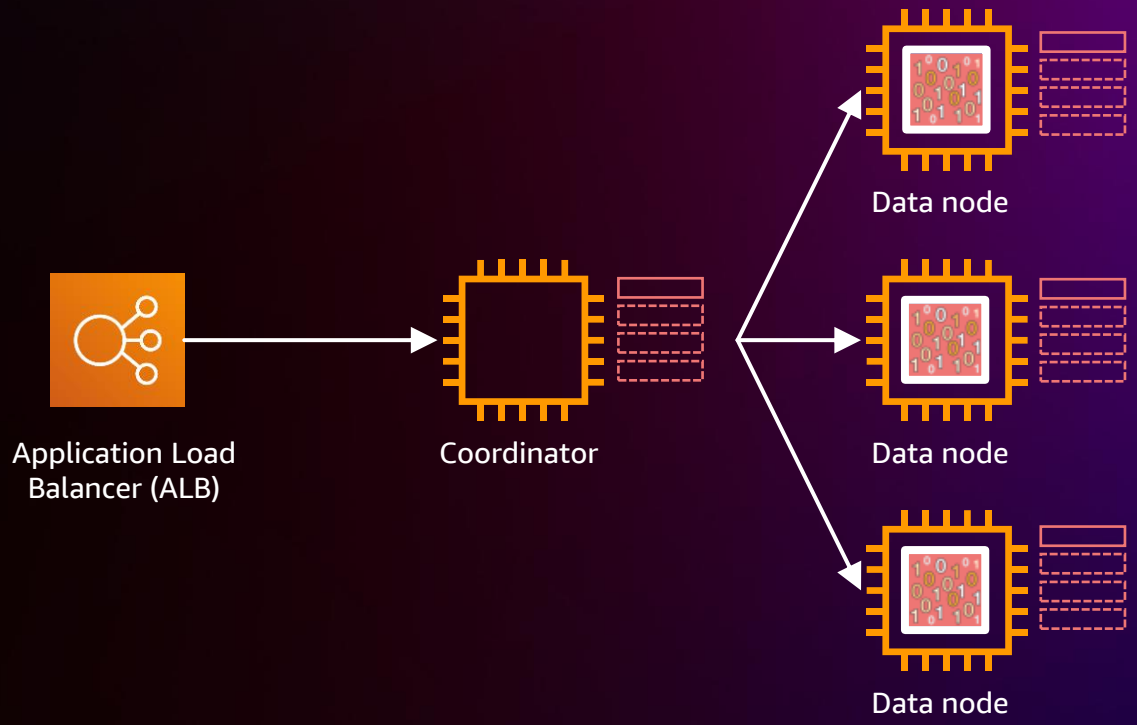
Max heap:
R series @ 2xlarge; M series @ 4xlarge;
C series @ 8xlarge; I series @ 2xlarge

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

OpenSearch Queues

Search queues, all instance types
1000 search slots
200 bulk slots

Full queues return 429 status code



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Up vs. out

Up

Pro

More processors

Larger disks

Less node-node chatter

Con

Java heap capped at 32 GB

Queues capped at 200/1000

Disk and network bandwidth capped

Out

Pro

More processors

More Java heap

More queue space

Con

Larger surface area for instance failure

More node-node chatter

Fan-in

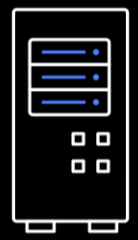
aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cost Optimization - Infrastructure

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cost model



Instance hours

You choose the instances to deploy for the data nodes in your cluster

OpenSearch Service charges per hour for these instances

Largest cost component



Amazon EBS

If you choose Amazon EBS-backed instances, OpenSearch Service charges per GB of data deployed



Amazon S3

When you use UltraWarm or Cold Storage, OpenSearch Service charges for actual usage of Amazon S3

Drives instance count



Data transfer

OpenSearch Service charges for data transfer at Amazon EC2 data transfer rates

aws

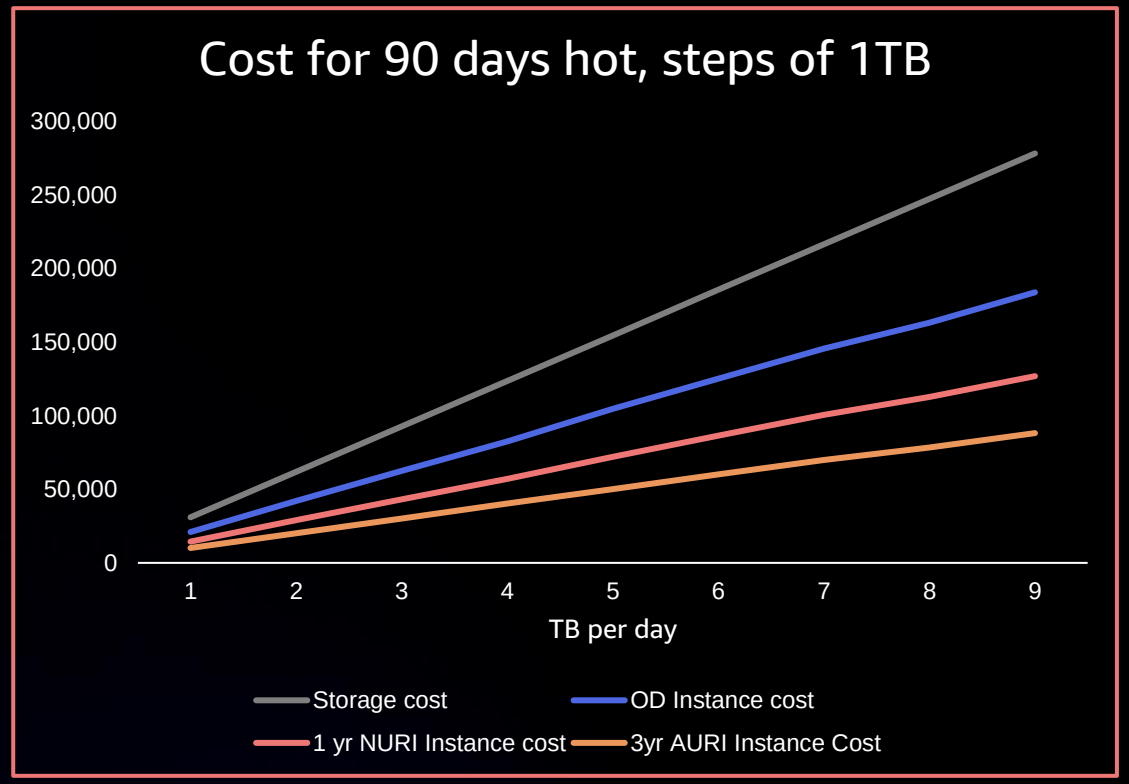
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Use reserved instances

Amazon OpenSearch Service supports reserved instances (RIs) at 1-yr and 3-yr commitment, with all- or no-up-front payment options

Instance savings vs. on demand

- ~30% for 1-year, no-up-front Ris
- ~50% for 3-year, all-up-front RIs



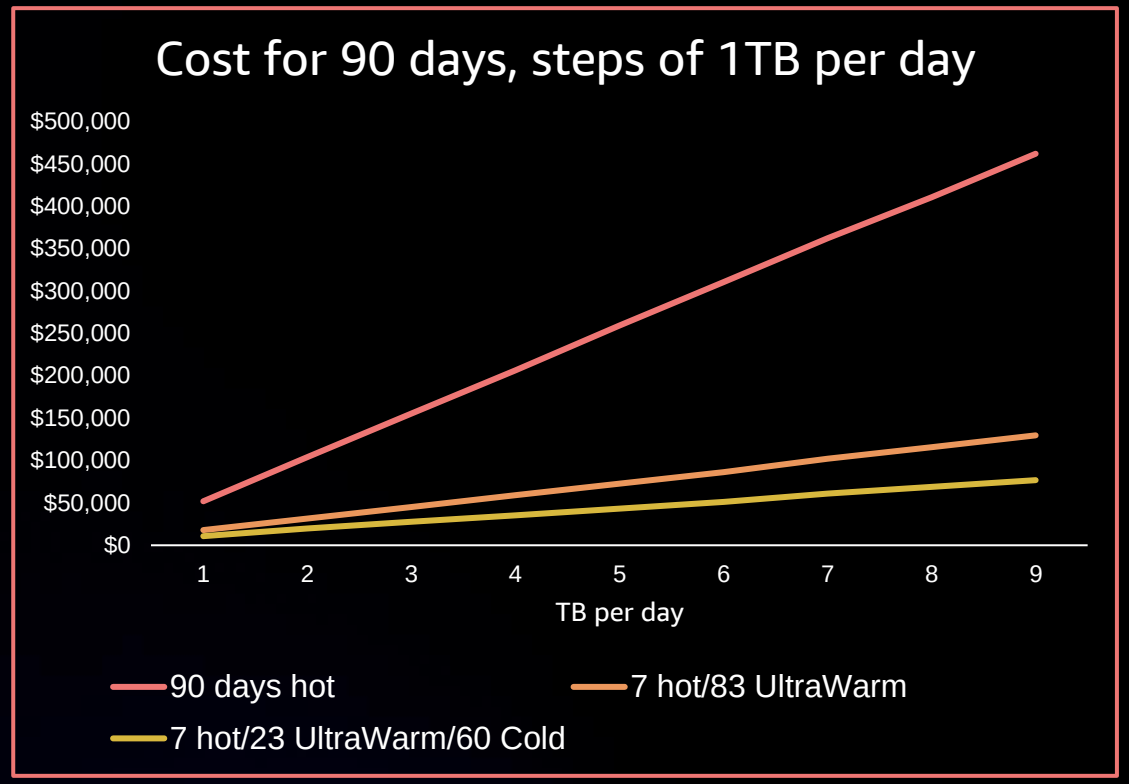
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Use UltraWarm and Cold Storage for logs

UltraWarm trades storage cost for performance

Data in UltraWarm is immutable

Designed for interactive access to small slices of log data



72% and 83% total cost reduction vs. all-hot

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Use UltraWarm and Cold Storage for logs

UltraWarm trades storage cost for performance

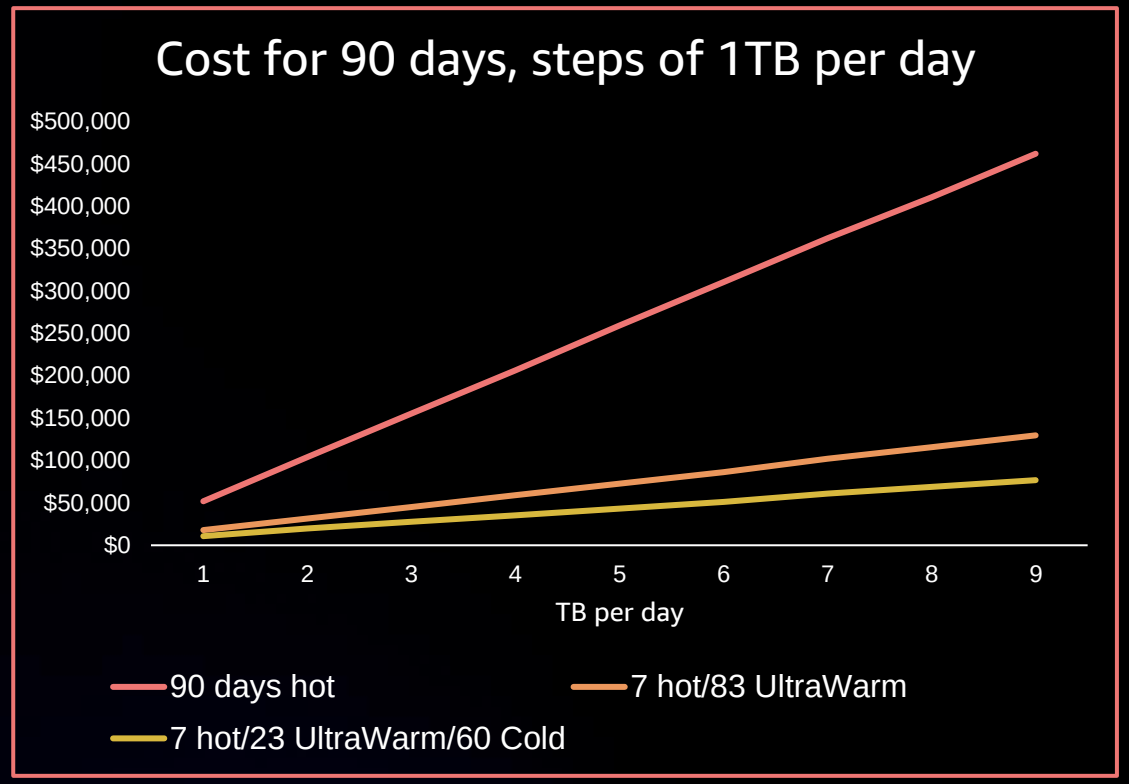
Data in UltraWarm is immutable

Designed for interactive access to small slices of log data

Cold Storage allows you to detach/reattach indices to UltraWarm

Reattach in seconds

Near-S3 cost for storage enables years of retention



72% and 83% total cost reduction vs. all-hot

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

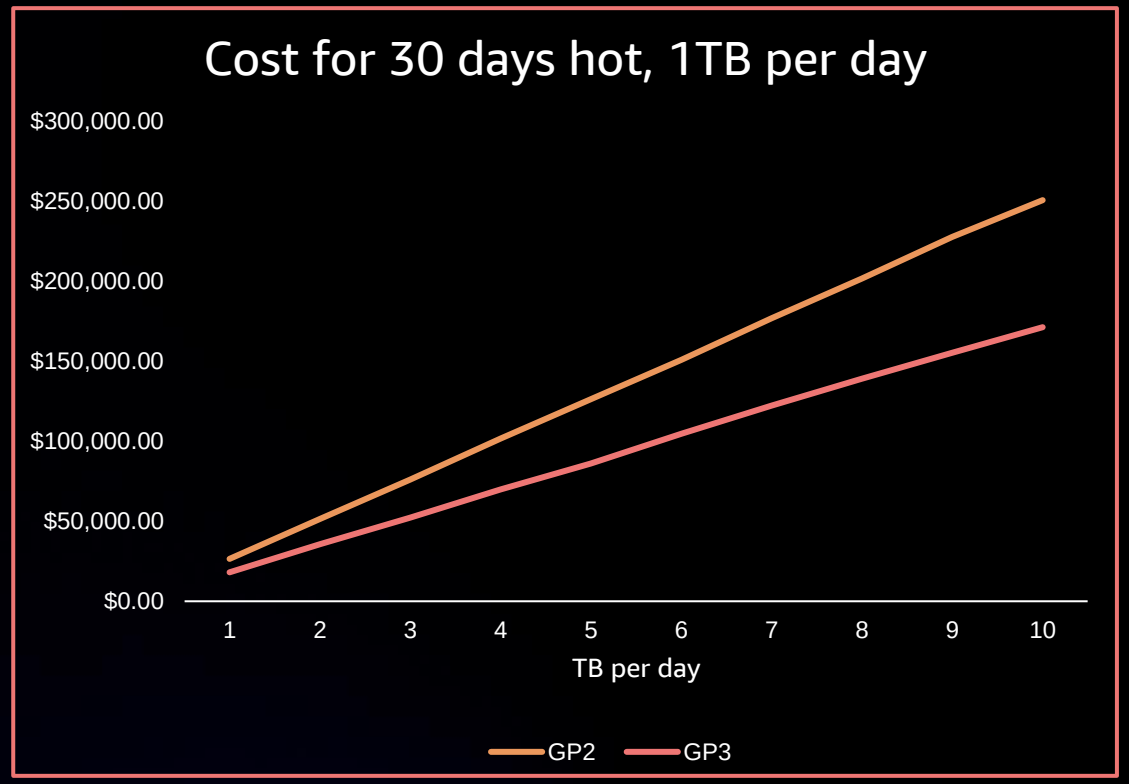
Use Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (GP3) volumes

2x denser storage on R5, R6g, M5, and M6g instance families

Available at a 9.6% lower price point per GB than GP2 volumes

Provision IOPS independent of storage capacity

Migrate with a configuration change

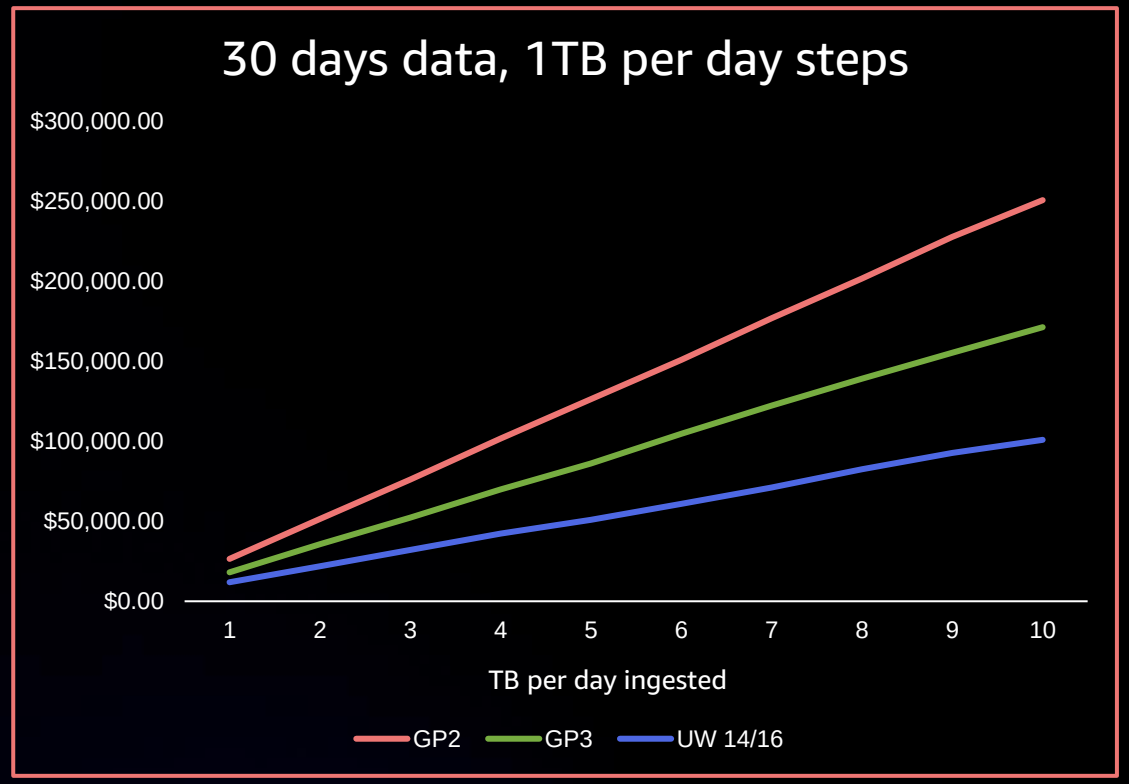


© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Combine storage options for best price/performance

Total cost for the domain

Comparing all-hot with 14 hot, 16 warm for UltraWarm



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cost Optimization: Workload

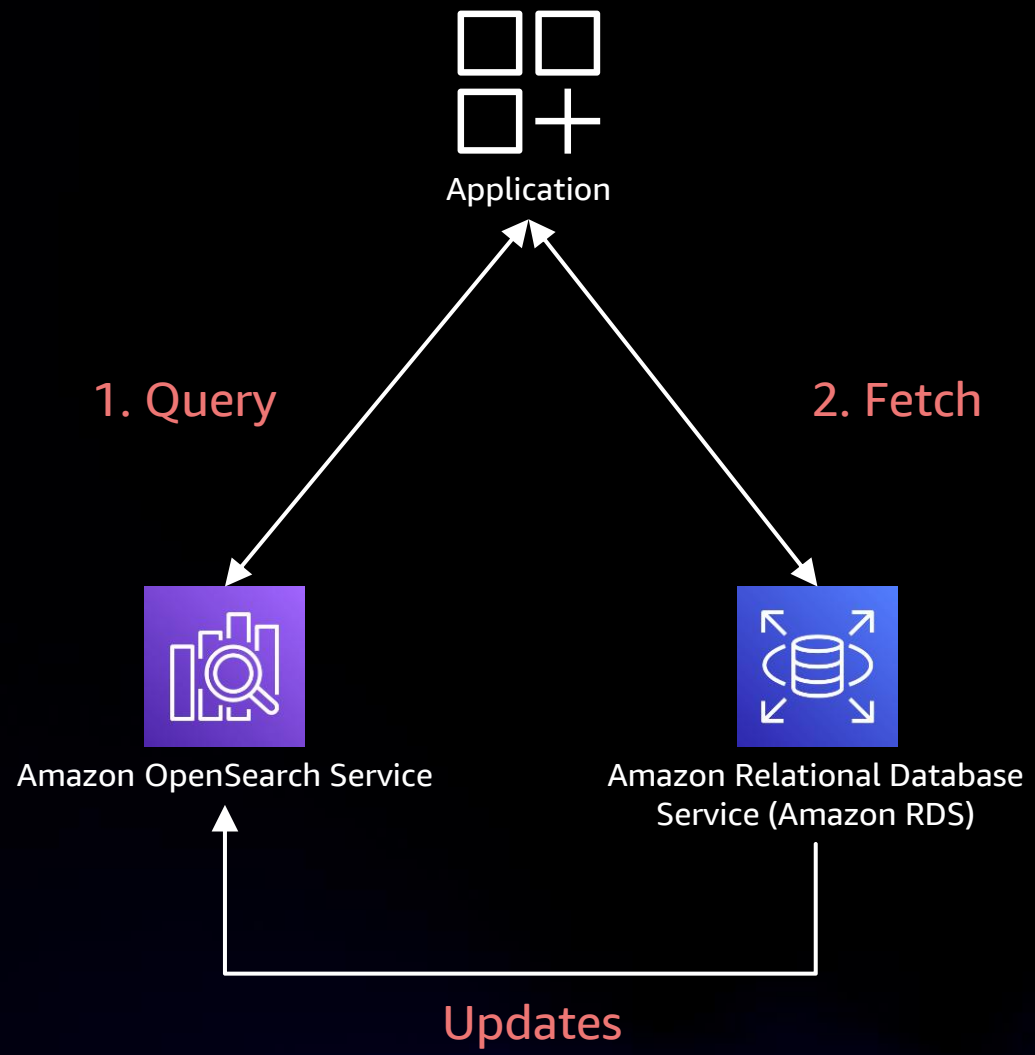
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Send only what you plan to search

Whether you're sending search or log data, remove any fields that are not the target of a query before you send the docs

Best practice for search – use in tandem with a system of record (SOR)

Search with OpenSearch, retrieve full records from SOR



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

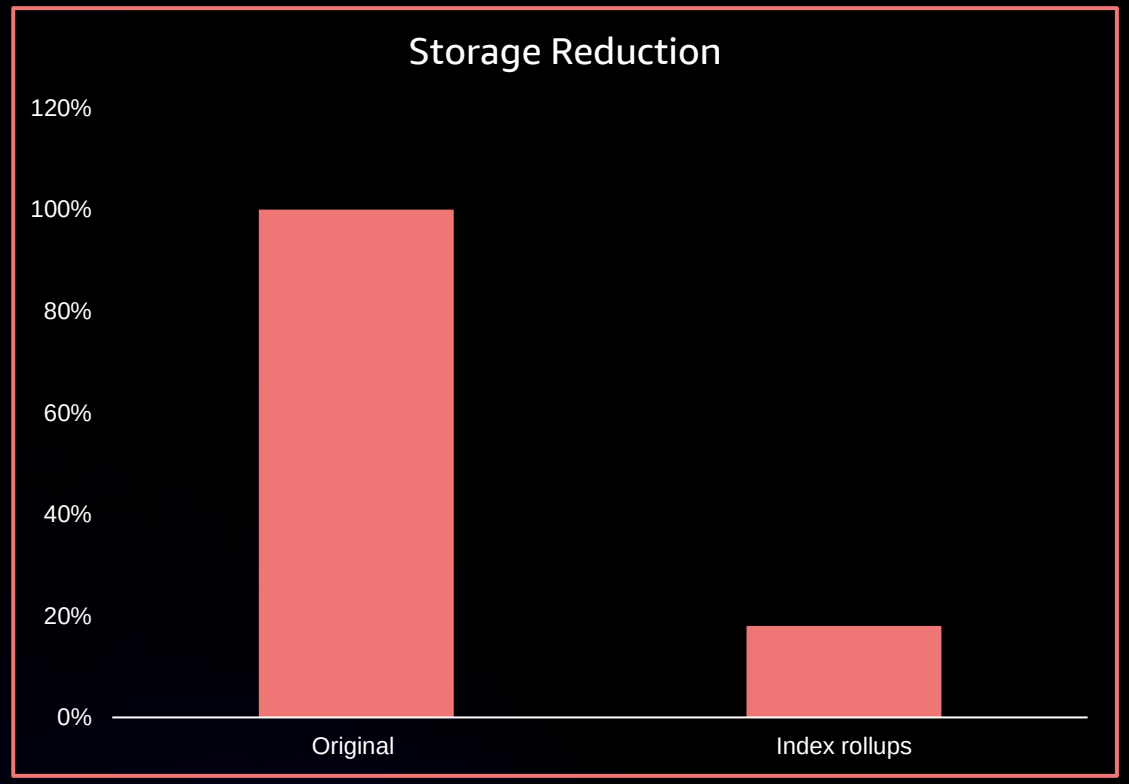
Index Rollups

Size and cost reduction based on granularity

Create a new index with fields that interest you

Original index has about 25 indices and the rollup index has 7 fields

```
1 health status index
    store.size pri.store.size
2 green  open   example_rollup
         1.7mb        876.7kb
3 green  open   opensearch_dashboards_sample_data_ecommerce
          9mb          4.3mb
```



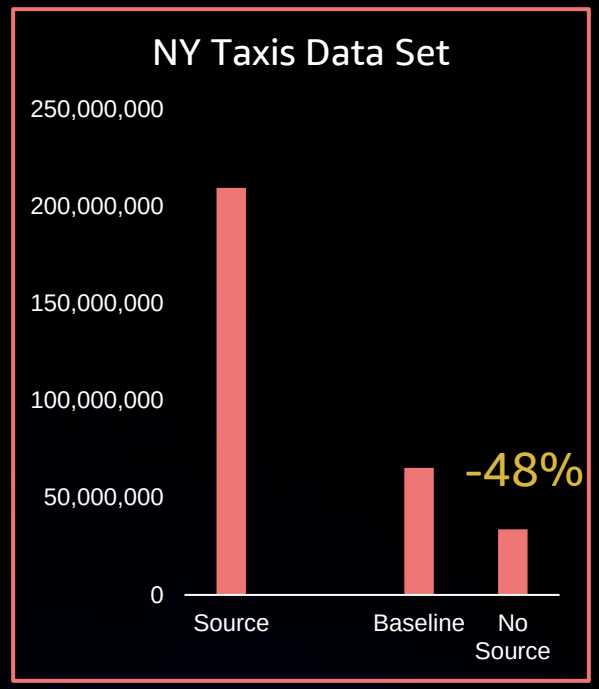
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

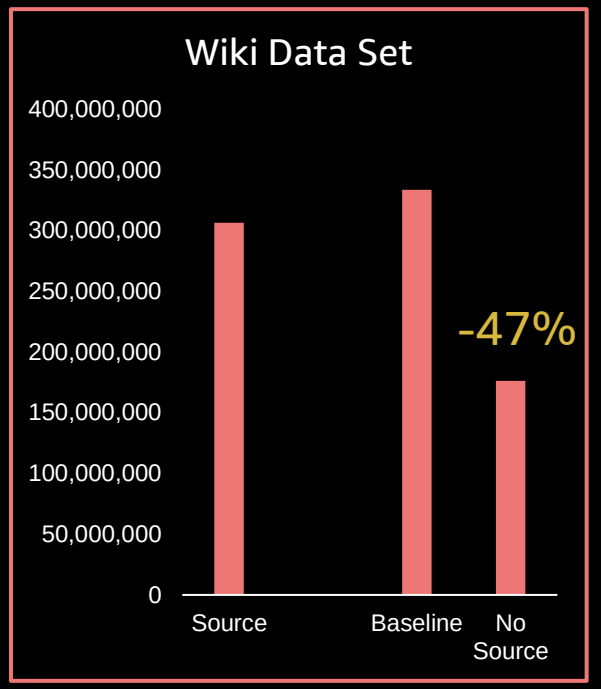
Remove _source

_source is used for updates, return values, reindexing, highlighting

Use selectively to reduce storage amounts

Fully disabling provides ~50% index size reduction





© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

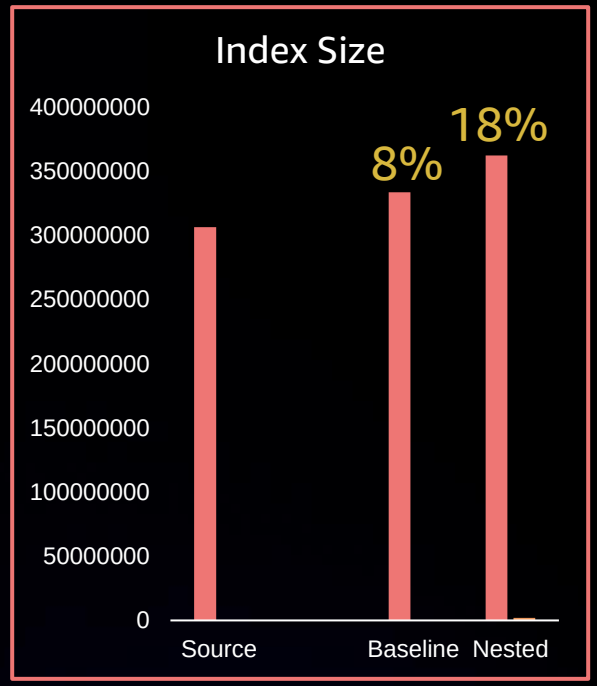
Nested vs. flat

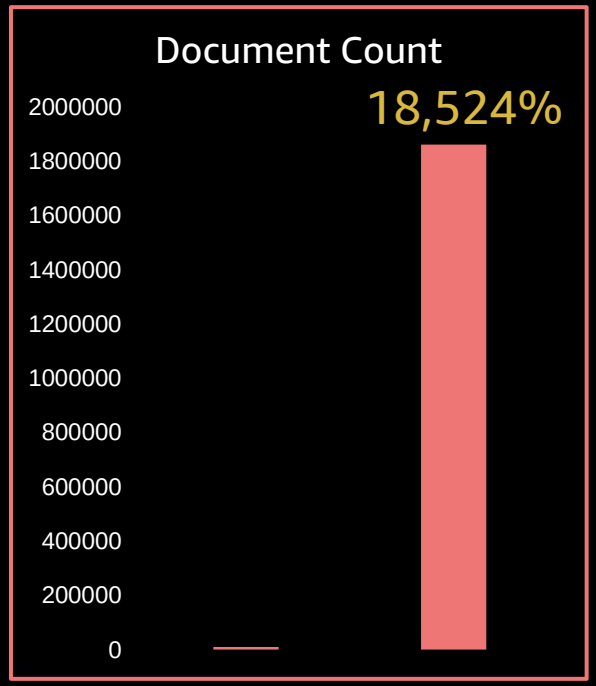
Use `nested` to prevent field cross-talk for array fields

Increases index size

Increase in doc count can increase latency

```
"links": [
  {"title": "Category:Chicago lawyers",
   "link": "Category:Chicago lawyers"
  },
  {"title": "Category:Illinois politicians",
   "link": "Category:Illinois politicians"
  }
]
```





aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

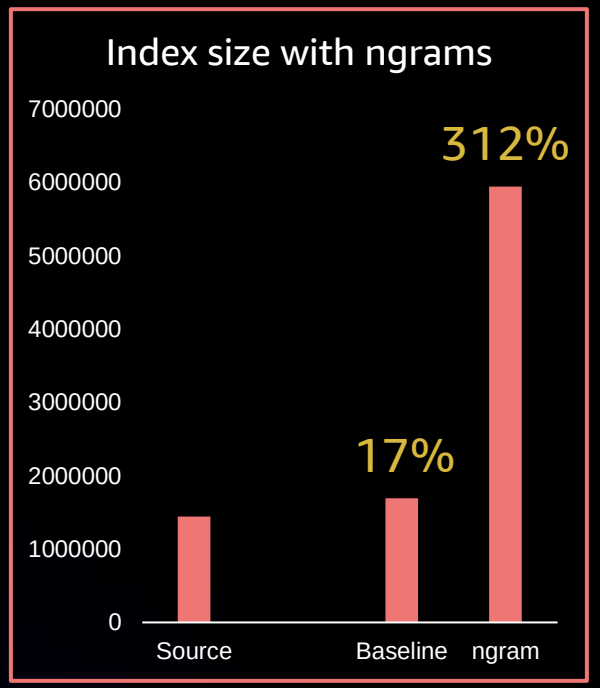
N-grams

Ngrams
recursion -> re ec cu ur rs si io on
rec ecu cur urs rsi sio ion
…

Use for: interior string matching, Asian language

Edge ngrams
pretzel key -> pr pre pret pretz pretze
pretzel ke key

Use for: Prefix matching for suggestions



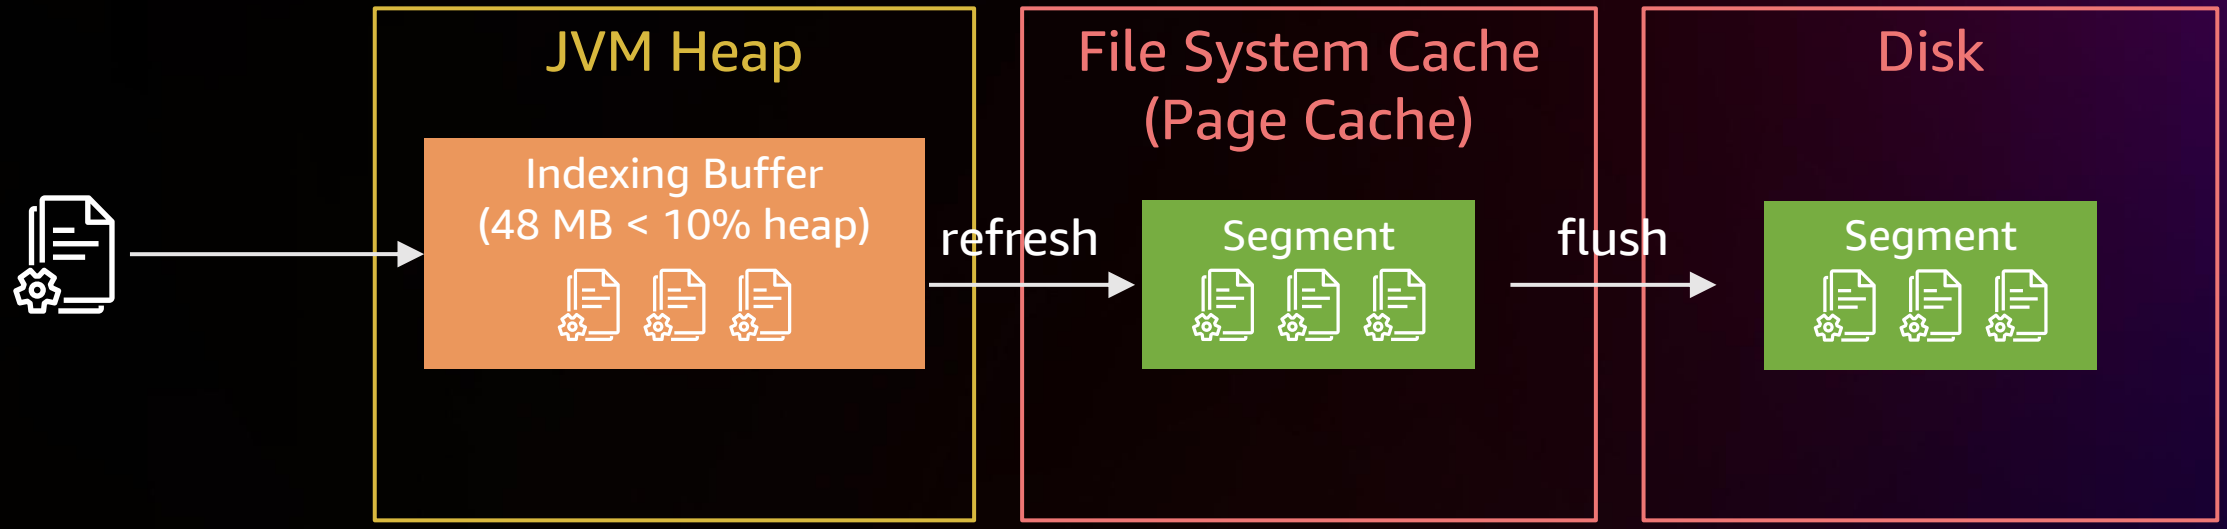
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Adjust refresh interval

Increase refresh interval (default 1s) to increase throughput

Reduce merge frequency to improve indexing performance and search performance

Trade-off with real-time requirement



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!



Please complete the session survey in the **mobile app**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.