



Elasticsearch Serverless: the Transition from Stateful to Stateless

Iraklis Psaroudakis, Principal Software Engineer
April 12th, Devvxx Greece 2025

Iraklis Psaroudakis



Principal Software Engineer
at [Elastic](#), focusing on distributed

Previously:
[Oracle Labs](#), graph-based analytics
PhD at [EPFL](#), scaling up analytics
ECE degree at [NTUA](#) in Athens

Agenda:

- What is Elastic
- What is Elasticsearch
- Shared-nothing stateful architecture
- Serverless Elasticsearch
- New stateless architecture
- Walkthrough of how data is stored in stateful vs serverless Elasticsearch
- Batching data
- Autoscaling

This talk contains personal views and is not officially endorsed

Meet Elastic — The Search AI Company

Elastic helps everyone find answers that matter.
From all data. In real time. At scale.



Founded in 2012
NYSE: ESTC



2800+
employees



40+
Countries with employees



4B+
downloads



54%
Of Fortune 500 companies
trust Elastic



Examples:  Microsoft  T-Mobile  ebay  Uber  Audi

Performance that Delivers Relevant Results in Real-time



Elasticsearch



- Distributed, scalable, highly available, resilient search & analytics engine
- HTTP based JSON interface
- Based on [Apache Lucene](#)
- Not only grep or SQL's LIKE = '%quick%'
 - Ranked results (BM25, recency, popularity), fuzzy matching
 - Complex search expressions
 - Spelling, Synonyms, Phrases, Stemming
- Timeseries, geospatial
- Vector search, (G)AI/ML, RAG search apps

github.com/elastic/elasticsearch



db-engines.com/en/ranking/search+engine

Rank			DBMS	Database Model
Mar 2025	Feb 2025	Mar 2024		
1.	1.	1.	Elasticsearch	Multi-model ⓘ
2.	2.	2.	Splunk	Search engine
3.	3.	3.	Apache Solr	Search engine, Multi-model ⓘ
4.	4.	4.	OpenSearch	Multi-model ⓘ
5.	5.	↑6.	Sphinx	Search engine
6.	6.	↑7.	Algolia	Search engine
7.	7.	↑8.	Microsoft Azure AI Search	Search engine, Multi-model ⓘ

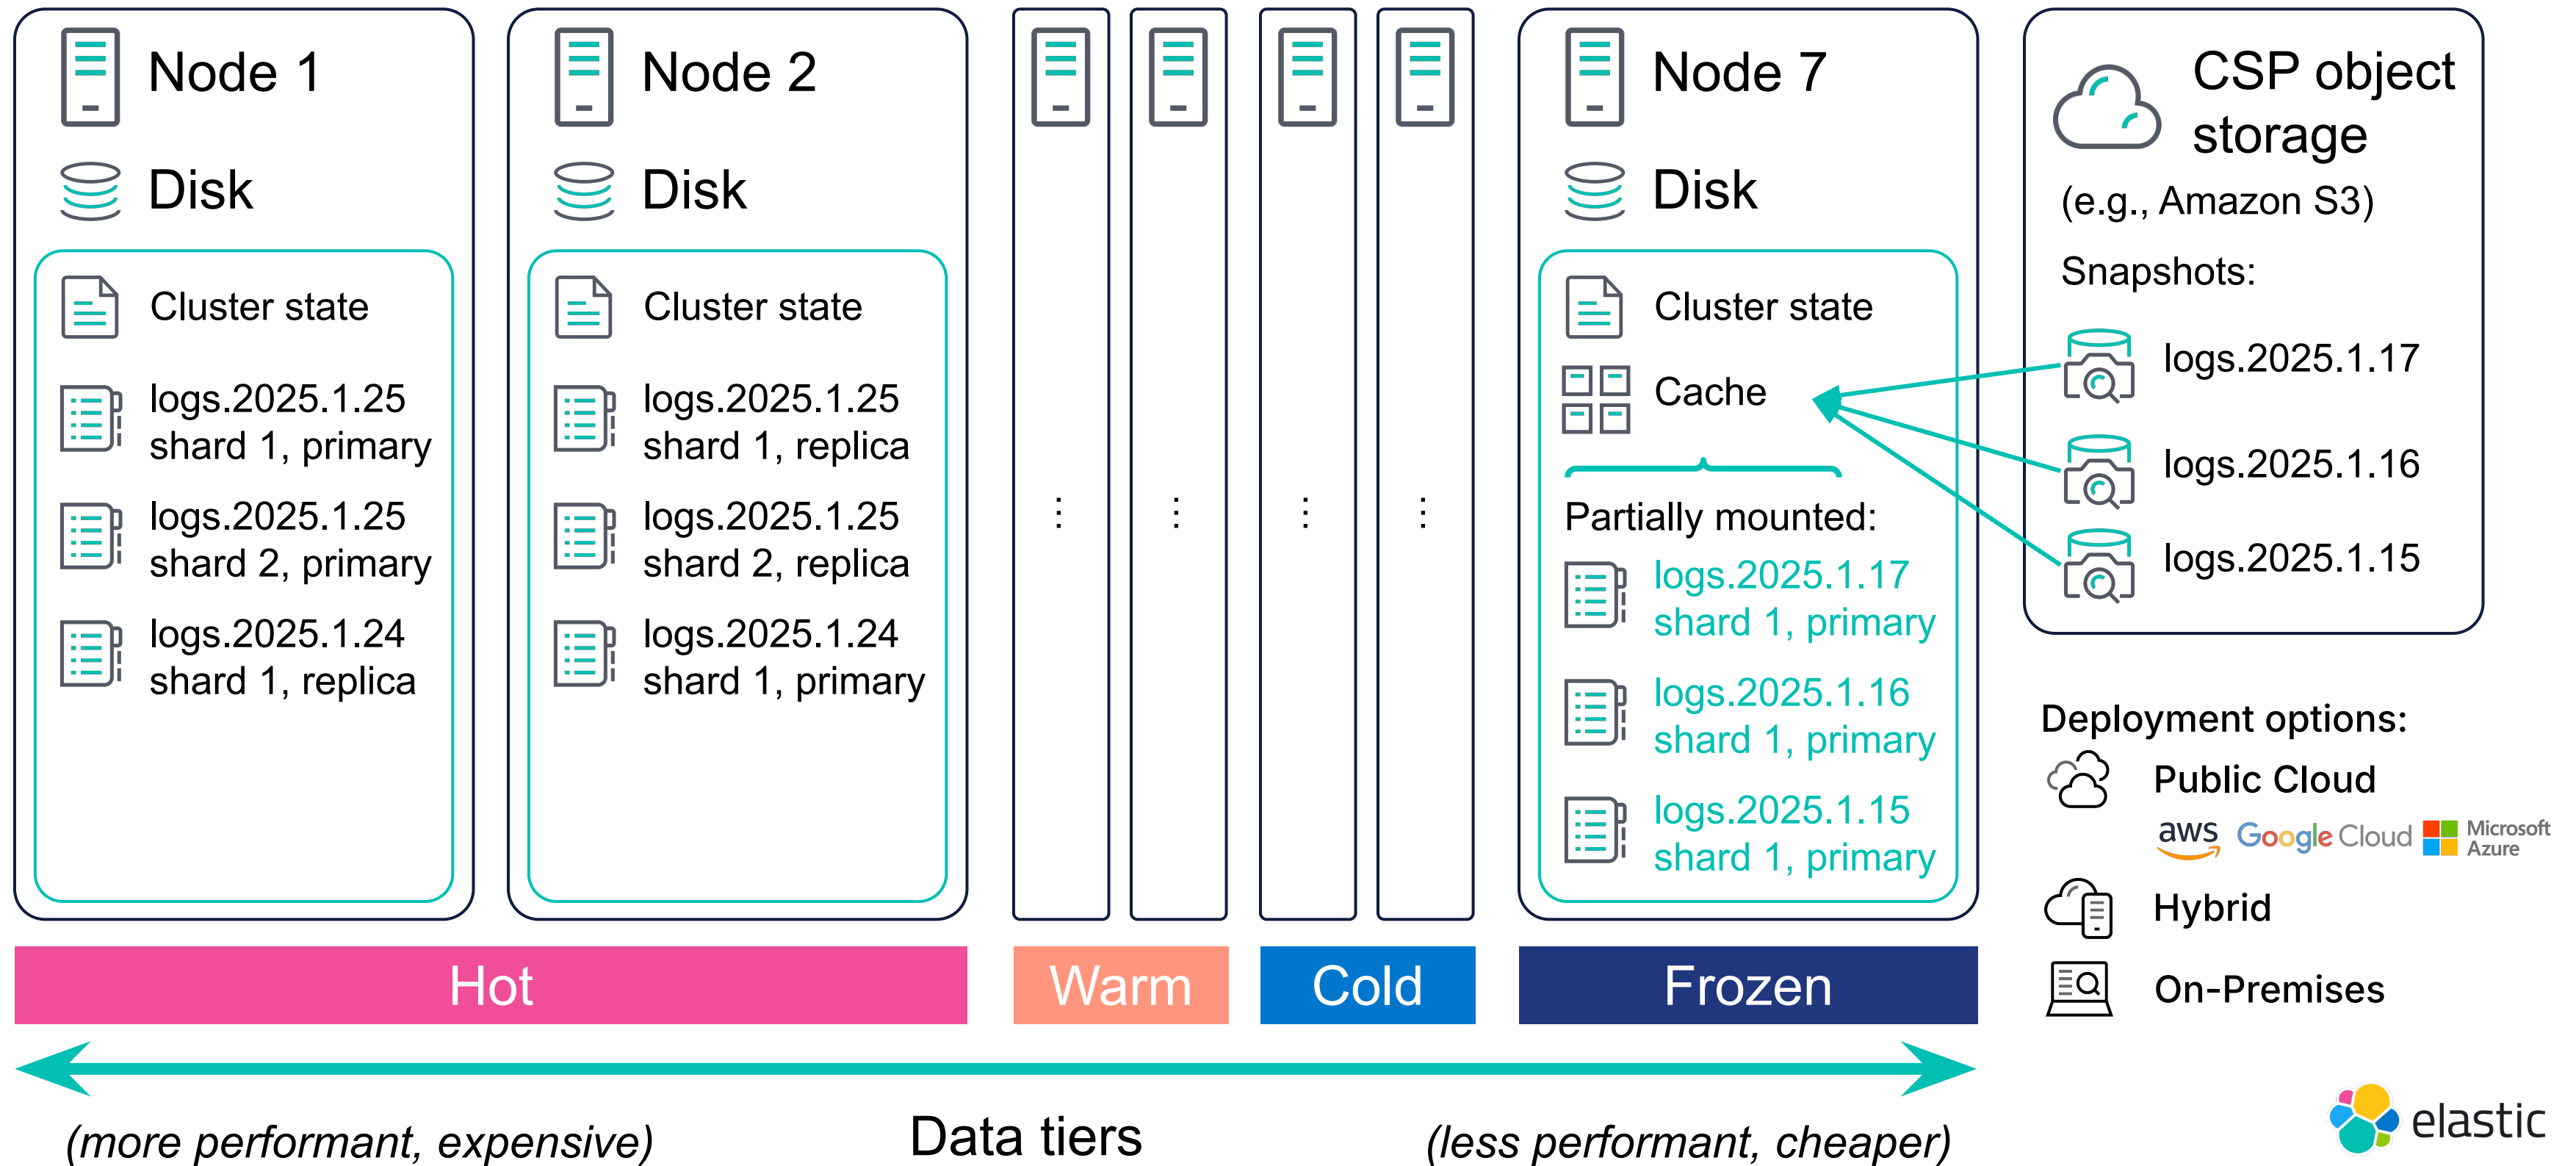
Elasticsearch in a nutshell

Devoxx 2023 talk: kingherc.com/archives/36



- An **index** comprises of Lucene segments (~ inverted indices)
- ES scales with index **shards** (primary & replica)
- Searches & aggregations across distributed shards
- Shards recover from disk & other nodes, and relocate
- **Cluster state** contains core metadata (e.g., shard assignments)
 - Master node updates it with a consensus protocol
- ES scales time-series data (e.g., logs) with data streams
 - Chain of older read-mostly indices and one write index
- Index Lifecycle Management (ILM) passes indices through data tiers
 - **Hot**, performant tier for recent updateable data
 - **Frozen**, cheaper, partial caching of static snapshots on an object store

Shared-nothing stateful architecture



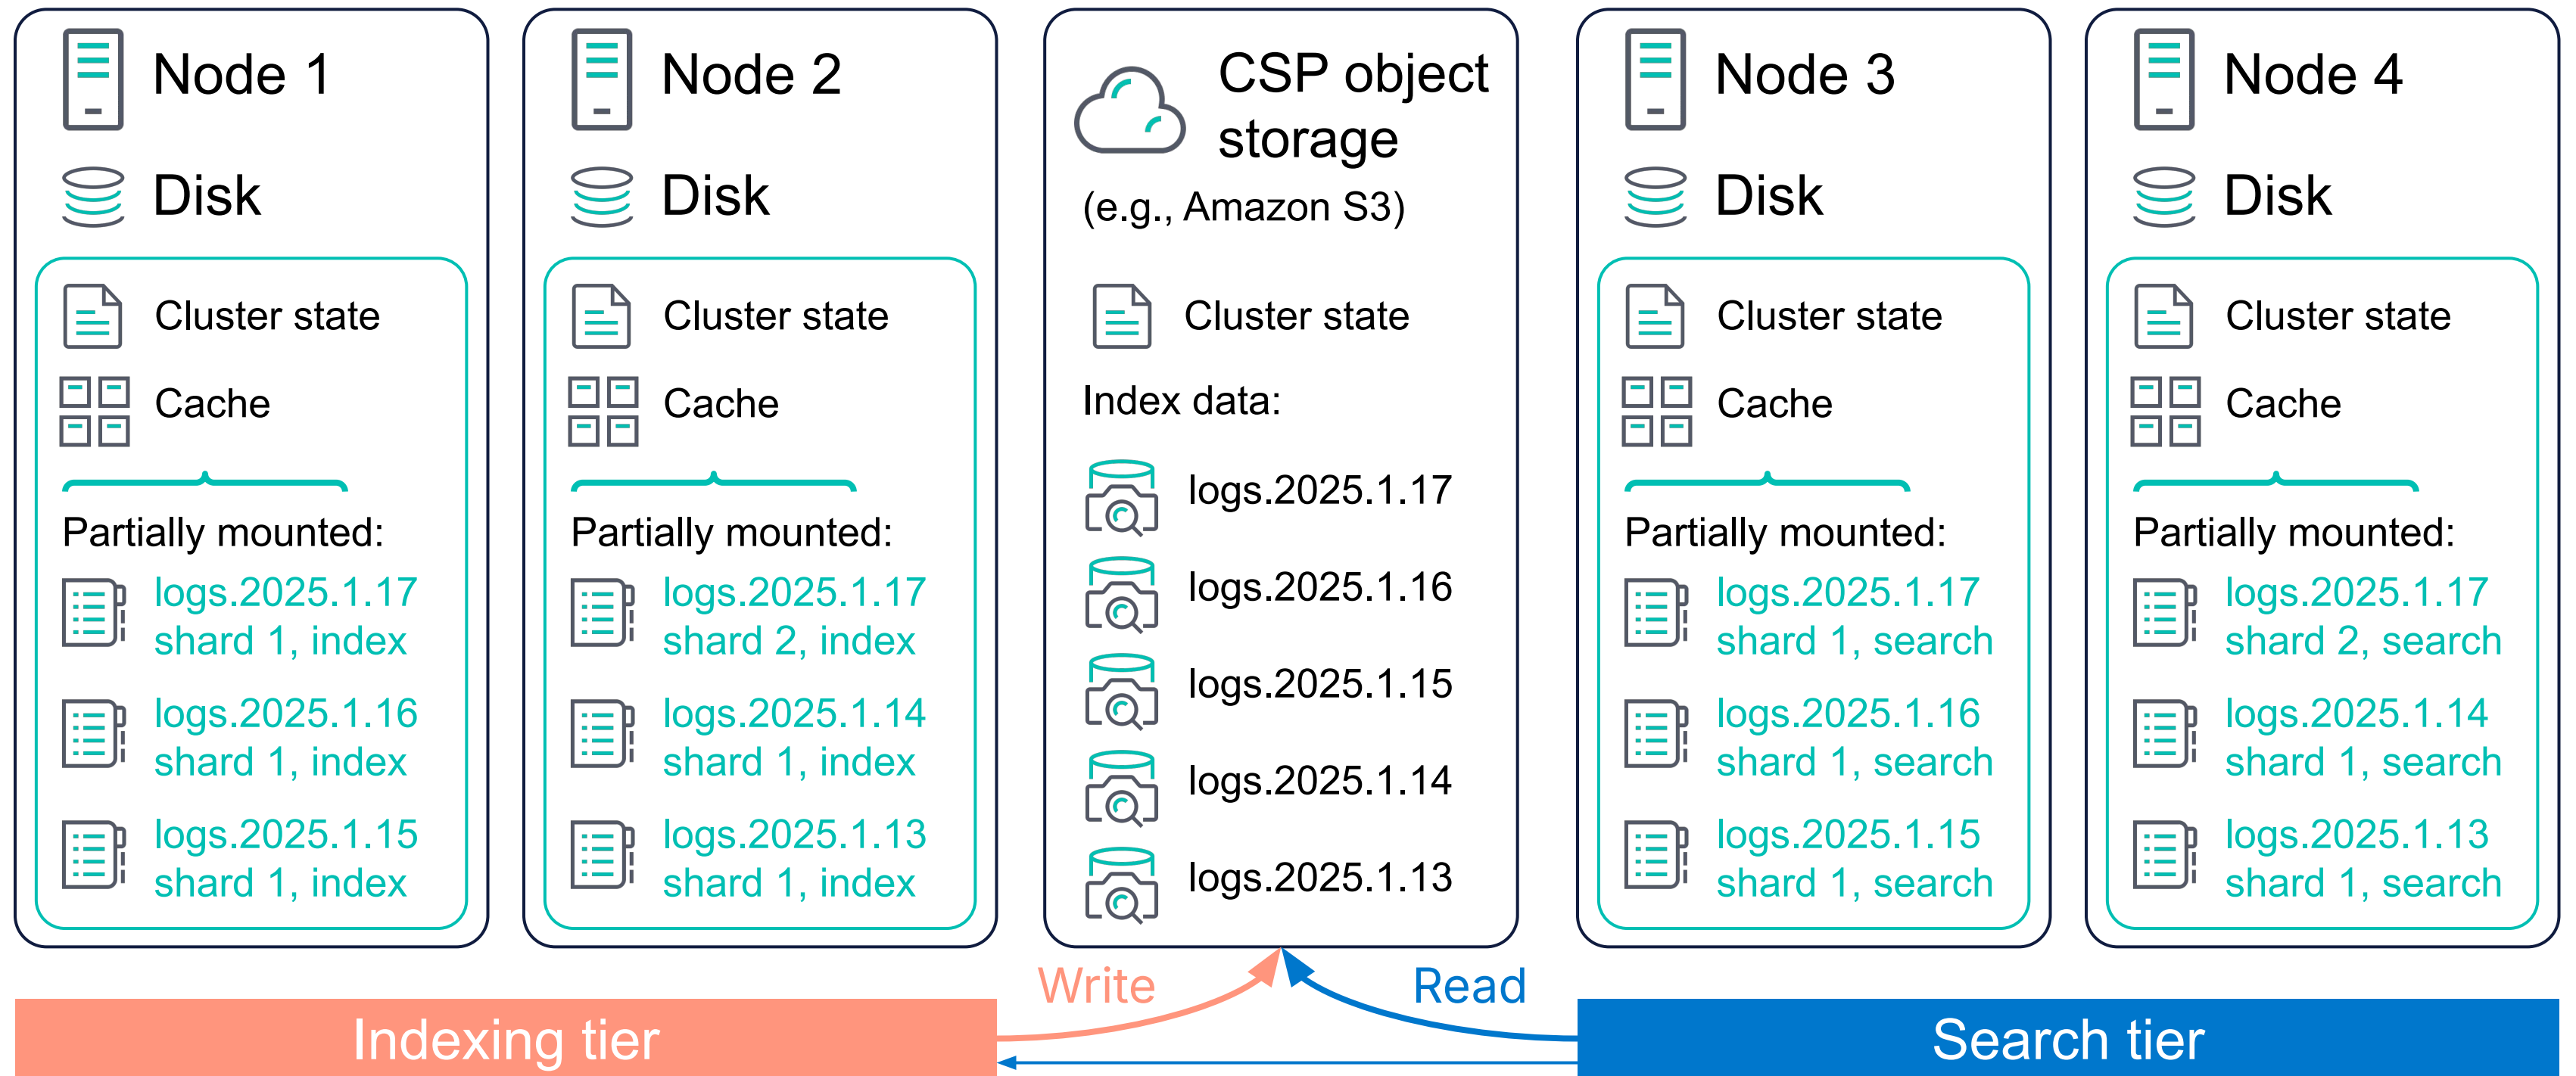
Disadvantages of stateful

- User defines cluster for their performance
 - Considers datase & workload size, and H/W (RAM, CPU, disk)
- Provisioning CPU & RAM coupled with storage → underutilization
- CPU shared between ingestion & search workloads → overprovisioning
- User considers different data tiers and ILM policies
 - To balance cost, performance, and high availability
- Adding/removing nodes moves data w/ inter-node traffic → slow, expensive
 - Leads to infrequent upgrade cadence, lowering security (e.g., CVEs)
 - Admins need to deal with upgrade issues, e.g., BWC issues


Serverless Elasticsearch

- SaaS (software as a service)
- Exploits CSP: resilience & scalability over commodity H/W
- Stateless: decouples compute from storage
 - Offload data to an affordable highly available object store → no replicas
 - Partial caching to dynamically load data as required by ingestion or searches
- Same API & read-after-write semantics as stateful ES
- Pay-as-you-go model for data retention, indexing, searching
- Without worrying about cluster specifics, sizing, upgrading, HA
- 2 simplified data tiers: (a) indexing, (b) search
 - Autoscaling based on workload

New stateless architecture



Thin shards



- Loaded with metadata until data is indexed/searched
- Indexing generates files, uploads them and deletes them when unused
- Searching loads from object store only the required data into the cache
- Indexing & searching virtually limitless data
- Not dependent on disk, not restricted by disk capacity → stateless
 - Durability is guaranteed by the object store
- Shard relocation (and recovery) is very fast
 - Indexing shard has a hand-off phase to transition workload
 - Removing/adding nodes is quick → autoscaling
 - Transparent frequent rolling upgrades to a cluster

Stateful example

To understand data inside a shard

Data inside a shard (stateful)

Node 1

New docs [1, 5]

Index "logs.2025.1.17", shard 1, primary

Lucene in-memory indexing buffers



Translog
generation:
Docs:



1

[1, 5]


Data inside a shard (stateful)

Node 1

New docs [6, 12]

Index "logs.2025.1.17", shard 1, primary

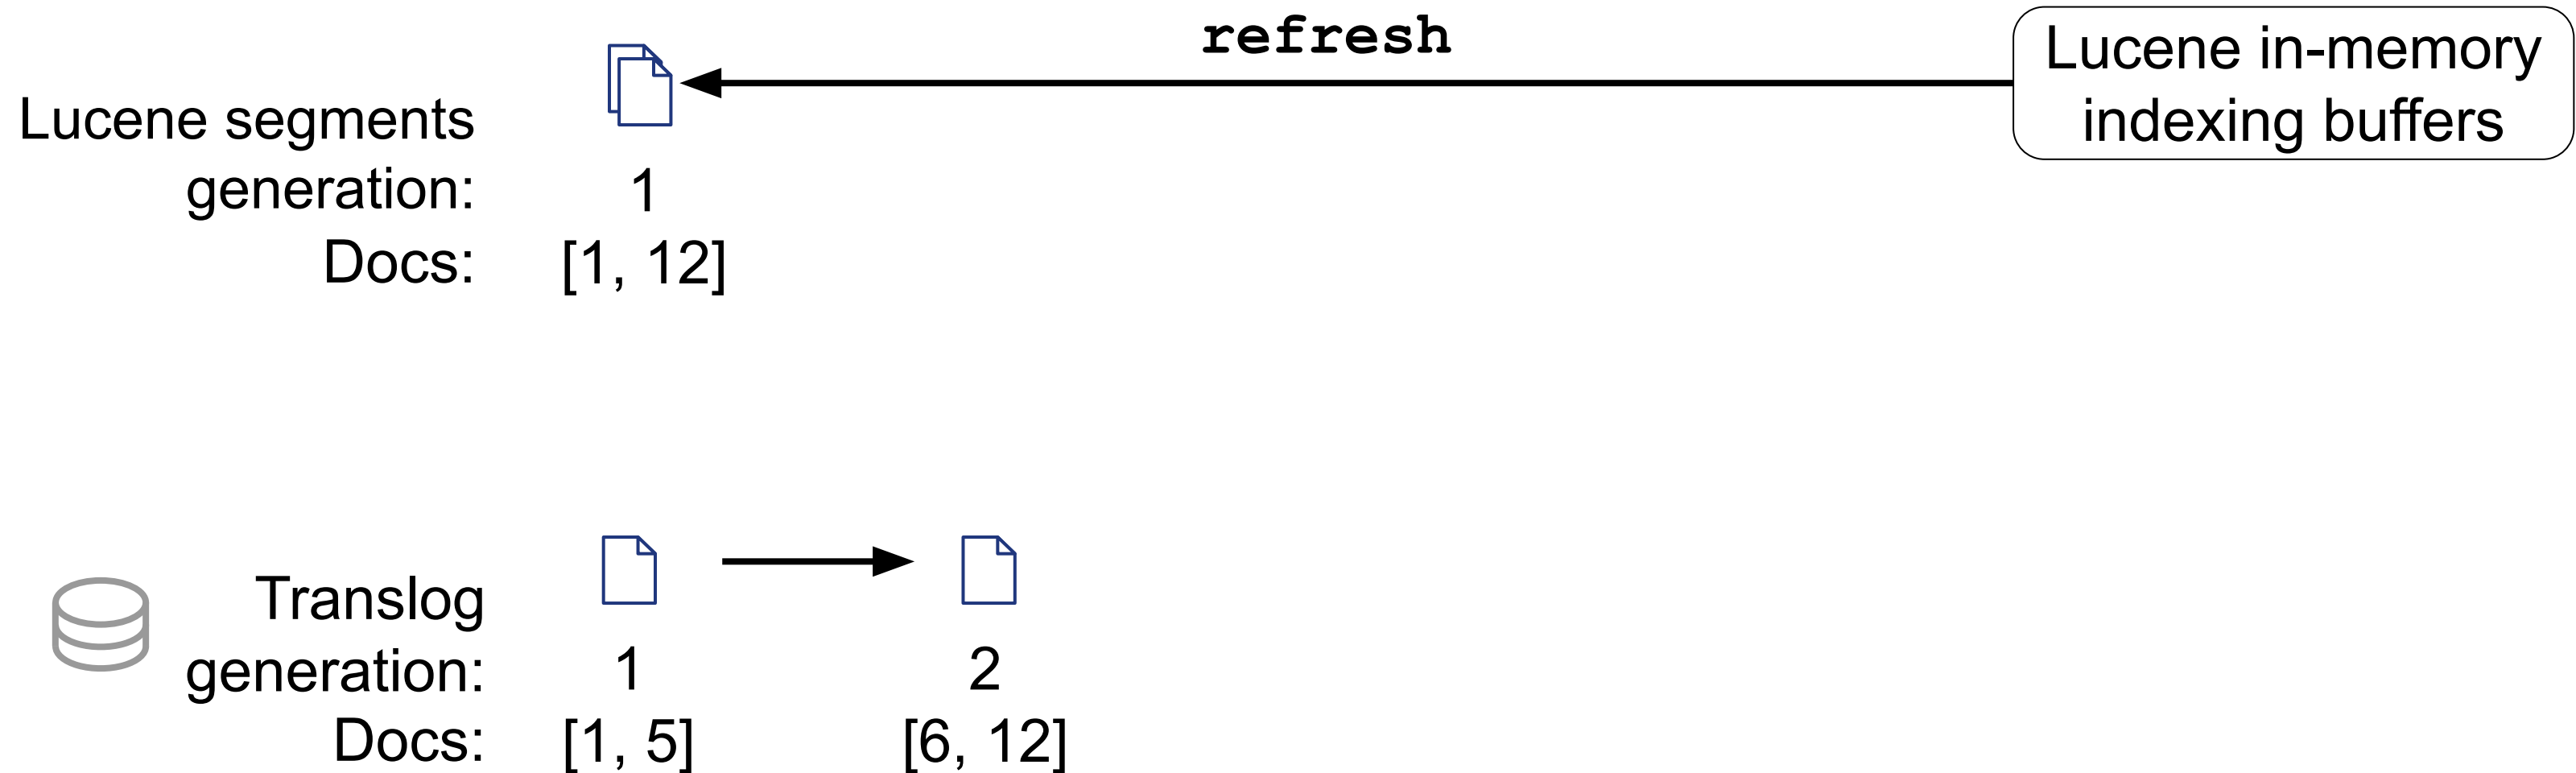
Lucene in-memory indexing buffers



Data inside a shard (stateful)

Node 1

Index "logs.2025.1.17", shard 1, primary




Data inside a shard (stateful)

Node 1

Search documents

Index "logs.2025.1.17", shard 1, primary

Lucene segments
generation: 1
Docs: [1, 12]




Lucene in-memory indexing buffers

Translog
generation: 1
Docs: [1, 5]

→

2
[6, 12]



Data inside a shard (stateful)

Node 1

Index "logs.2025.1.17", shard 1, primary

Lucene segments
generation: 1
Docs: [1, 12]



1

[1, 12]



Translog
generation:
Docs:



1

[1, 5]



2

[6, 12]



3

[13, 15]

New docs [13, 15]



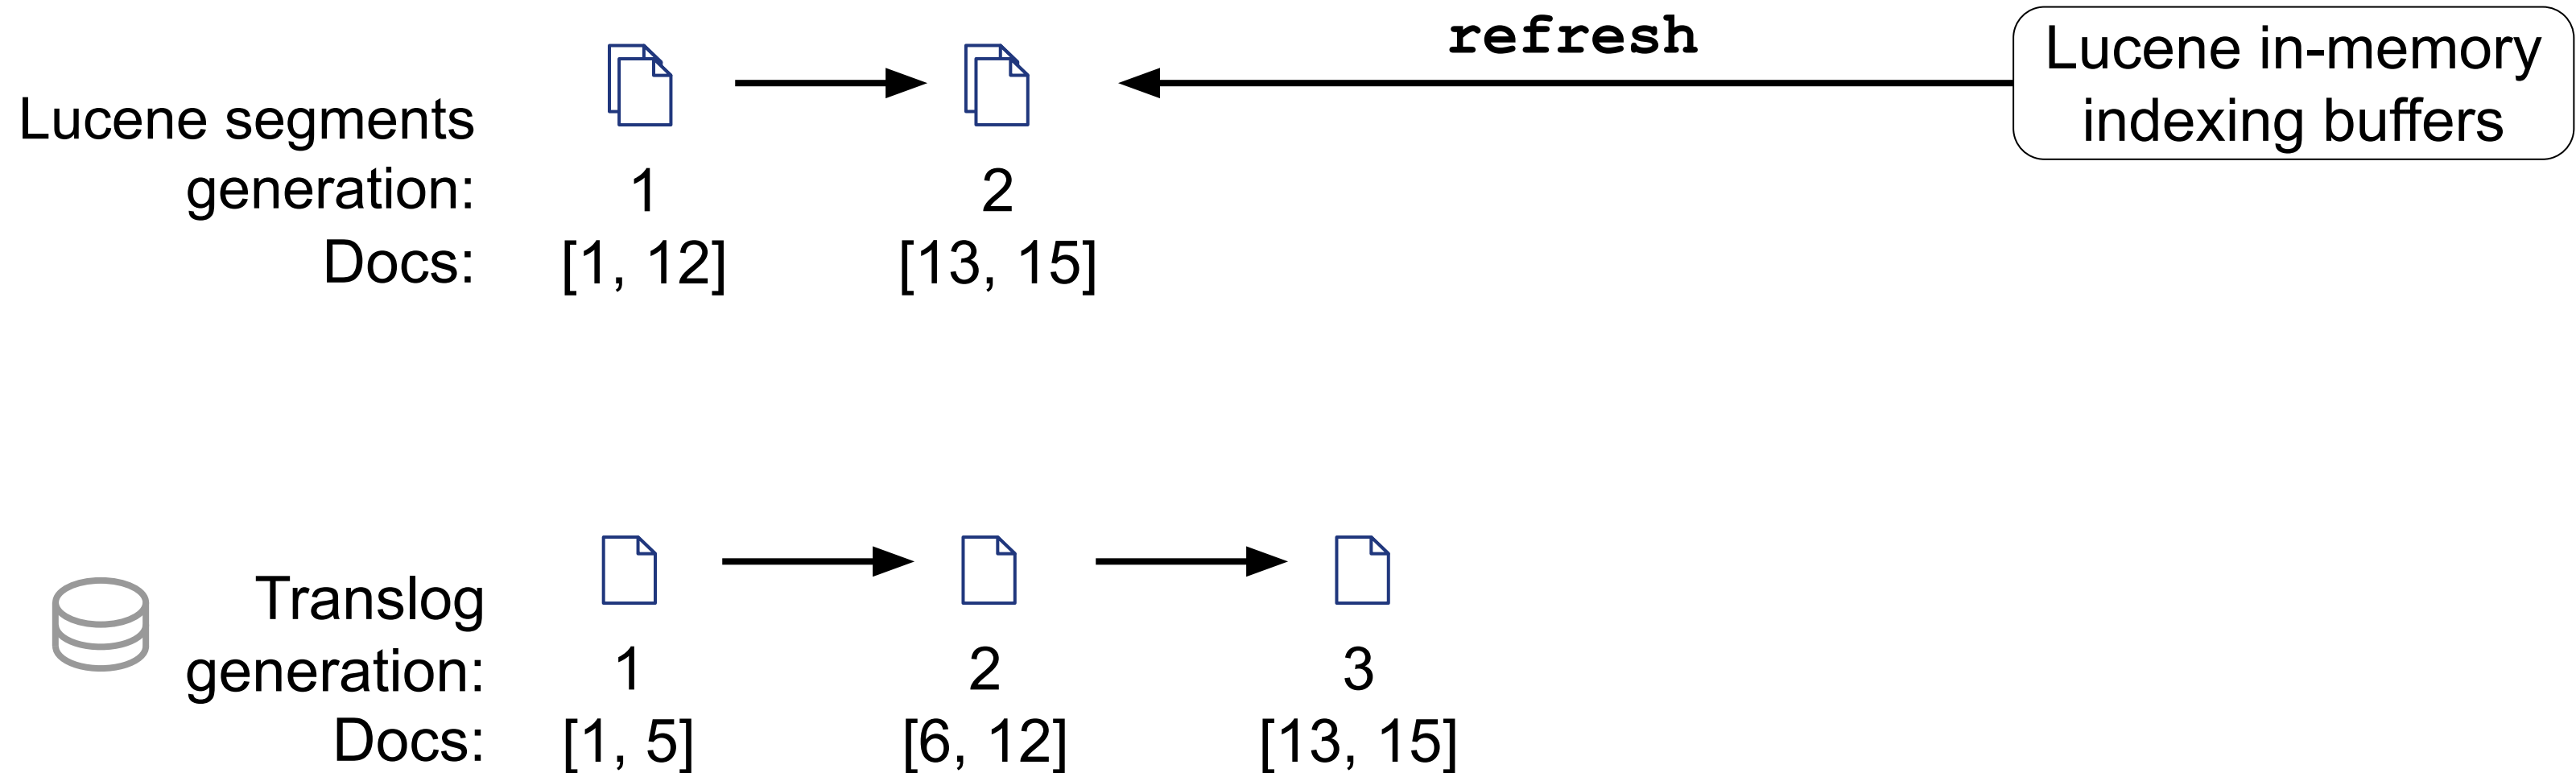
Lucene in-memory
indexing buffers



Data inside a shard (stateful)

Node 1

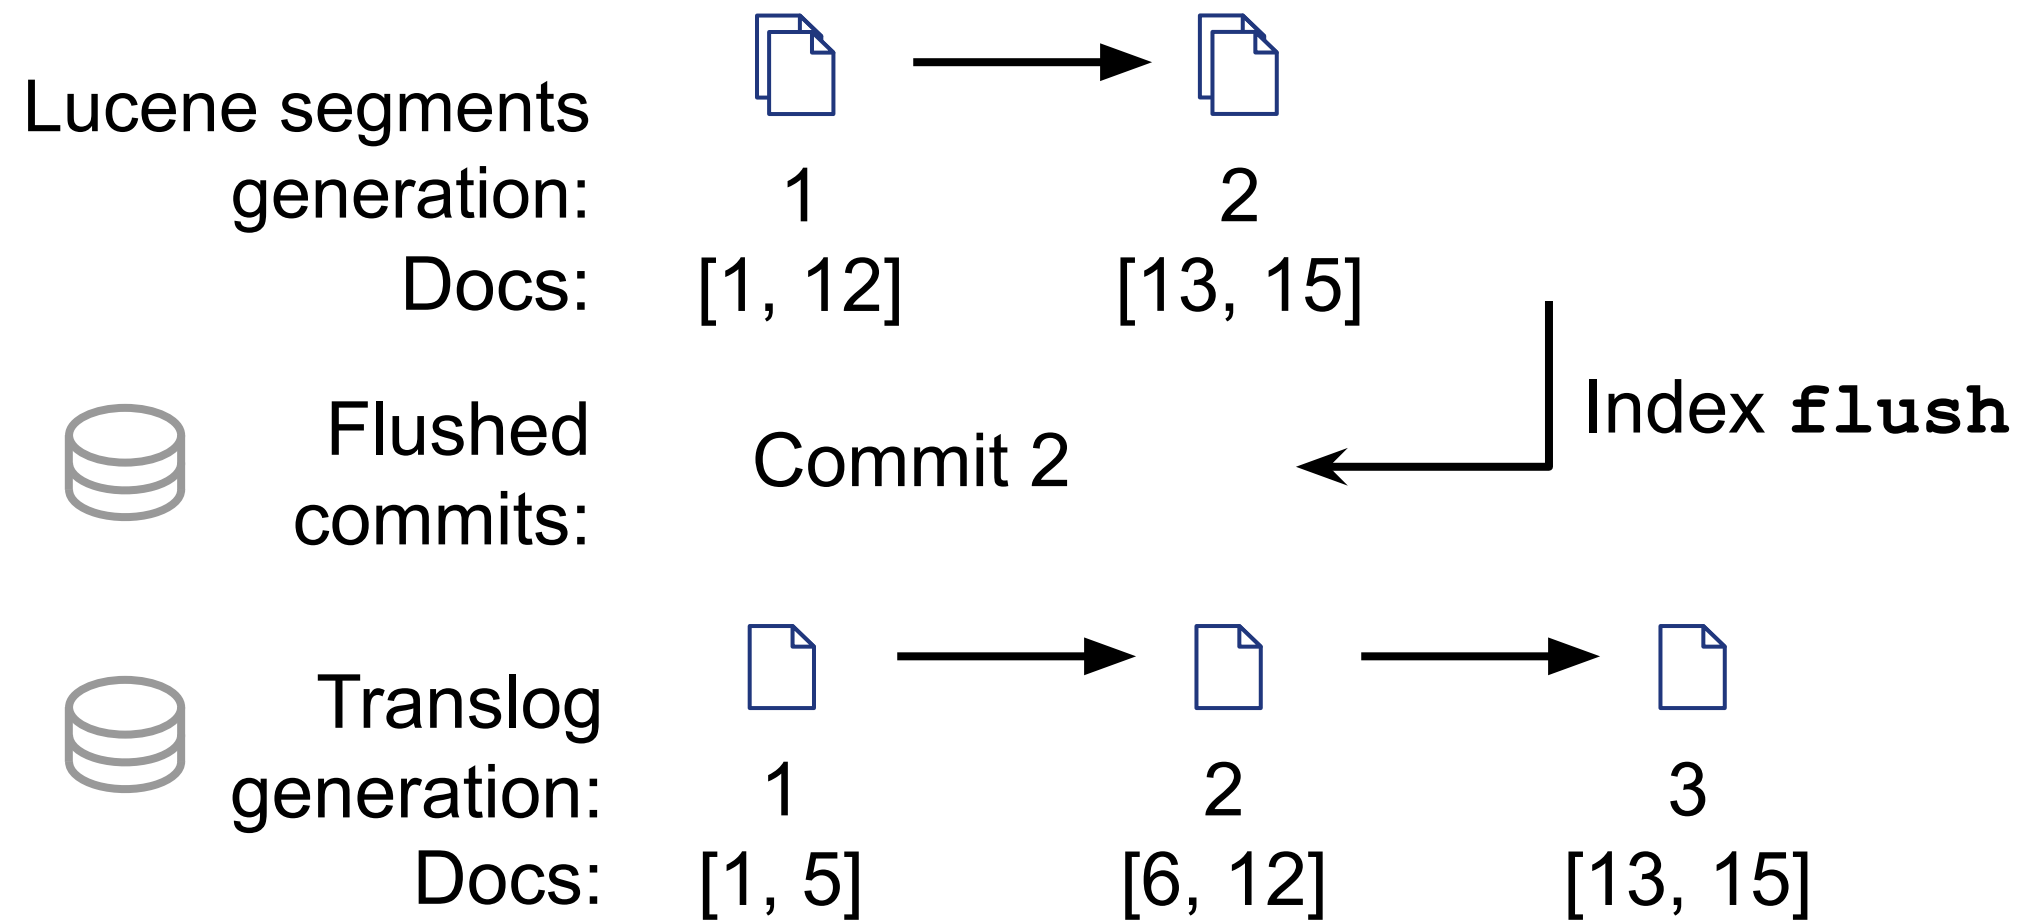
Index "logs.2025.1.17", shard 1, primary



Data inside a shard (stateful)

Node 1

Index "logs.2025.1.17", shard 1, primary




Lucene in-memory indexing buffers

Data inside a shard (stateful)

Node 1

Index "logs.2025.1.17", shard 1, primary



Flushed commits:

Commit 2



Translog generation:
Docs:

New docs [16, 20]

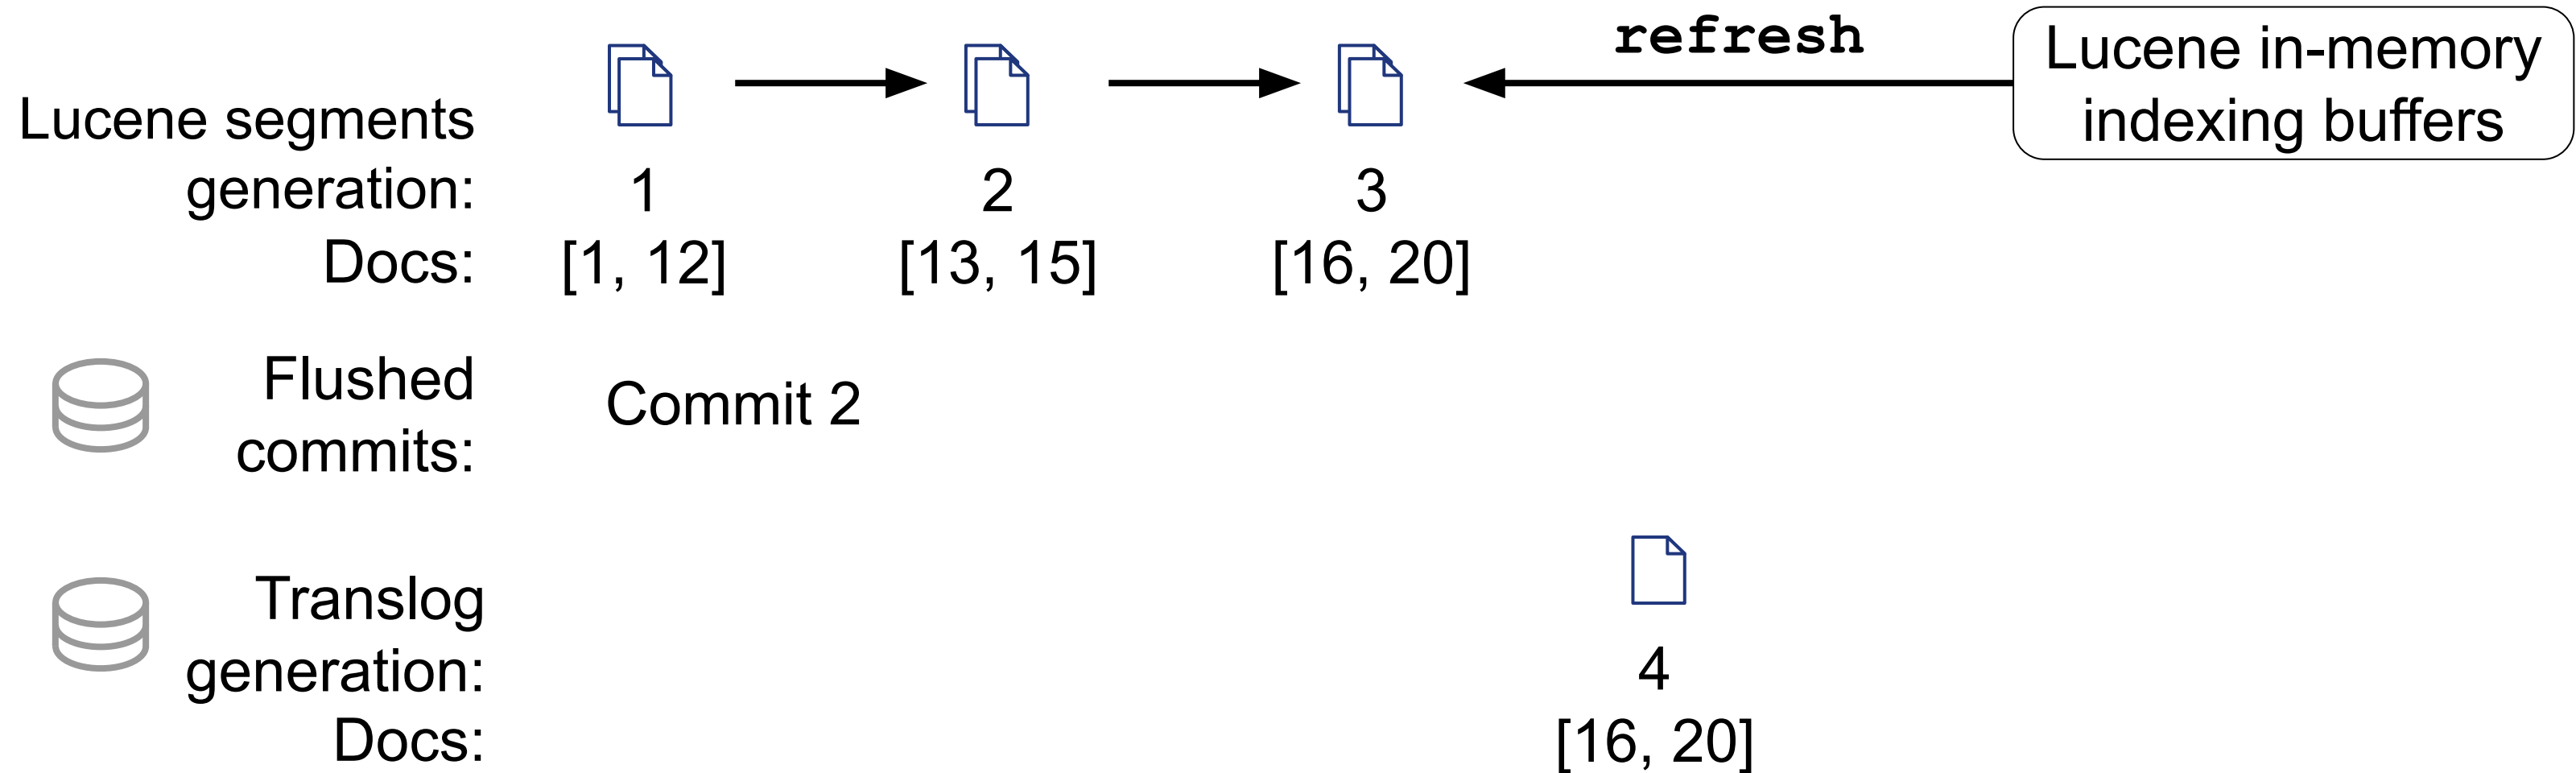
Lucene in-memory indexing buffers



Data inside a shard (stateful)

Node 1

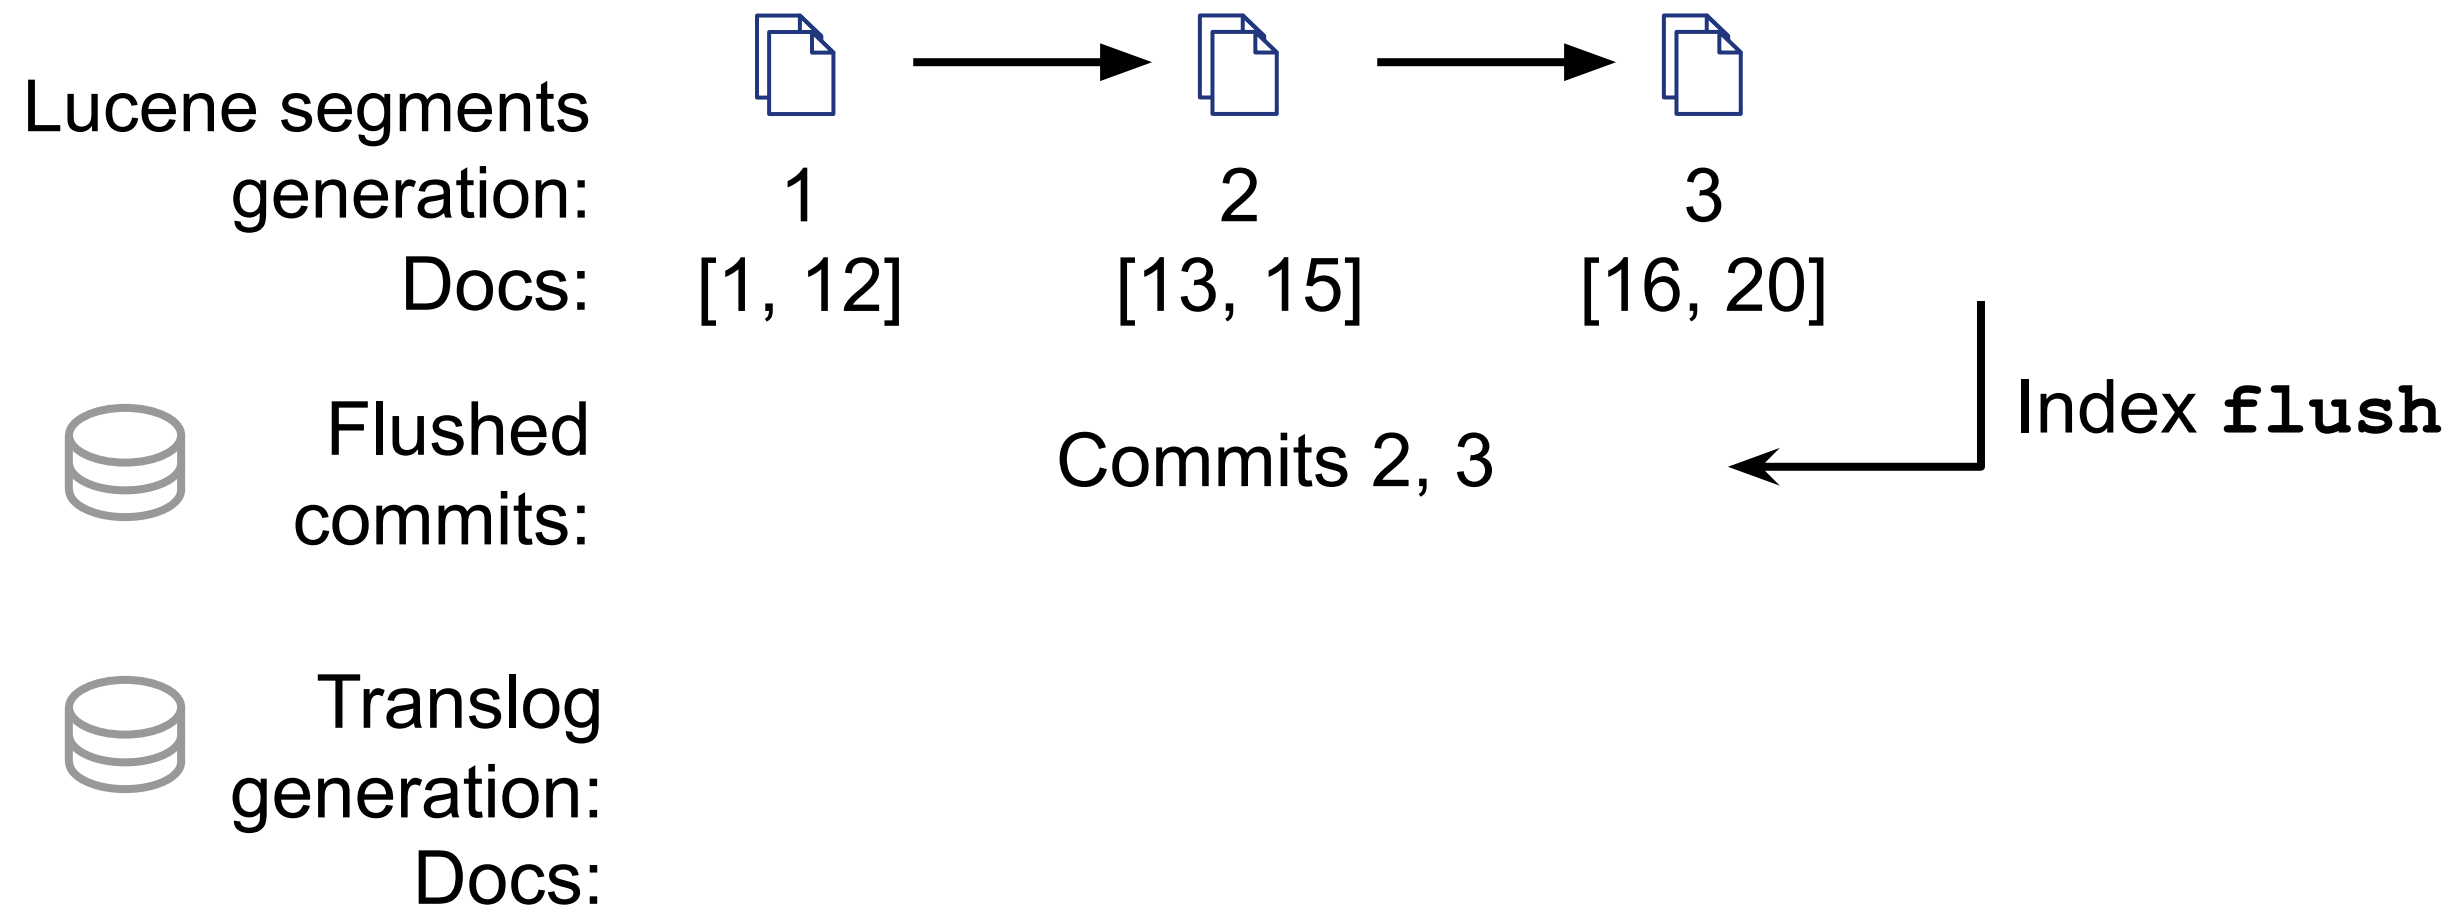
Index "logs.2025.1.17", shard 1, primary



Data inside a shard (stateful)

Node 1

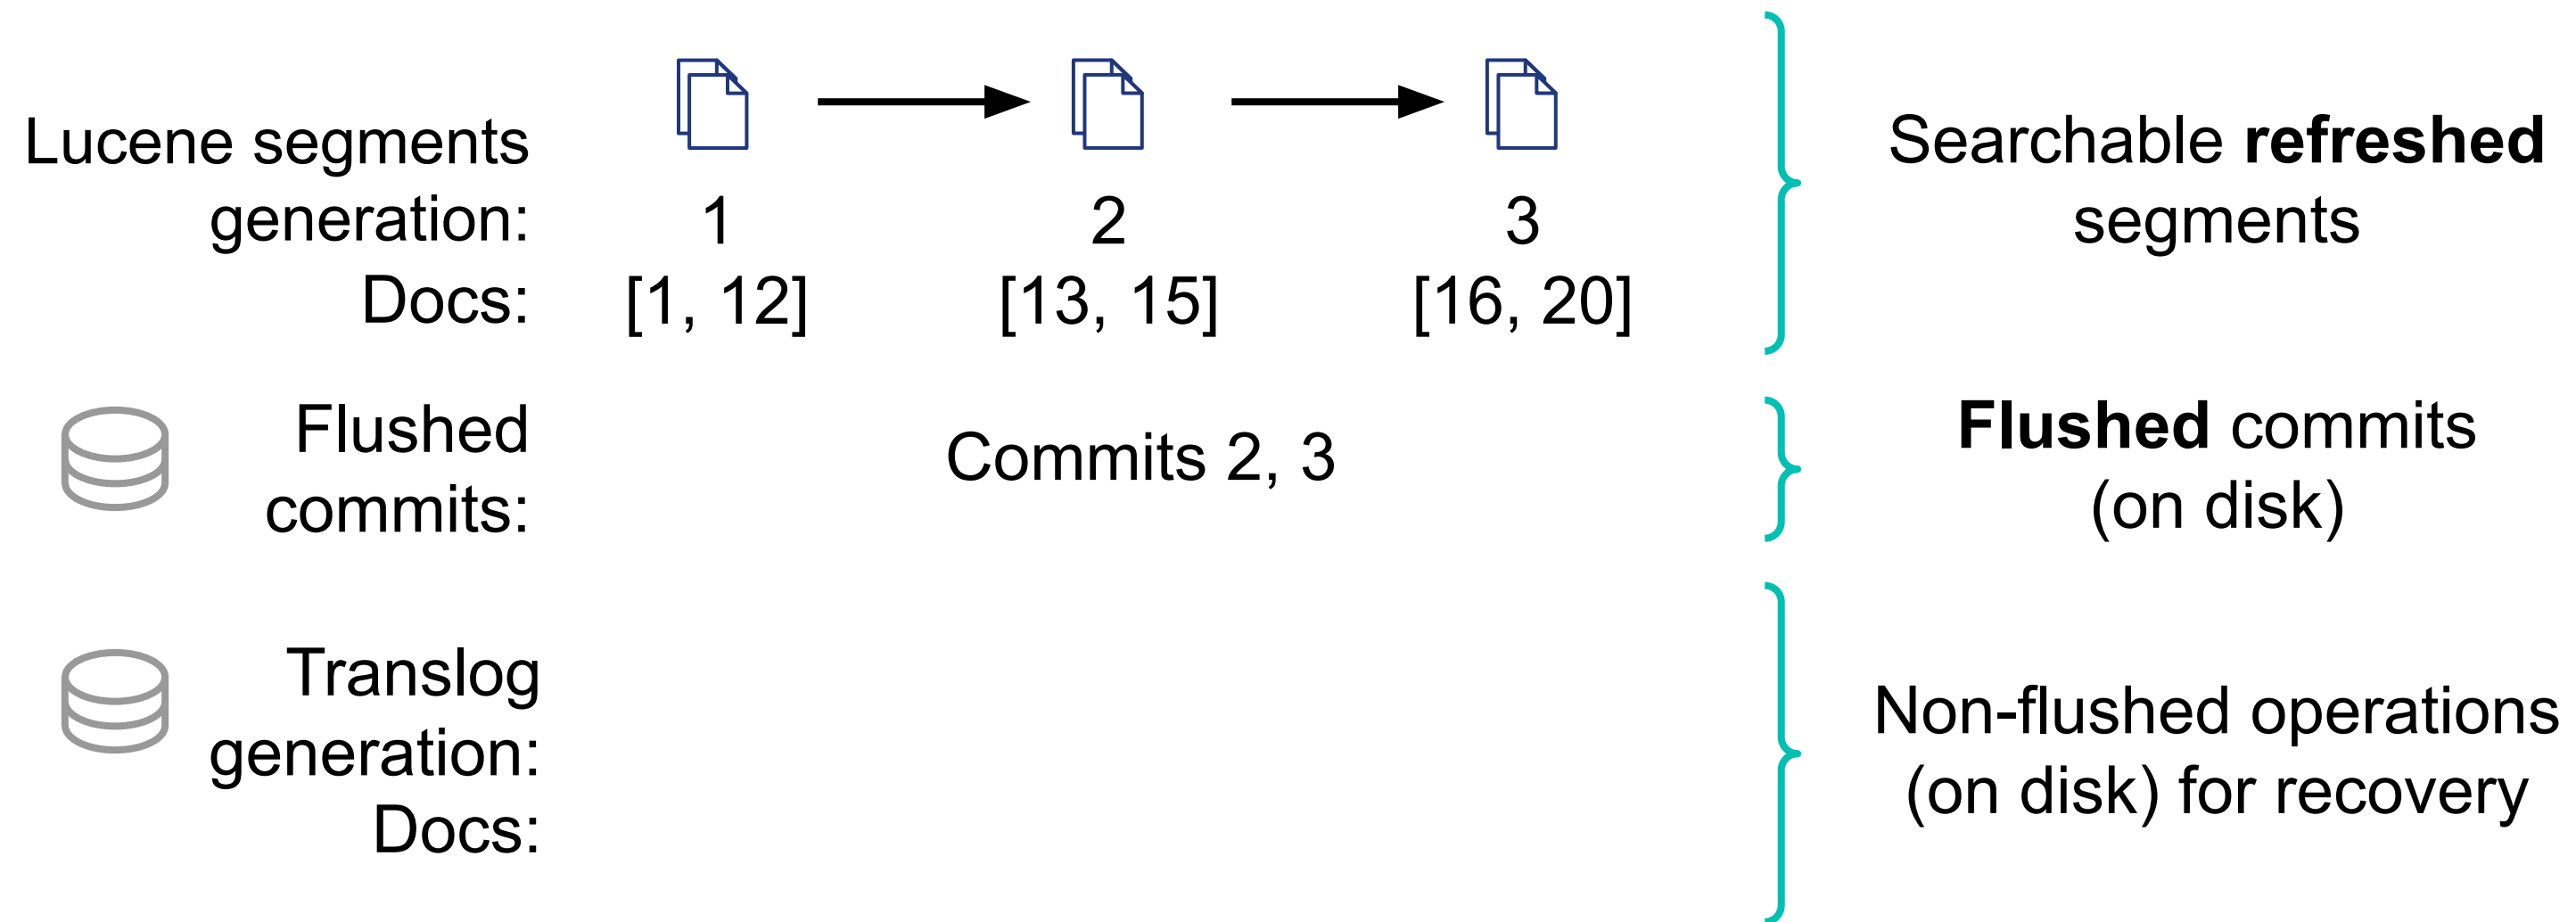
Index "logs.2025.1.17", shard 1, primary



Data inside a shard (stateful)

Node 1

Index "logs.2025.1.17", shard 1, primary



Stateless example

How we offload shard data to the object store

Offloading index data to the object store (stateless)

Node 1 (indexing)

New docs [1, 5]

Index "logs.2025.1.17", shard 1

Lucene segments
generation:
Docs:



Flushed
commits:



Translog
generation:
Docs:



1

[1, 5]

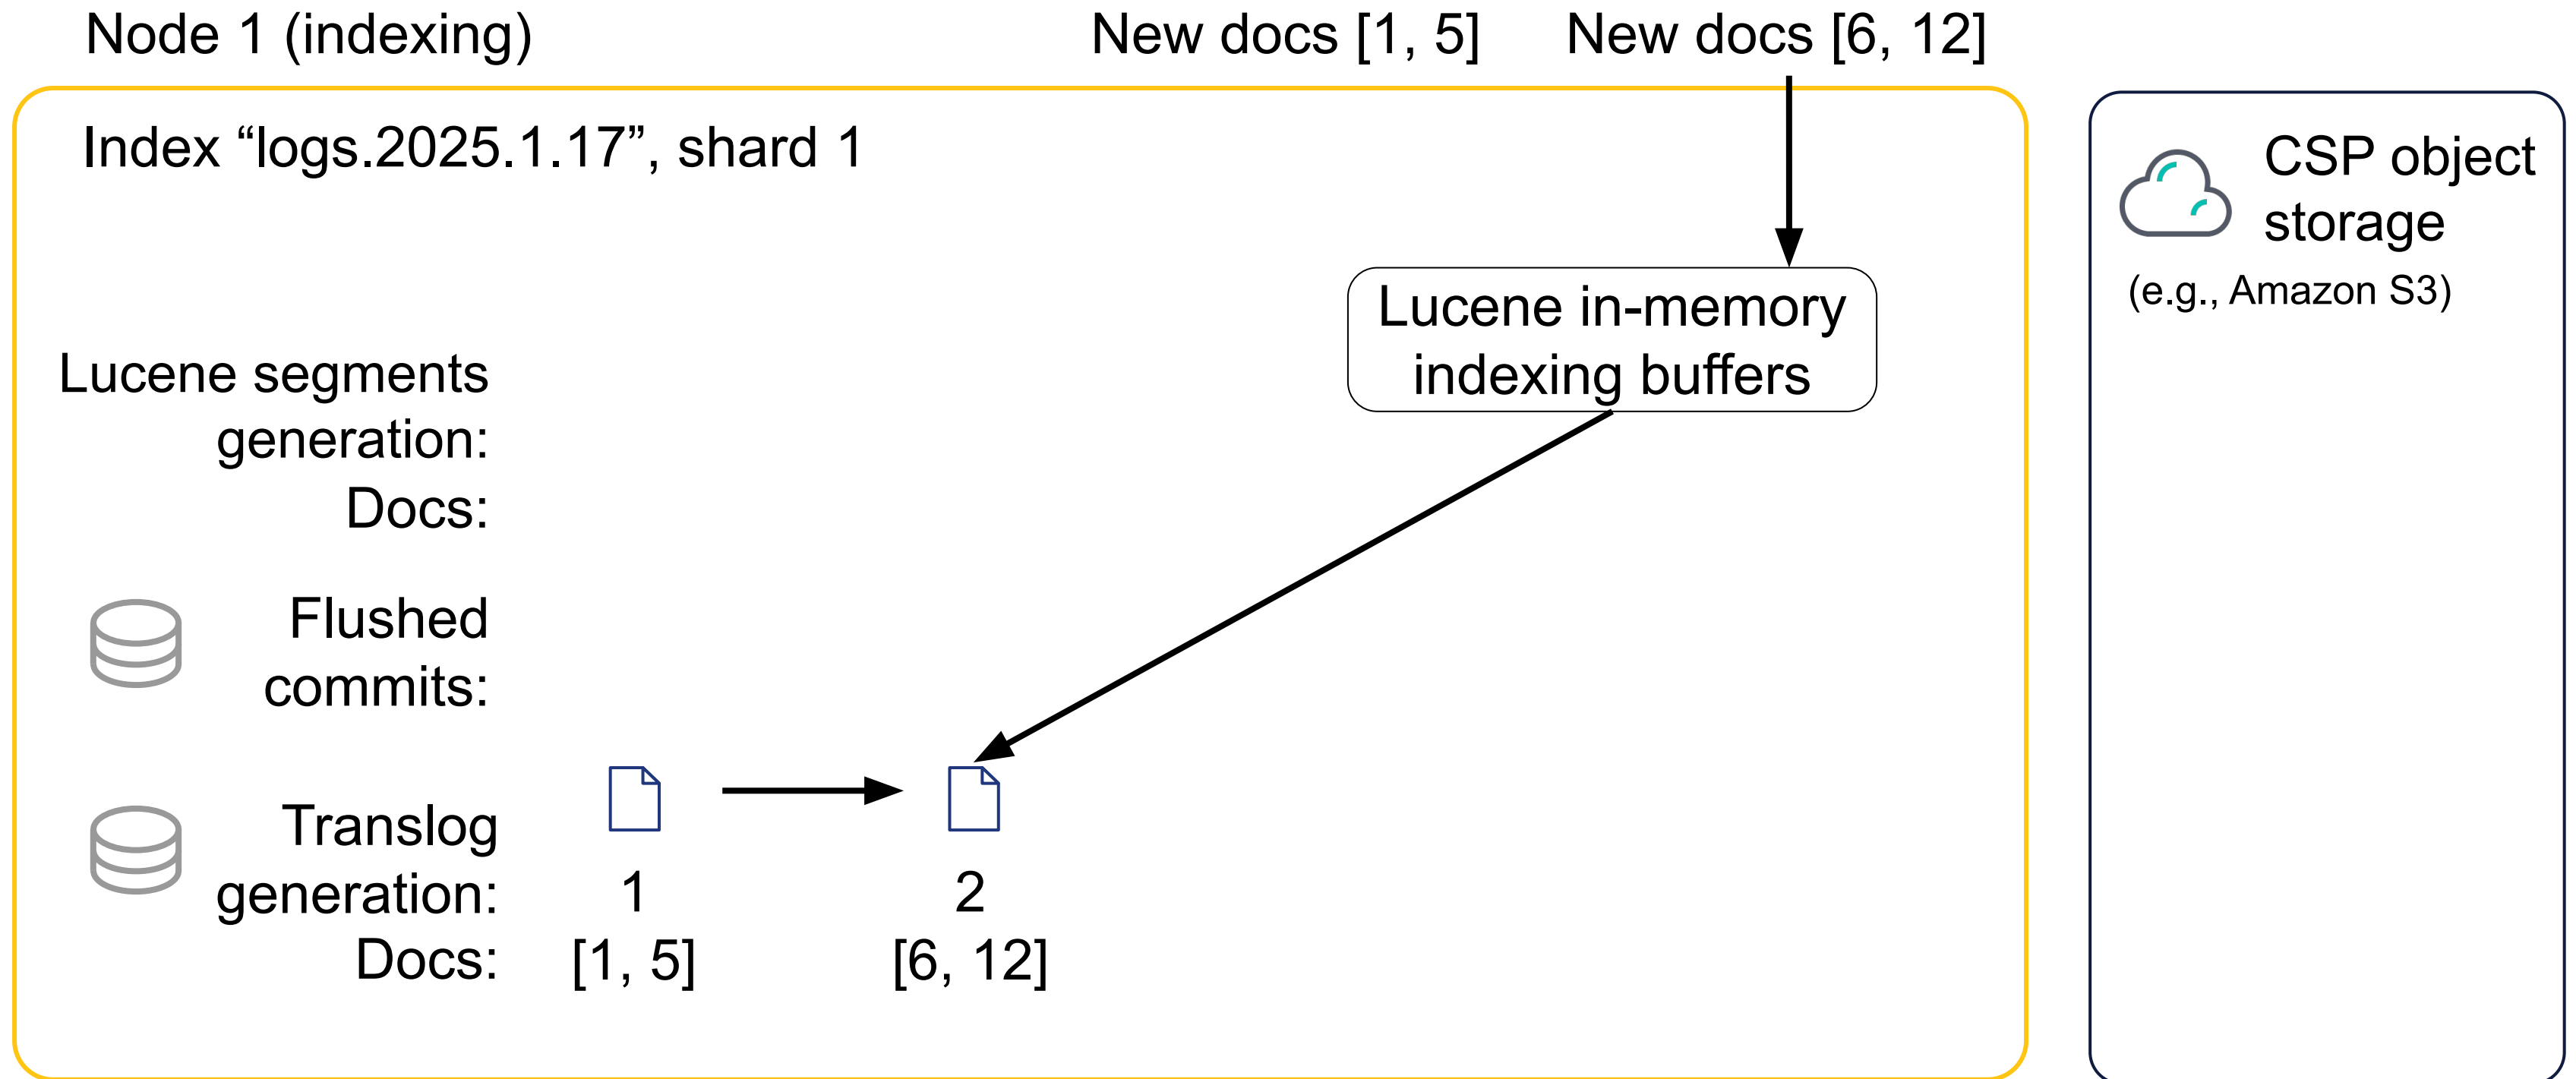
Lucene in-memory
indexing buffers



CSP object
storage

(e.g., Amazon S3)

Offloading index data to the object store (stateless)



Offloading index data to the object store (stateless)

Node 1 (indexing)

New docs [1, 5]

New docs [6, 12]

Index "logs.2025.1.17", shard 1

Lucene segments
generation:
Docs:



Flushed
commits:



Translog
generation:
Docs:



1

[1, 5]



2

[6, 12]

Lucene in-memory
indexing buffers

Upload (< 200ms)



CSP object
storage

(e.g., Amazon S3)

Translog data:

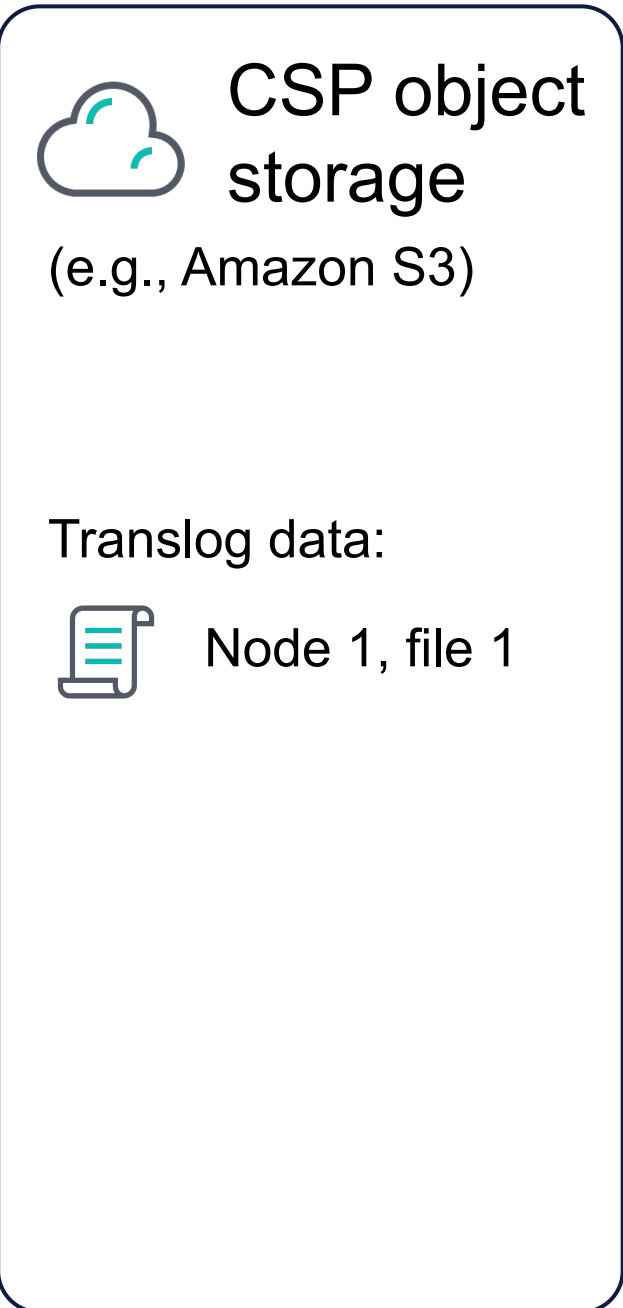
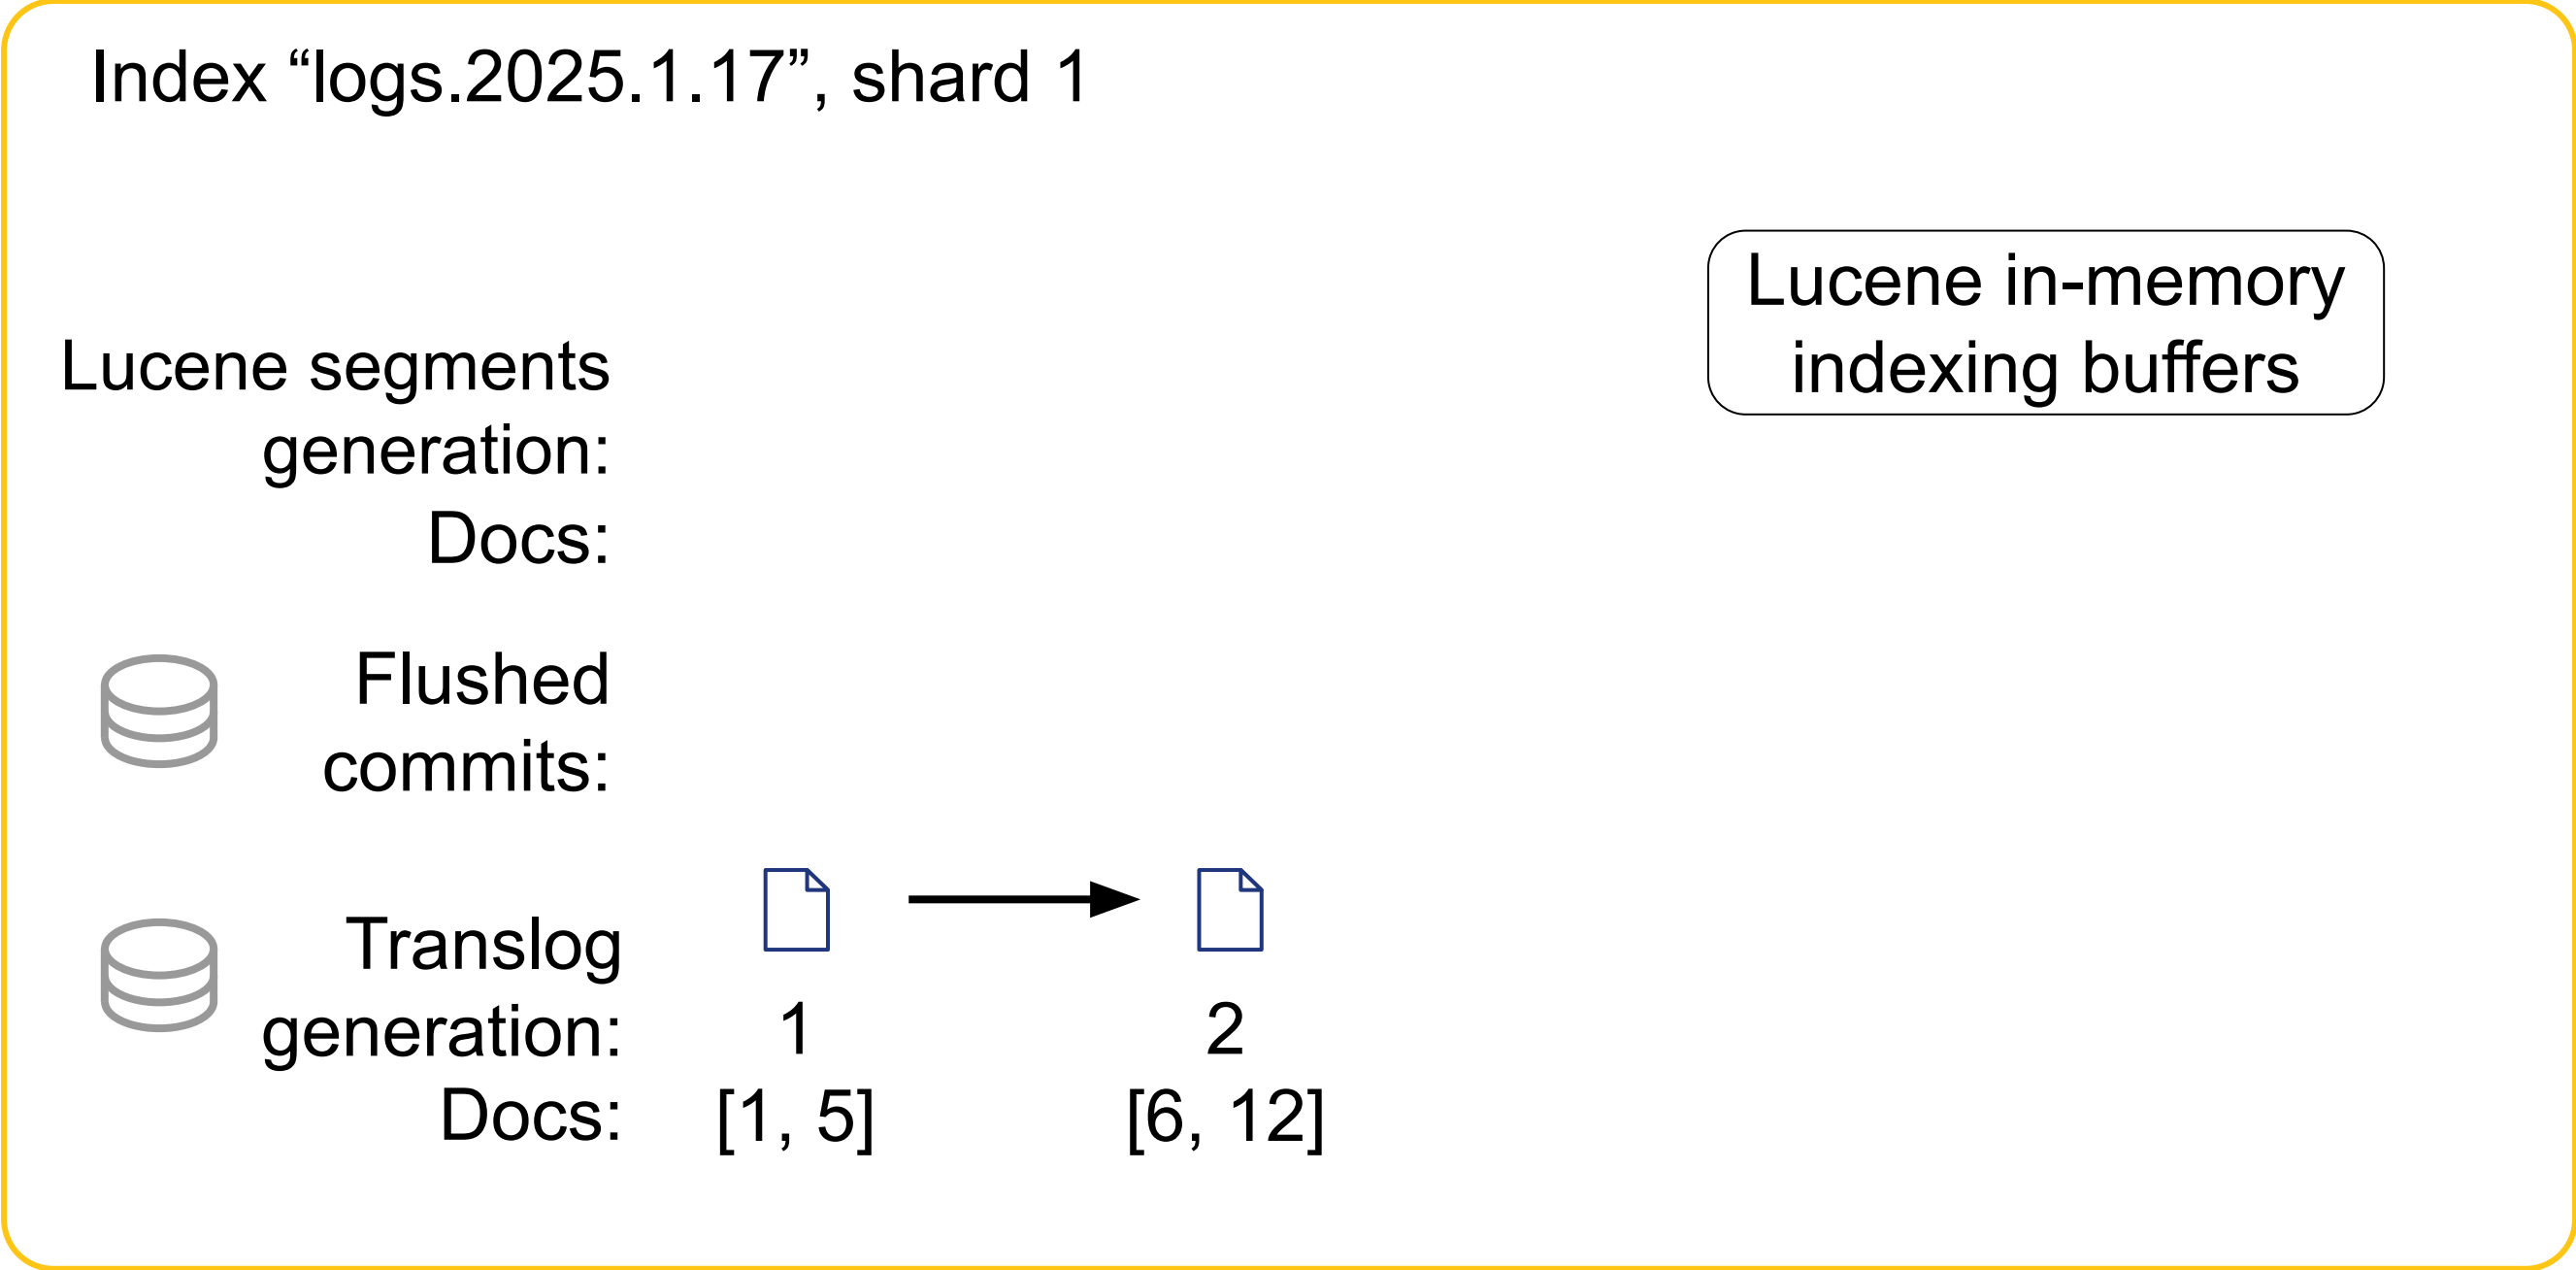


Node 1, file 1

Offloading index data to the object store (stateless)

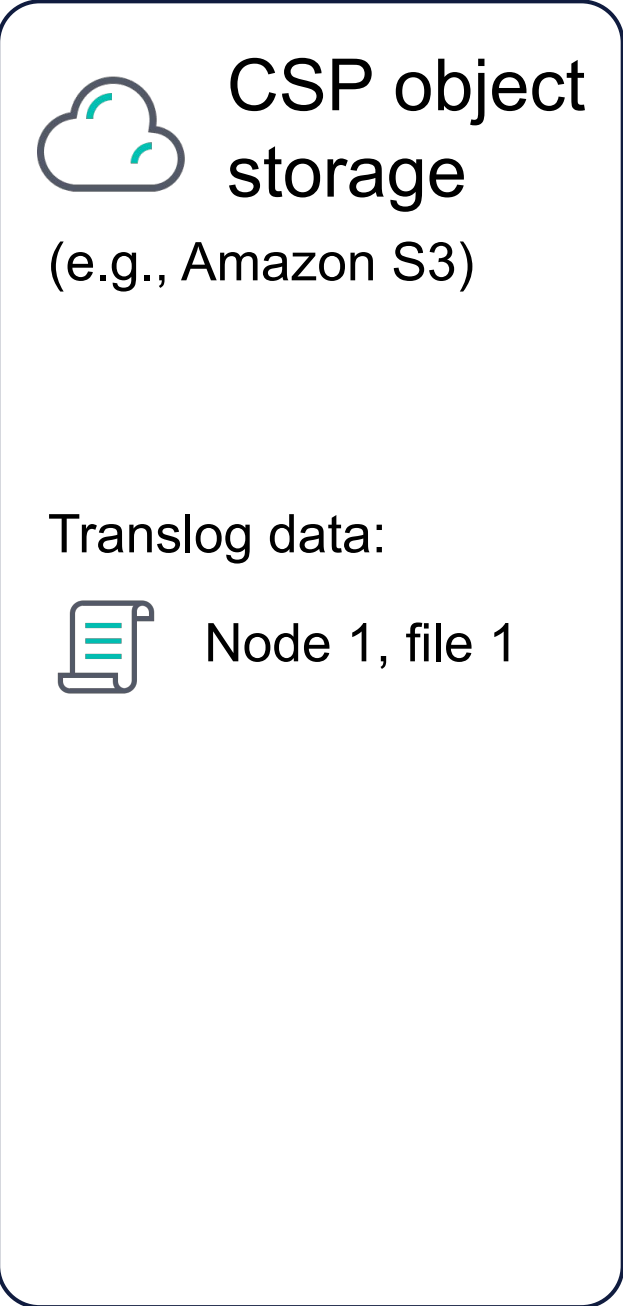
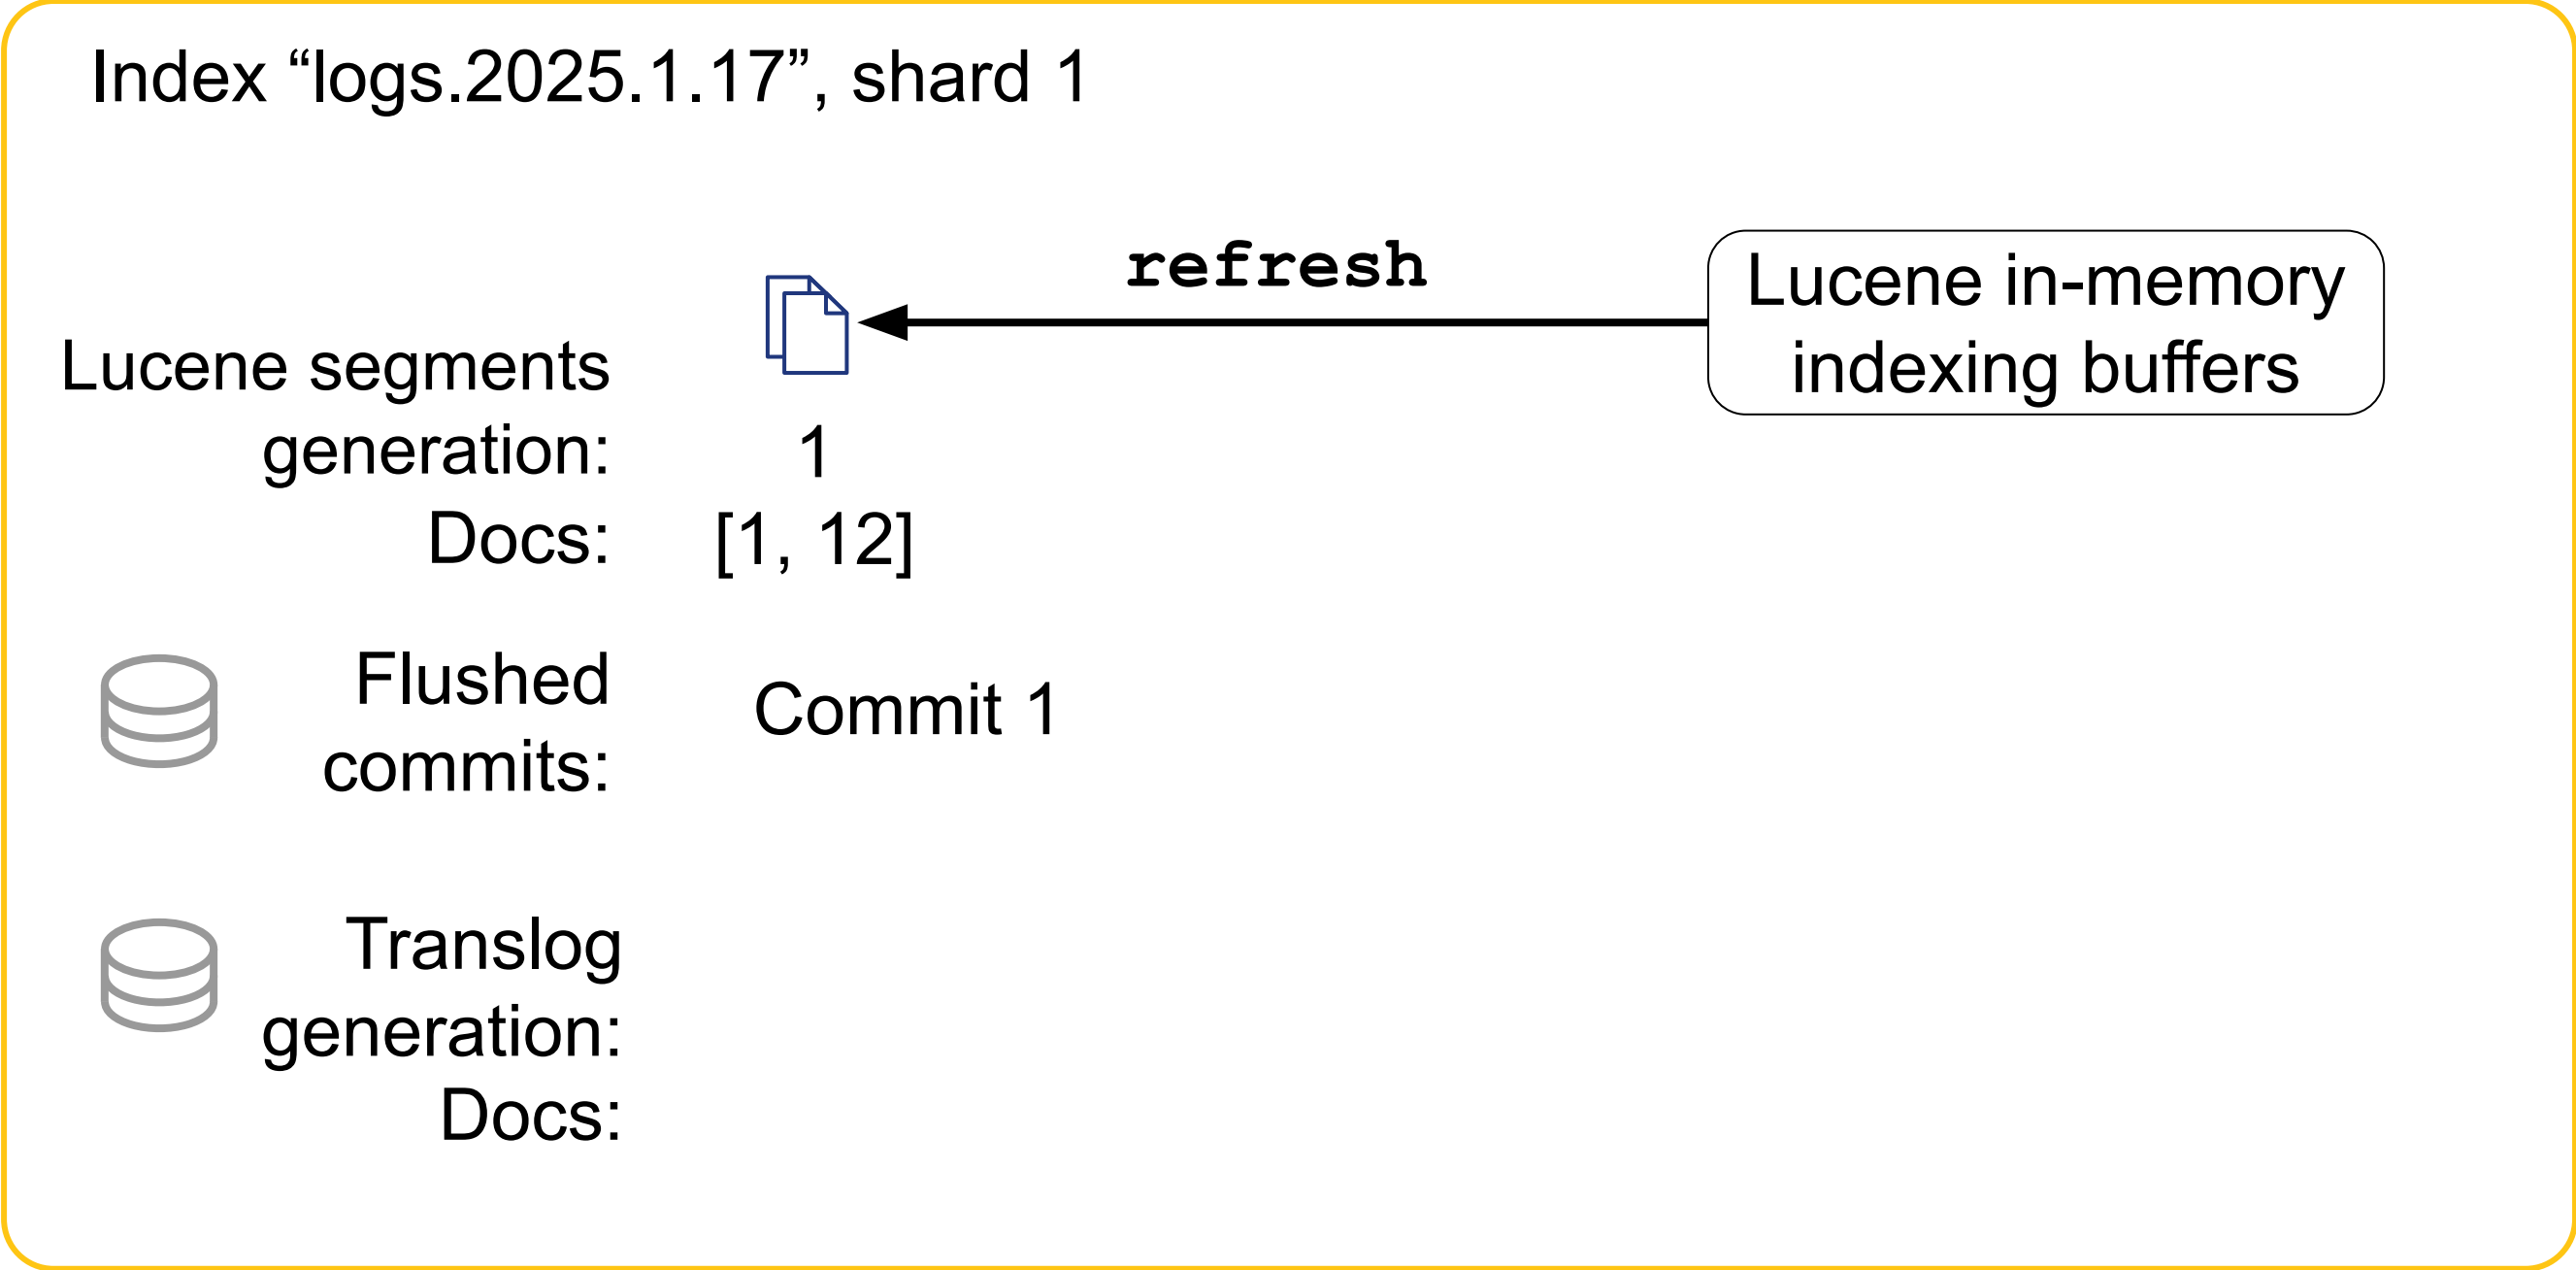
Node 1 (indexing)

Acknowledged writes



Offloading index data to the object store (stateless)

Node 1 (indexing)



Offloading index data to the object store (stateless)

Node 1 (indexing)

Index "logs.2025.1.17", shard 1

Lucene segments generation: 1
Docs: [1, 12]

Flushed commits: Commit 1

Translog generation:
Docs:

New docs [13, 15]

Lucene in-memory indexing buffers

3
[13, 15]

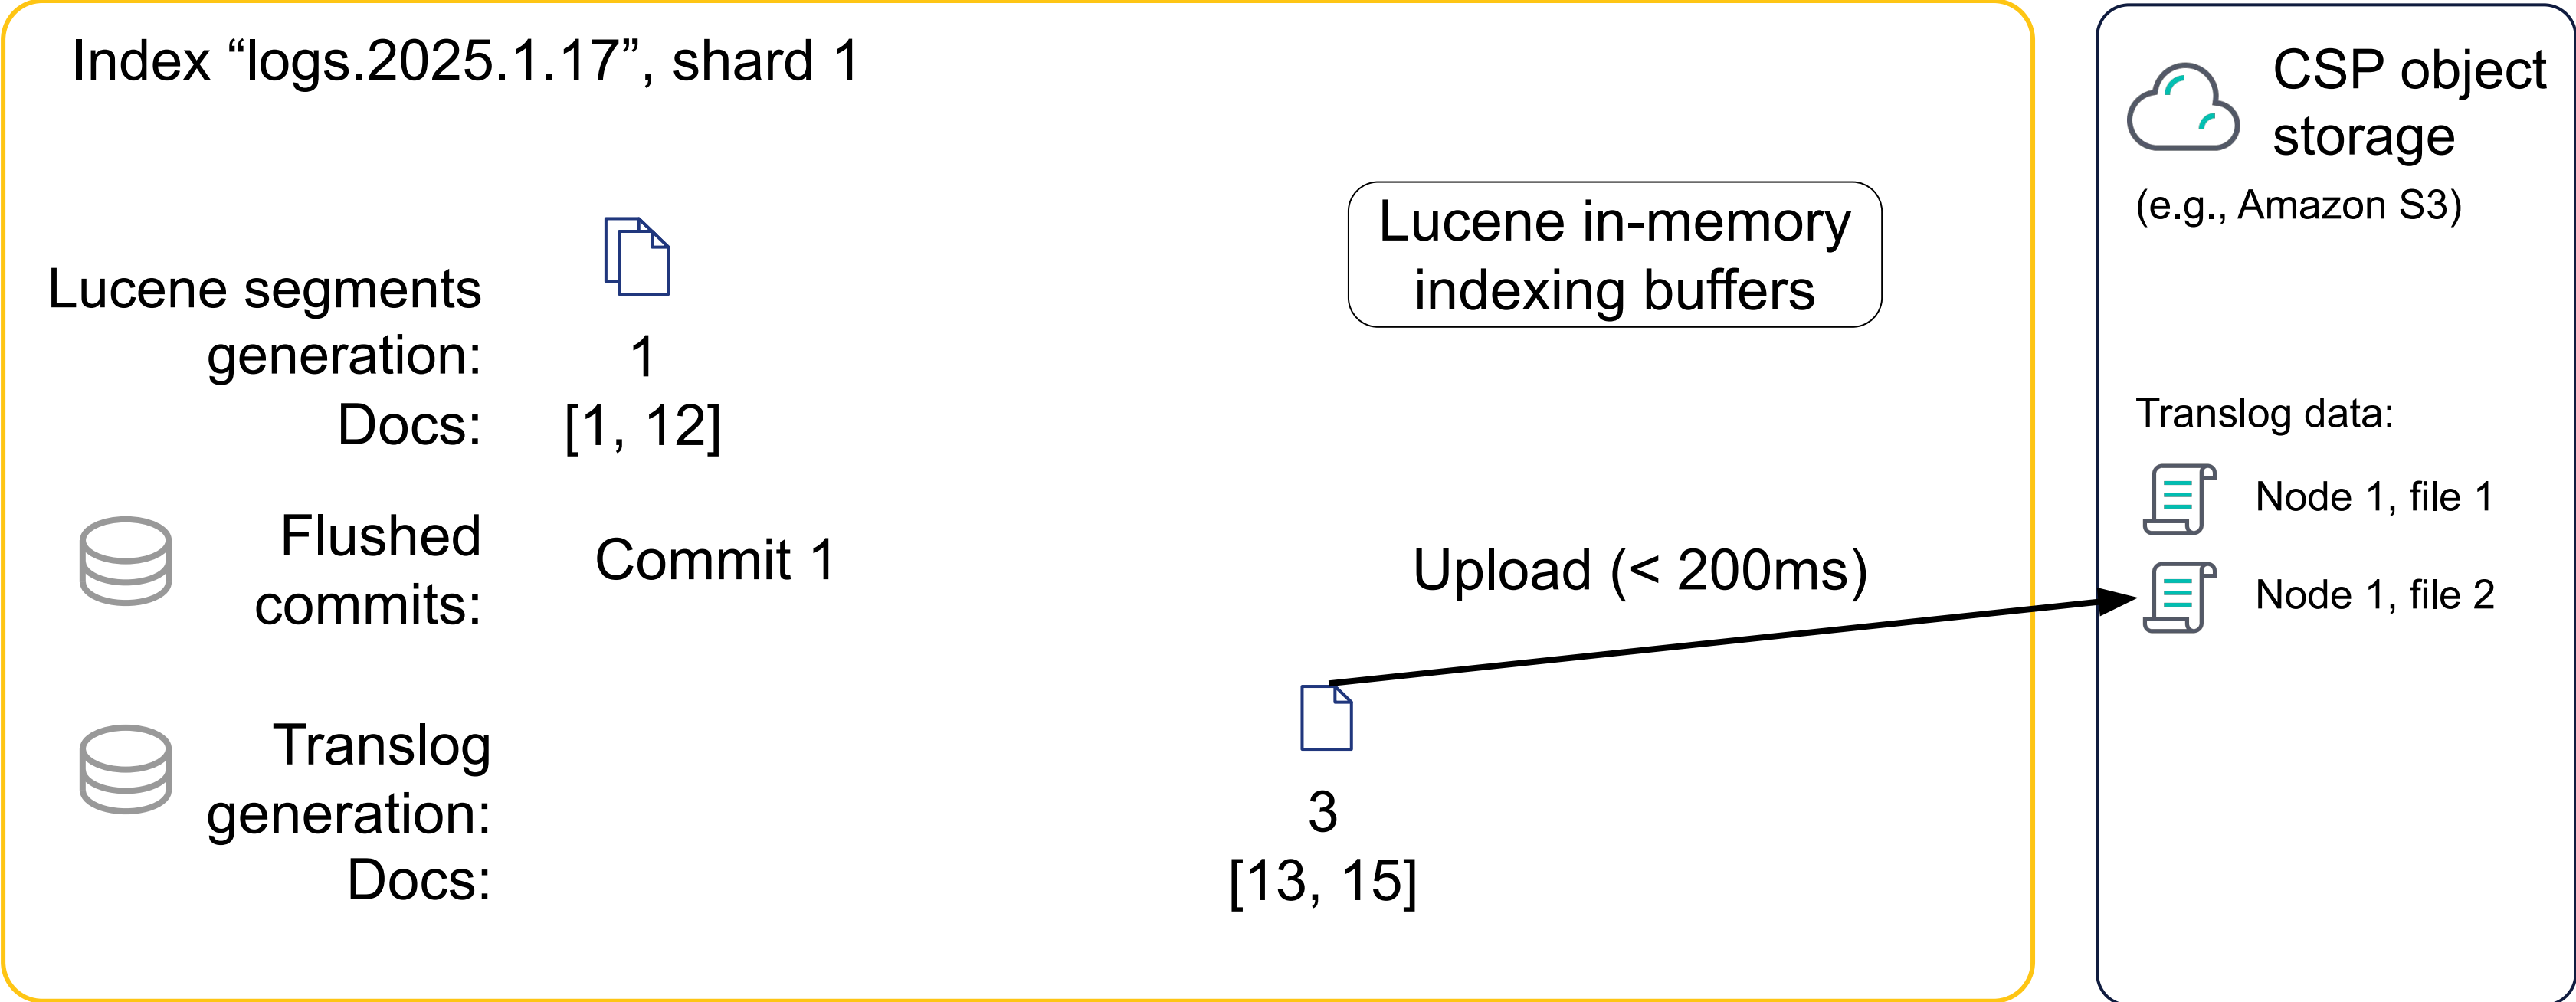
CSP object storage
(e.g., Amazon S3)

Translog data:
Node 1, file 1

Offloading index data to the object store (stateless)

Node 1 (indexing)

New docs [13, 15]



Offloading index data to the object store (stateless)

Node 1 (indexing)

Acknowledged writes

Index "logs.2025.1.17", shard 1

Lucene segments generation: 1
Docs: [1, 12]

Flushed commits: Commit 1

Translog generation: 3
Docs: [13, 15]

Lucene in-memory indexing buffers

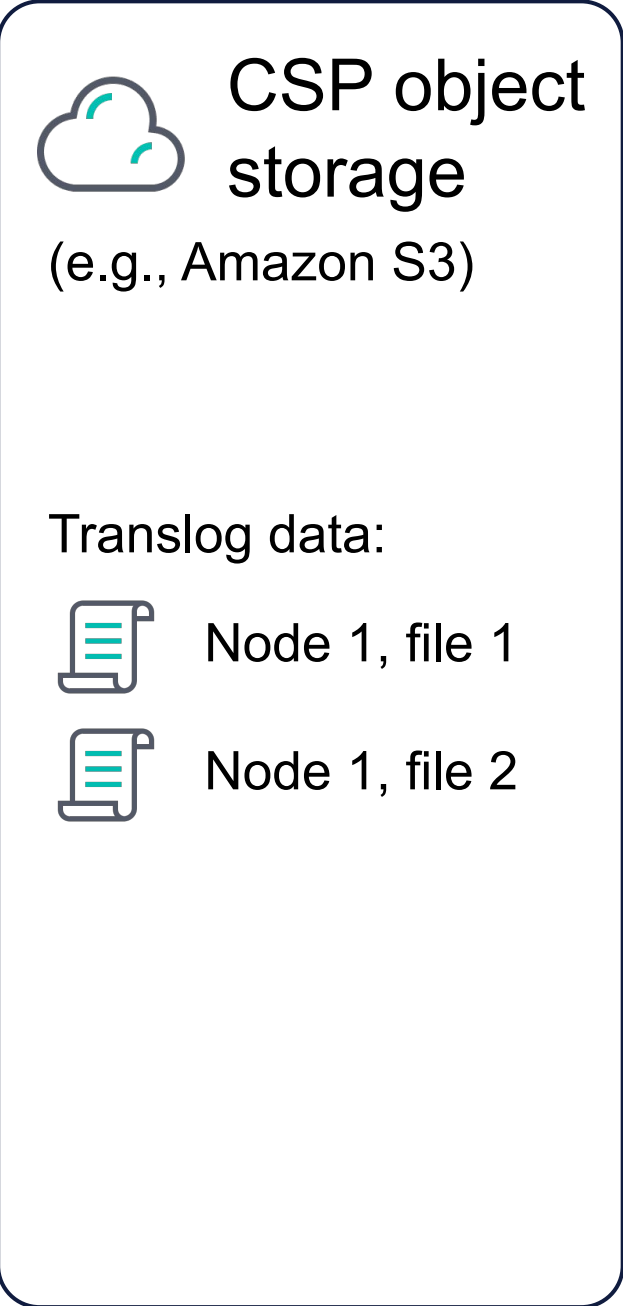
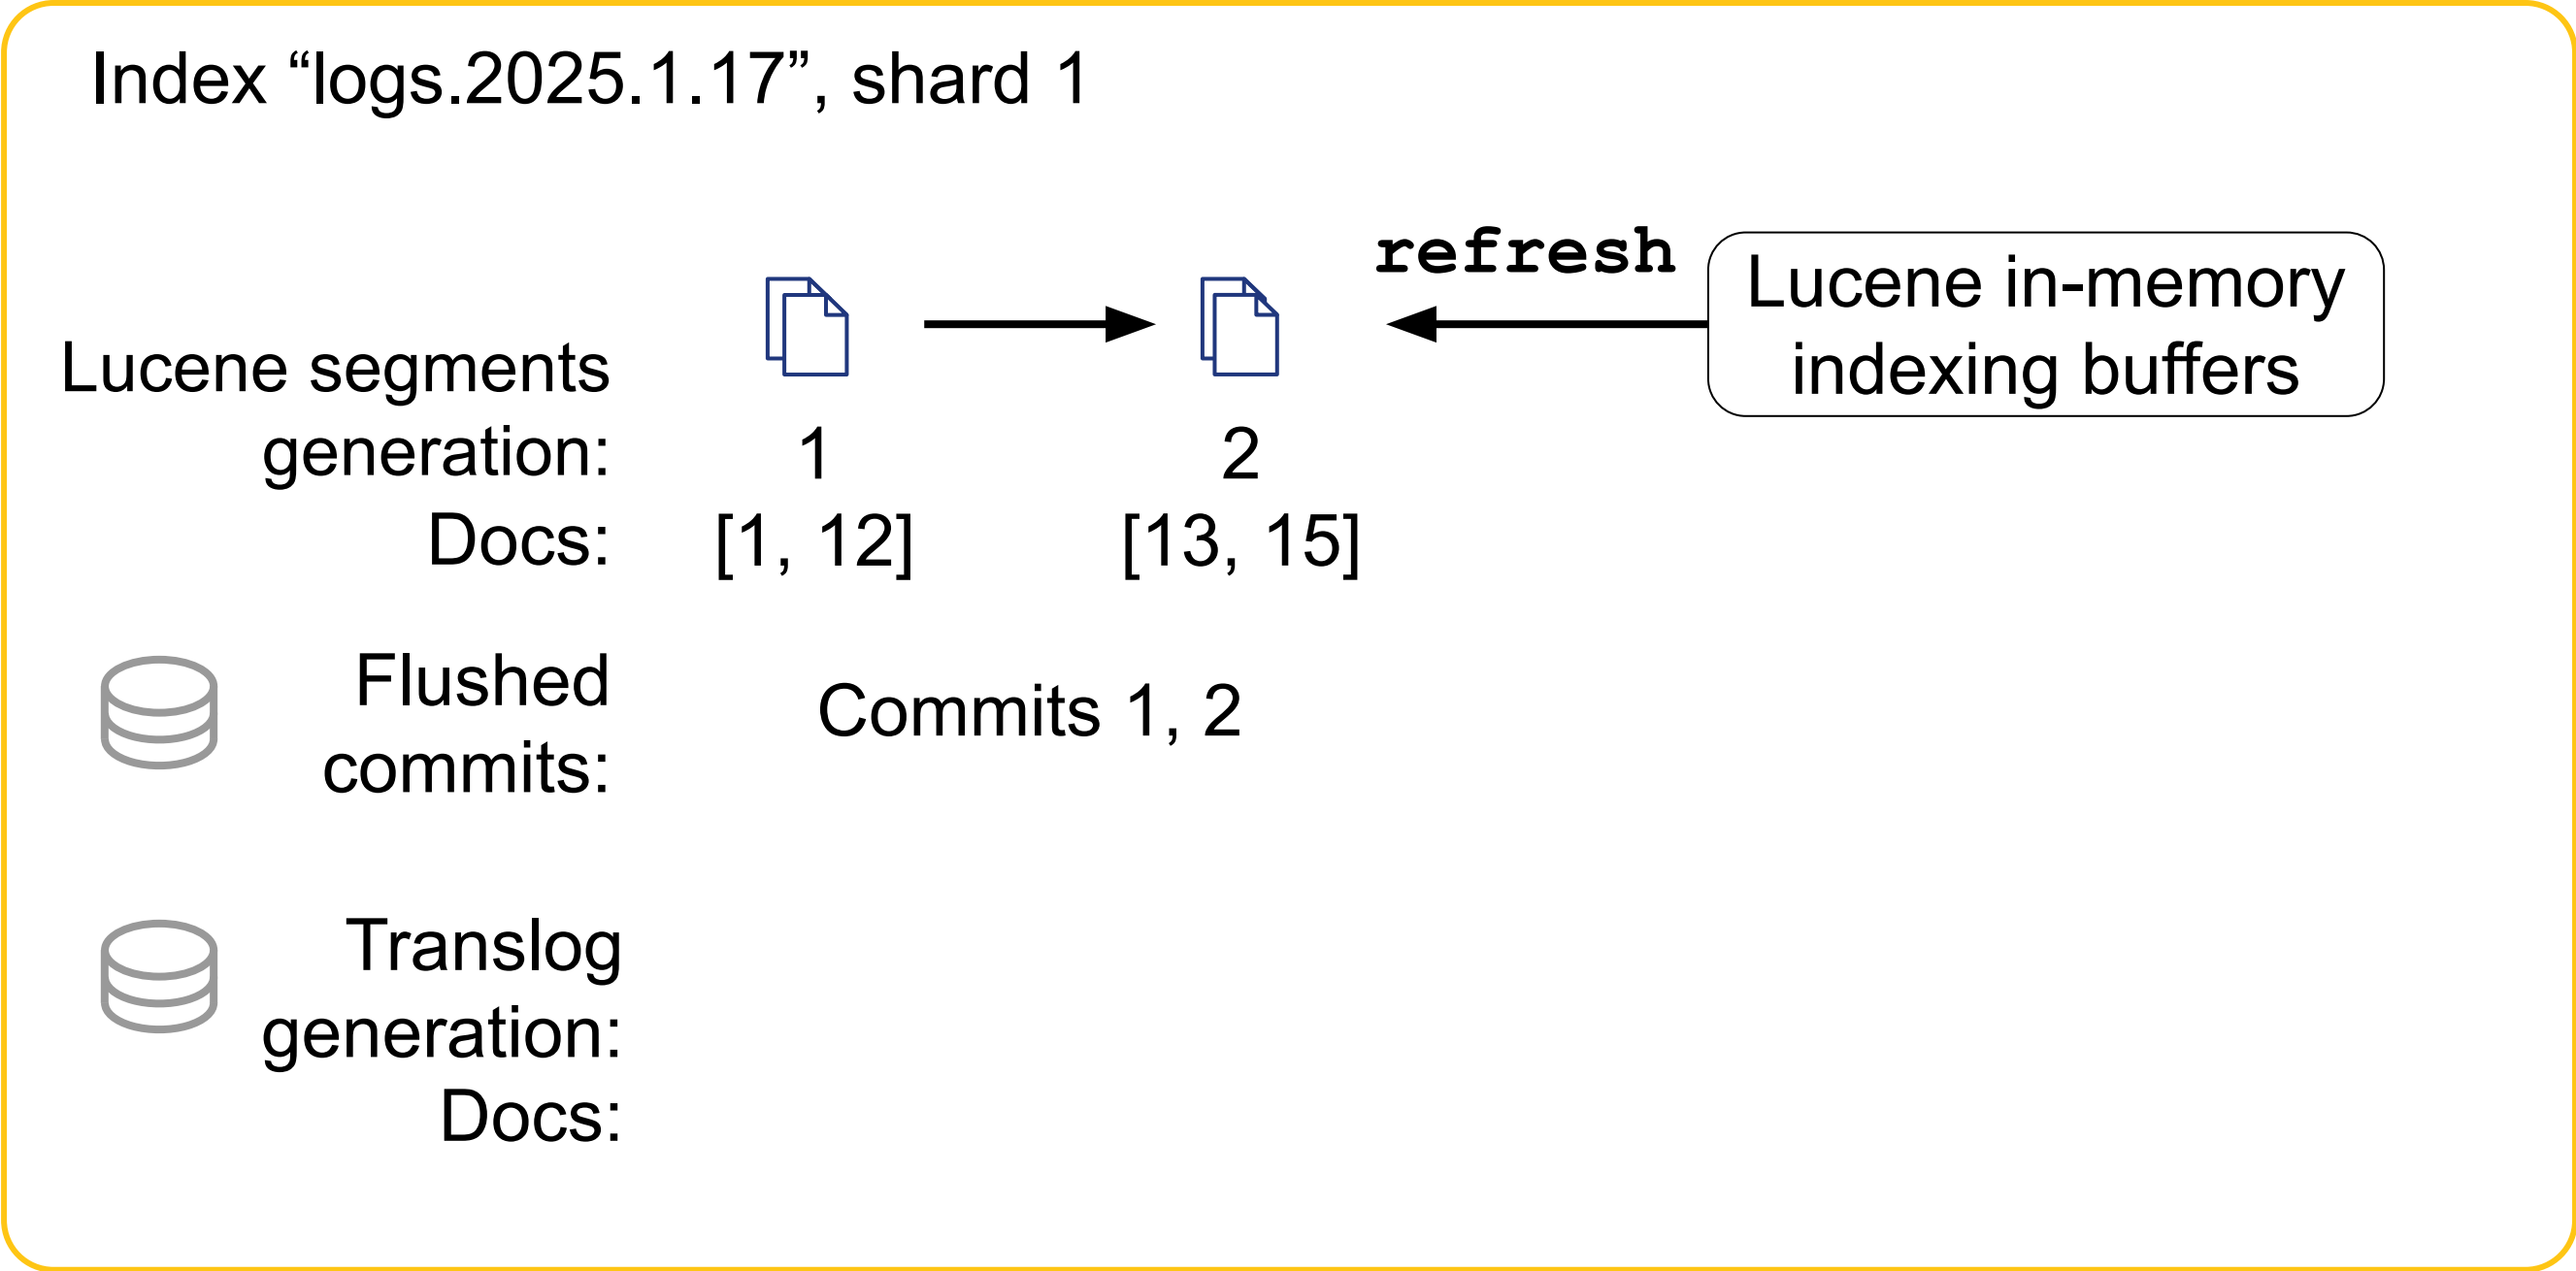
CSP object storage
(e.g., Amazon S3)

Translog data:

- Node 1, file 1
- Node 1, file 2

Offloading index data to the object store (stateless)

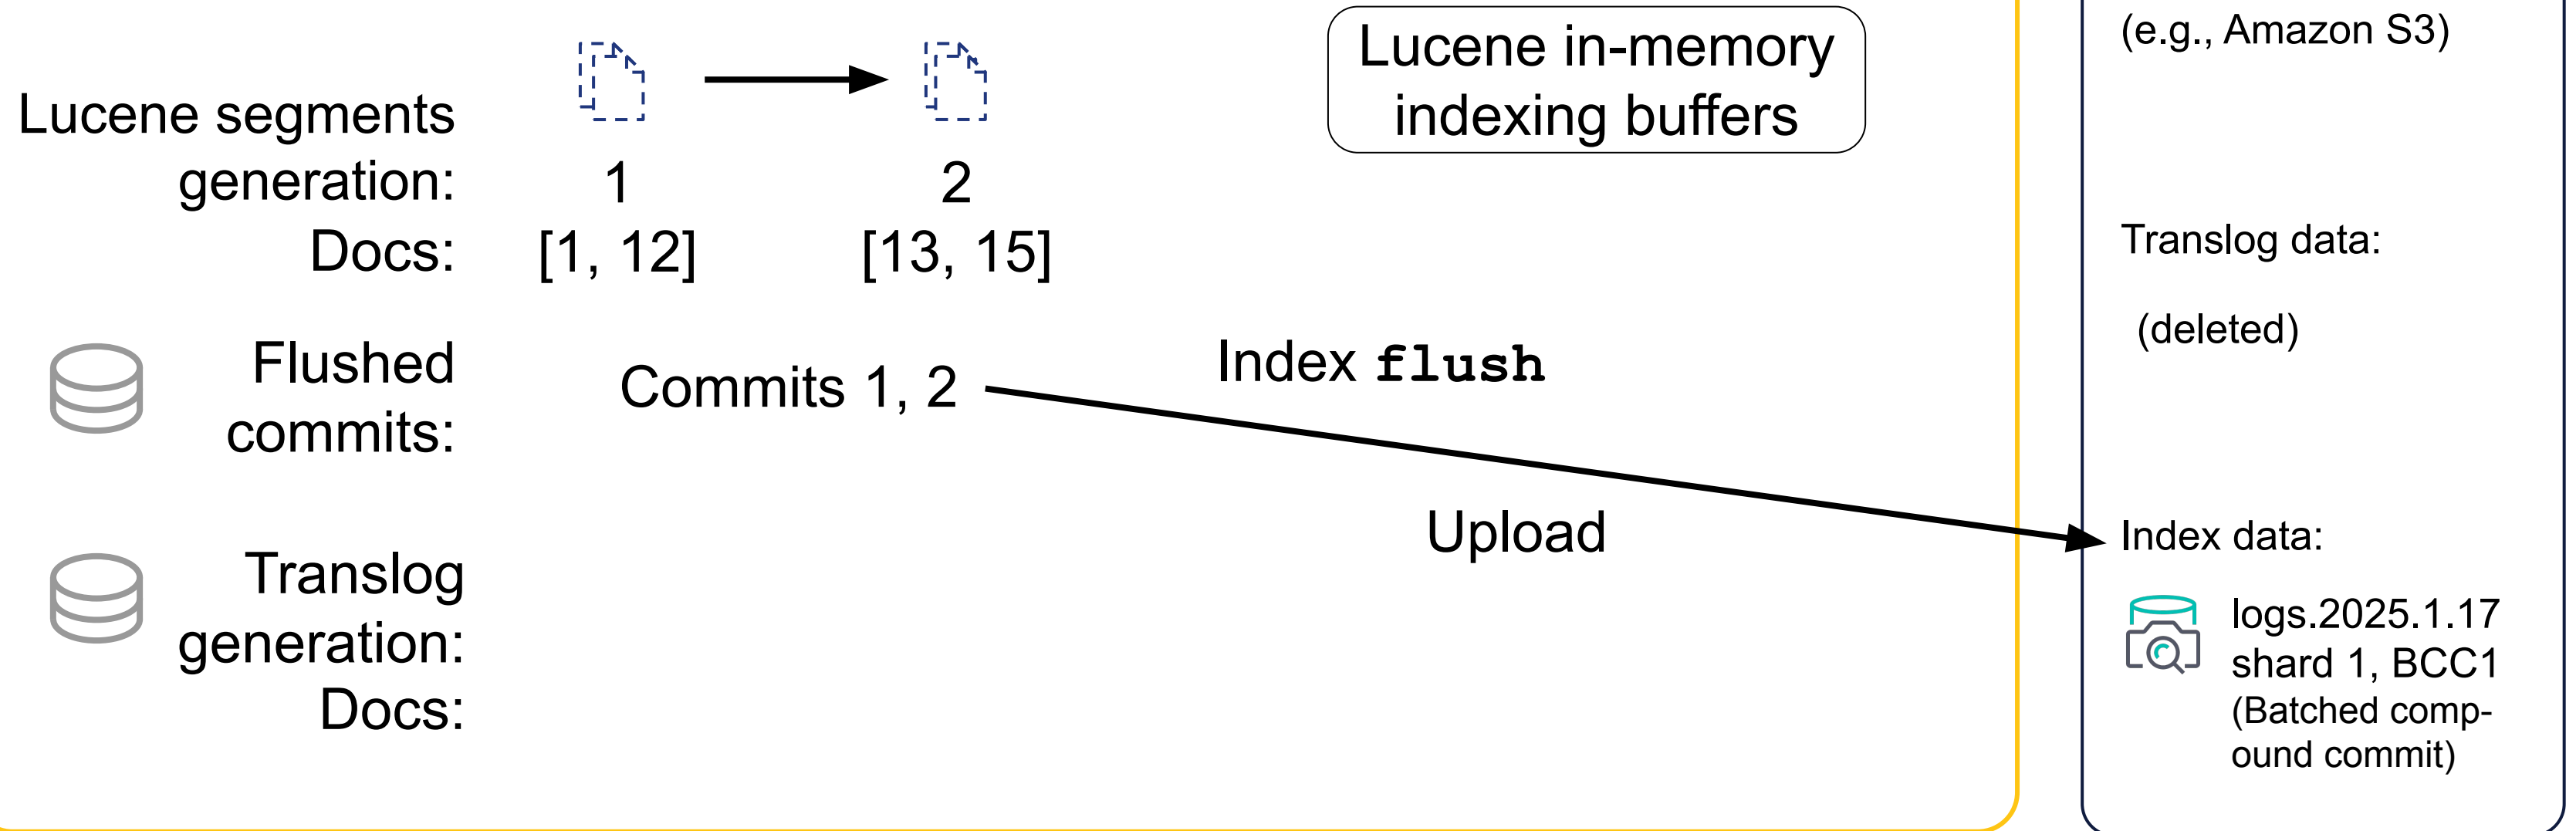
Node 1 (indexing)



Offloading index data to the object store (stateless)

Node 1 (indexing)



Index "logs.2025.1.17", shard 1




Offloading index data to the object store (stateless)

Node 2 (**search**)

Index "logs.2025.1.17", shard 1



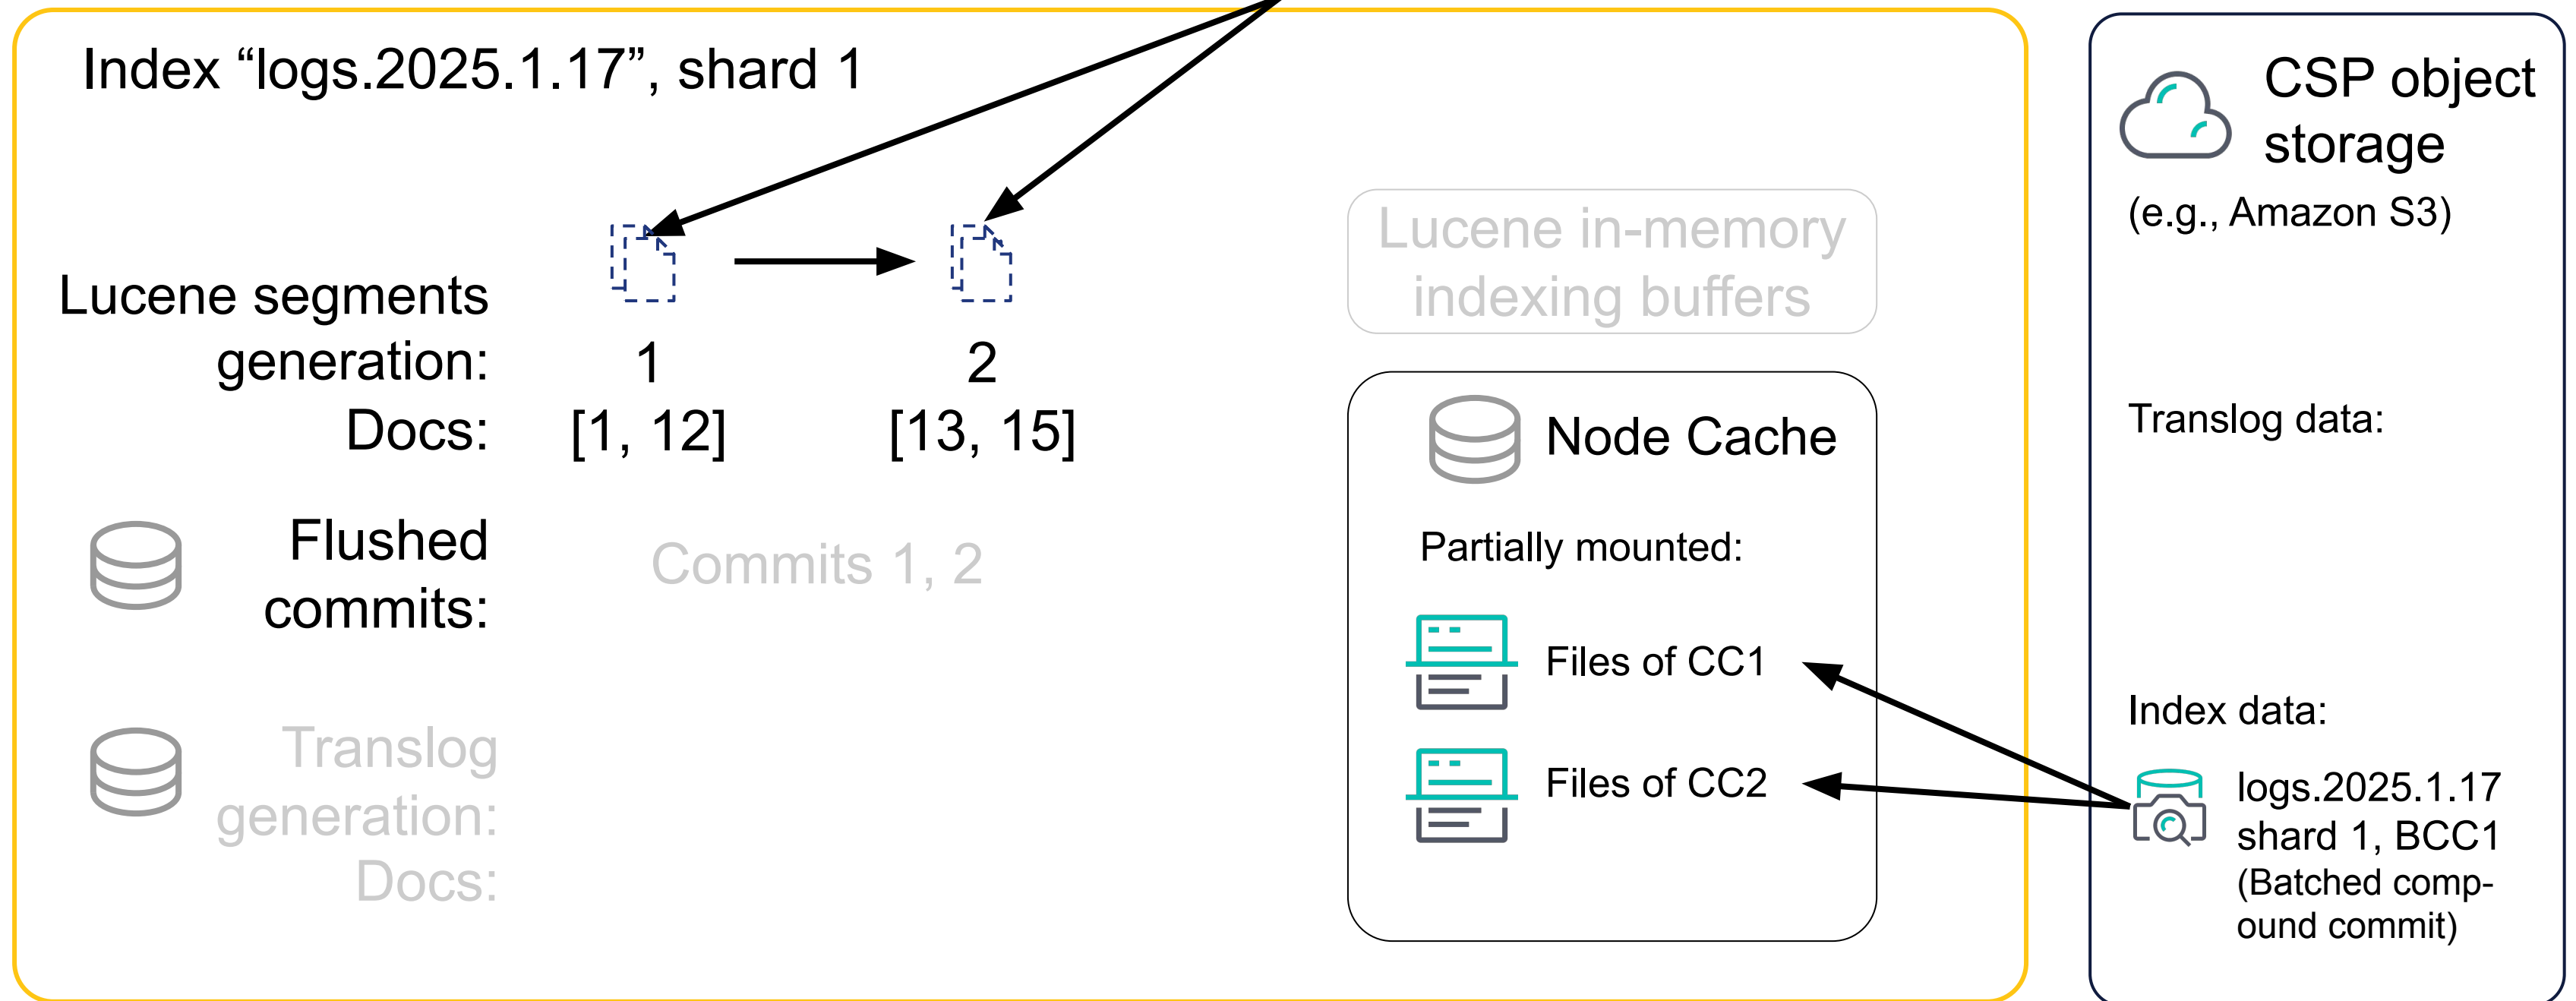
Translog data:



Offloading index data to the object store (stateless)


Node 2 (**search**)

Search documents 2 & 14




Batched Compound Commit

BCC gen=1 blob file:



BCC gen=3 blob file:

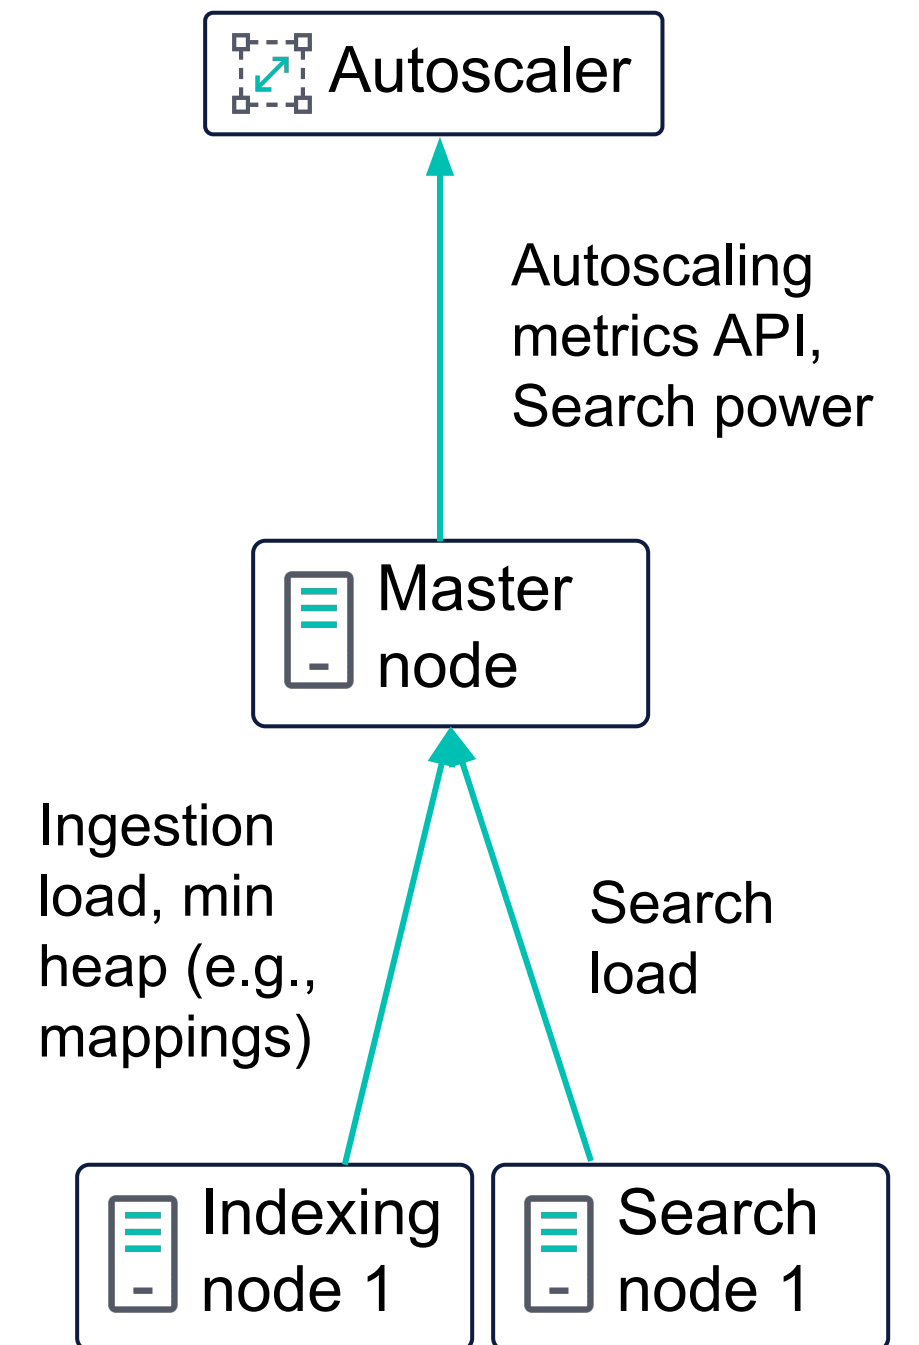


References

- Minimizes PUT/GETs to the CSP to only BCC-level blob
- To search refreshed, but not yet uploaded data, a search node may request data from the last "virtual" (being built) BCC on the indexing node.

Autoscaling

- Autoscaler adjust cluster size up/down to match search and ingestion load
- Nodes report:
 - Estimated minimum heap memory
 - CPU load & queuing, reflects I/O
- Search power (user range setting):
 - #VCUs (fraction of RAM, CPU, disk cache) allocated to boosted data (new time-based docs, and non-time-based docs) → responsiveness
- Autosharding data streams



Summary

- Elasticsearch is an unparalleled scalable search & analytics engine
- ES has been stateful, relying on disk durability
- New **Serverless ES** is stateless; durability on CSP **object store**
 - Showed how we upload BCCs and translogs cost-efficiently
- Just two **tiers, indexing and search** → independent workloads
- Partial caching enables indexing and searching limitless data
- Thin shards that can relocate/recover fast
 - Frequent automatic hassle-free **upgrades**, better security
 - Workload-based **autoscaling**
- Users **just use the same API**
- Pay-as-you-go model for data retention, indexing, searching

Community, culture and careers

- Start a trial at cloud.elastic.co
- Register for an upcoming event at elastic.co/events
- Vibrant community
 - Community portal → www.elastic.co/community
 - Elastic Community on Slack → ela.st/slack
 - Community videos → ela.st/community-youtube
 - Discussion forums → discuss.elastic.co
 - Elastic Contributor Program (e.g., earn training) → elastic.co/community/contributor
 - Community events & groups across the globe → community.elastic.co
 - Newsletter
- Careers → elastic.co/about/careers
 - [Elastic Source Code](#), remote, 2800+ employees across 40+ countries
- Subscribe for next meetup.com/greece-elastic event

Thank You!

Iraklis Psaroudakis
www.kingherc.com

