Road to Auto Scaling

Varun Thacker

Lucidworks

Apache Lucene/Solr Committer, and PMC member

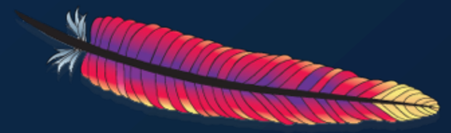

Agenda

- APIs
- Metrics
- Recipes
- Auto-Scale Triggers

SolrCloud Overview

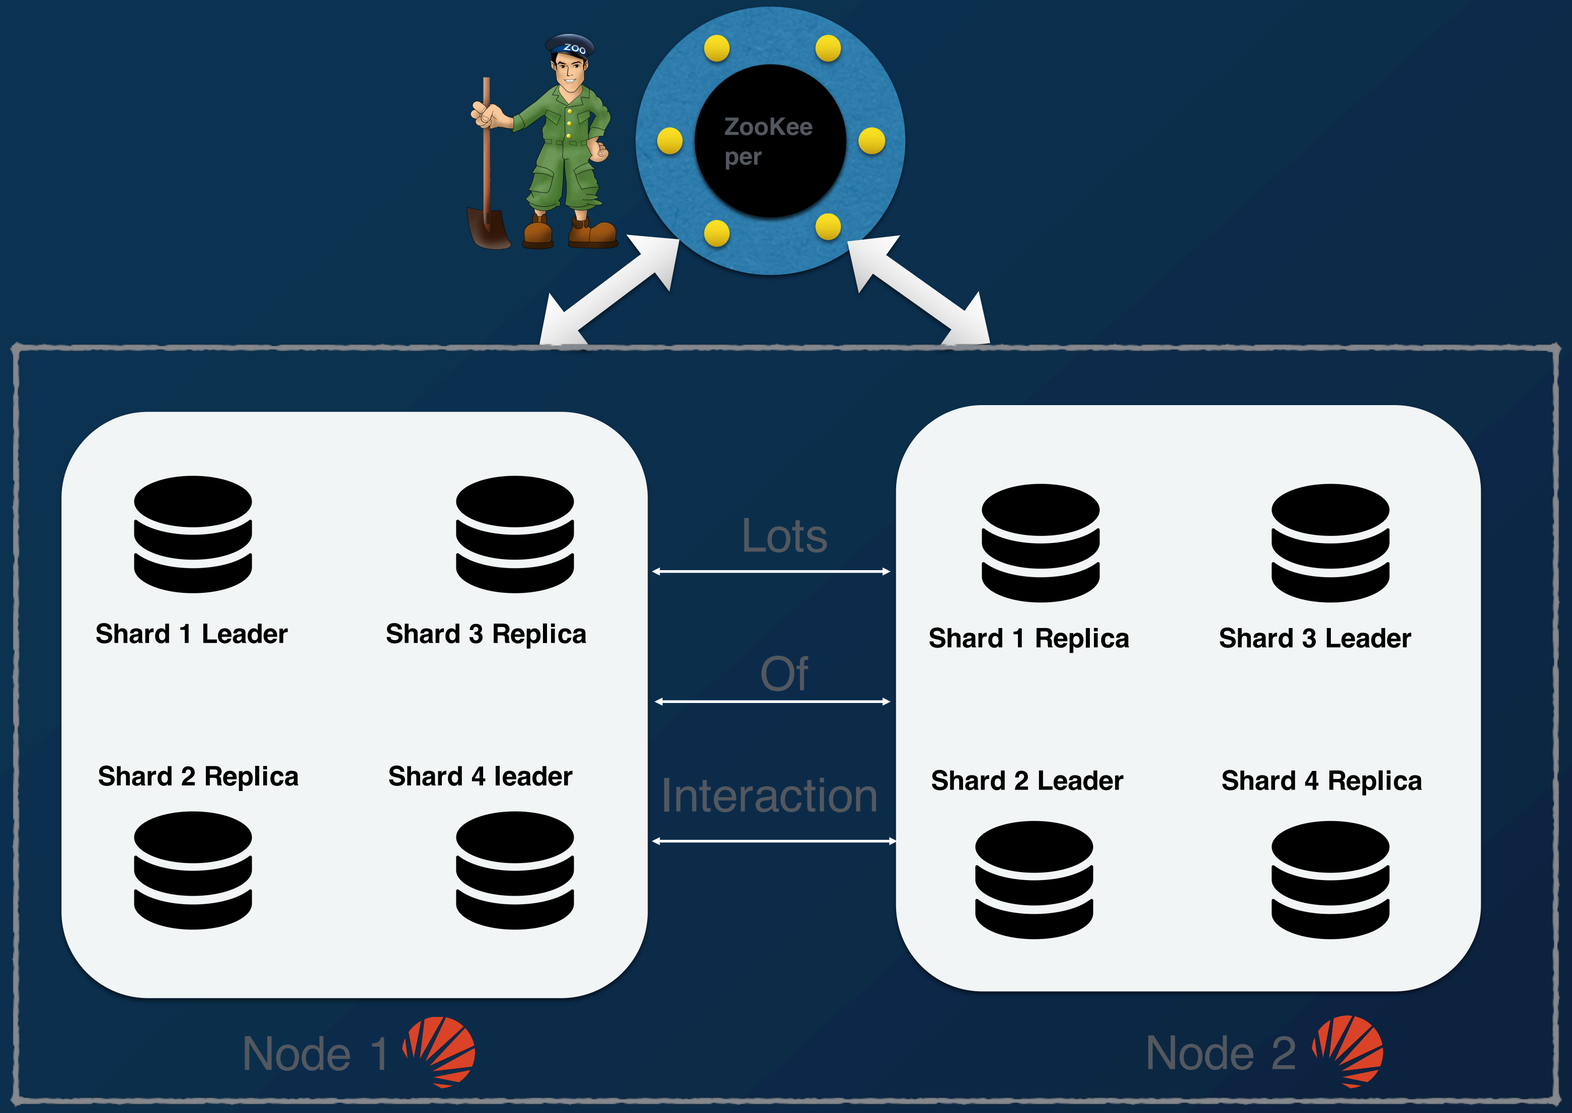

- No master node
- Per shard leader for writes
- All nodes are equal for queries

APIs

- Regularizing REST APIs
- JSON V2 for cluster and node APIs
- Introspection support via JSON
- Schema API
- Config API (solrconfig.xml)

API Examples

- curl -X GET http://localhost:8983/v2/collections/_introspect
- curl -X POST http://localhost:8983/v2/collections -d
  '{ create: { name: bbuzz, config: gettingstarted, numShards: 1, replicationFactor: 1 } }'

Lot’s of APIs

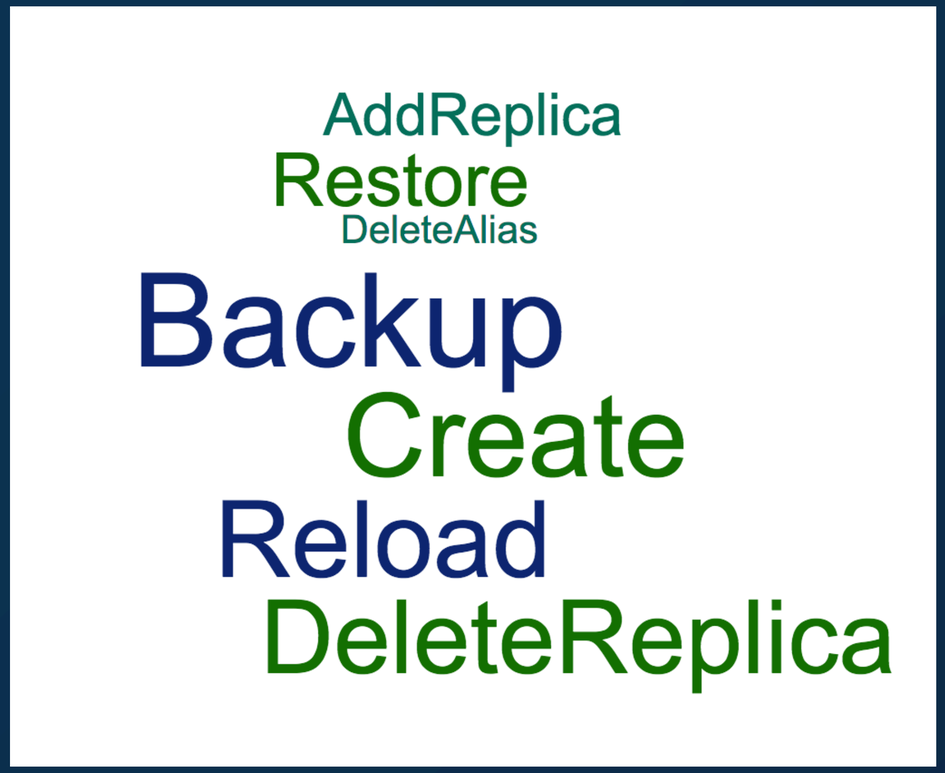

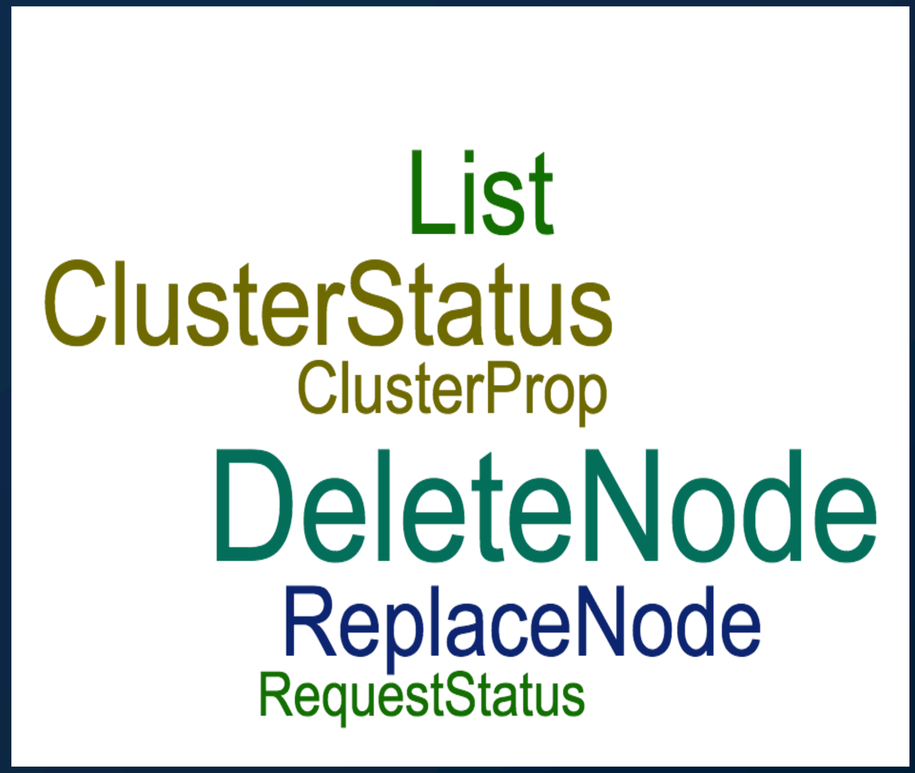

Still need to know when to call the APIs!

Metrics : Dropwizard

- counters
- meters
- histograms
- timers
- gauges

Registries

- Grouped related metrics
- Not persisted across JVM restarts

JVM

Node

Core

Jetty

Registry Summary

- JVM: OS memory, CPU , GC stats, Physical memory stats
- Node: authorization success/failure counts, API request times and counts
- Core: Request Handler metrics (request counts, percentiles) , Indexing stats
- Jetty: Thread pools, HTTP Response count, GET/PUT/DELETE operation counts

Reporters

- Pushes data to external systems
- Solr ships with reporters to ship to : Graphite, File via SLF4J
- Exposes all the metrics via JMX and a REST API

Writing a custom reporter

- Want to push data to your reporting tool?
- Dropwizard comes with reporter library for over 10+ reporter databases: http://metrics.dropwizard.io/3.1.0/manual/third-party/
- Very easy to wrap the library and load it as a solr reporter plugin
- Example for pushing data to influxDB : https://github.com/vthacker/solr-metrics-influxdb
- I used it to visualise data in Grafana in just a few hours!

Solr Replication Modes

SolrCloud Replication

- Solr’s default replication model is designed for consistency. Here’s how it works
- Shard leader accepts a document
- Shard leader writes to its transaction log
- Shard leader writes the document to it’s index , forwards requests to all replicas
- All replicas write to their transaction log then index locally
- Shard Leader waits for all the replicas to return and then acknowledges back to client
- If a replica fails to write the document the leader puts the replica in recovery

SolrCloud Replication

- How can I isolate reads and writes
- With Solr 7 you can create query only replicas

Recipes

Increase Query Throughput

- QUERY./select.requestTimes
- os.systemLoadAverage
- gc.ParNew.time

- Add more replicas and scale out horizontally
- Adding more replicas could also help reduce query latency
- Action: use a pullReplica (query only replica)

Improve Query Latency

- QUERY./select.requestTimes (p95_ms, 5minRate etc)
- QUERY.httpShardHandler.threadPool.httpShardExecutor.running

- Split Shard - With more shards you increase parallelism
- Or simply create a collection with more shards

Improve Indexing Throughput

- os.systemLoadAverage
- gc.ParNew.time
- Maybe nodes go into recovery when you push it too hard during indexing
- Merge 15 ^{min} mean rate

- More shards spread across more nodes
- Reduce replication during bulk index to quickly ingest data
- Use a combination of different replication models for your replicas

Autoscaling

Autoscaling

- A new set of features to help users manage and scale their clusters
- Replaces Solr’s existing Replica Placement Strategy with a more generic policy engine
- Events create actionable triggers

Policy

- A set of rules at the cluster and collection level
- Don't create any replicas on the overseer node : {'nodeRole':'!overseer', 'replica':0}
- Don't have more than 2 replicas on any node : {'replica':'<2', 'shard': '#EACH', 'node': '#ANY'}
- You can define them in Solr 7 and the Collection APIs (Create Collection, Add Replica etc) will use it to balance a cluster

Events/Triggers

- Cluster Events: Events like nodeAdded, nodeLost, searchRate
- Triggers performs an action
- It can compute a plan or print out for the user to execute
- Work still in progress. Will likely come out in 7.1
- Solr Jira https://issues.apache.org/jira/browse/SOLR-9735 to follow the discussion and see the design document

What’s next

- Lucene/Solr 7 will release in the next couple of months
- Lucene/Solr Revolution in September and will have talks on Metrics, Replication Modes and Autoscaling in greater detail over multiple talks.

Lucidworks

Thank You


Lucidworks