

Cassandra At Bloomberg: Backing Up & Restoring Clusters with Medusa

Community Over Code NA 2024
October 7, 2024

Isaac Reath, Engineering Team Lead, NoSQL Infrastructure
Lindsey Zurovchak, Software Engineer

TechAtBloomberg.com

© 2024 Bloomberg Finance L.P. All rights reserved.

Engineering

Bloomberg

WHO ARE WE?



350,000+ SUBSCRIBERS



25,000+ EMPLOYEES IN 70+ COUNTRIES



'OPEN SOURCE FIRST' CULTURE



9,000+ SOFTWARE ENGINEERS

TechAtBloomberg.com

© 2024 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

NOSQL INFRASTRUCTURE



OFFERS APACHE CASSANDRA-AS-A-SERVICE FOR BLOOMBERG'S ENGINEERS



PART OF A LARGER MANAGED DATA SERVICES ORG



STARTED IN 2017

MANAGED CASSANDRA PLATFORM STATS

~300

CLUSTERS

~3PB

DATA SIZE

~4,000

NODES

WHY BACKUP CASSANDRA?

OPERATIONAL RESILIENCY

 APPLICATION LEVEL DATA LOSS

 INFRASTRUCTURE ISSUES

 DATA LIQUIDITY

BUSINESS REQUIREMENTS

 AUDIT REQUIREMENTS

 REGULATORY COMPLIANCE

 BUSINESS CONTINUITY

WHAT DOES CASSANDRA OFFER?

BACKUPS

```
nodetool snapshot -ttl 10d
```

```
cat /etc/cassandra/cassandra.yaml | grep "auto_snapshot:"  
auto_snapshot: true
```

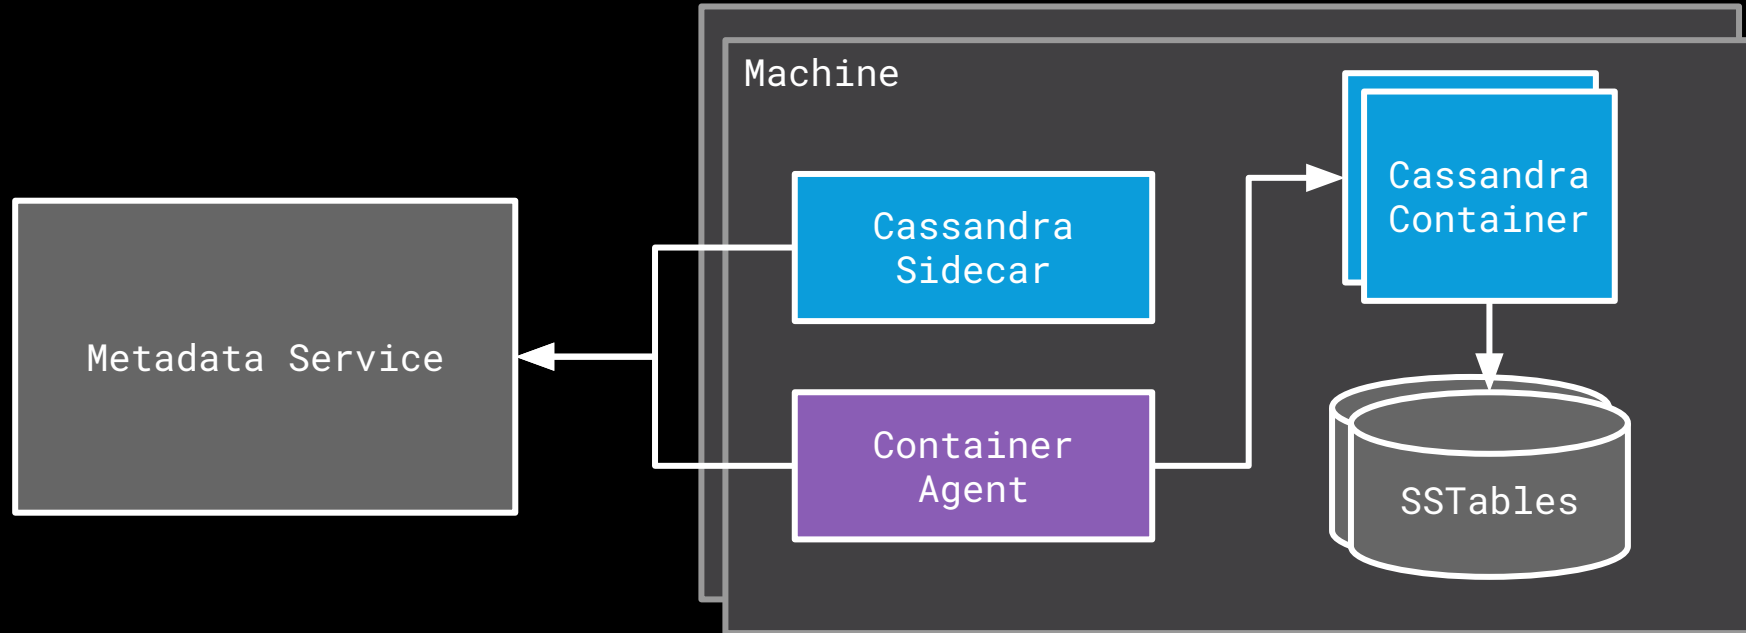
```
cat /etc/cassandra/cassandra.yaml | grep "incremental_backups"  
incremental_backups: true
```

RESTORES

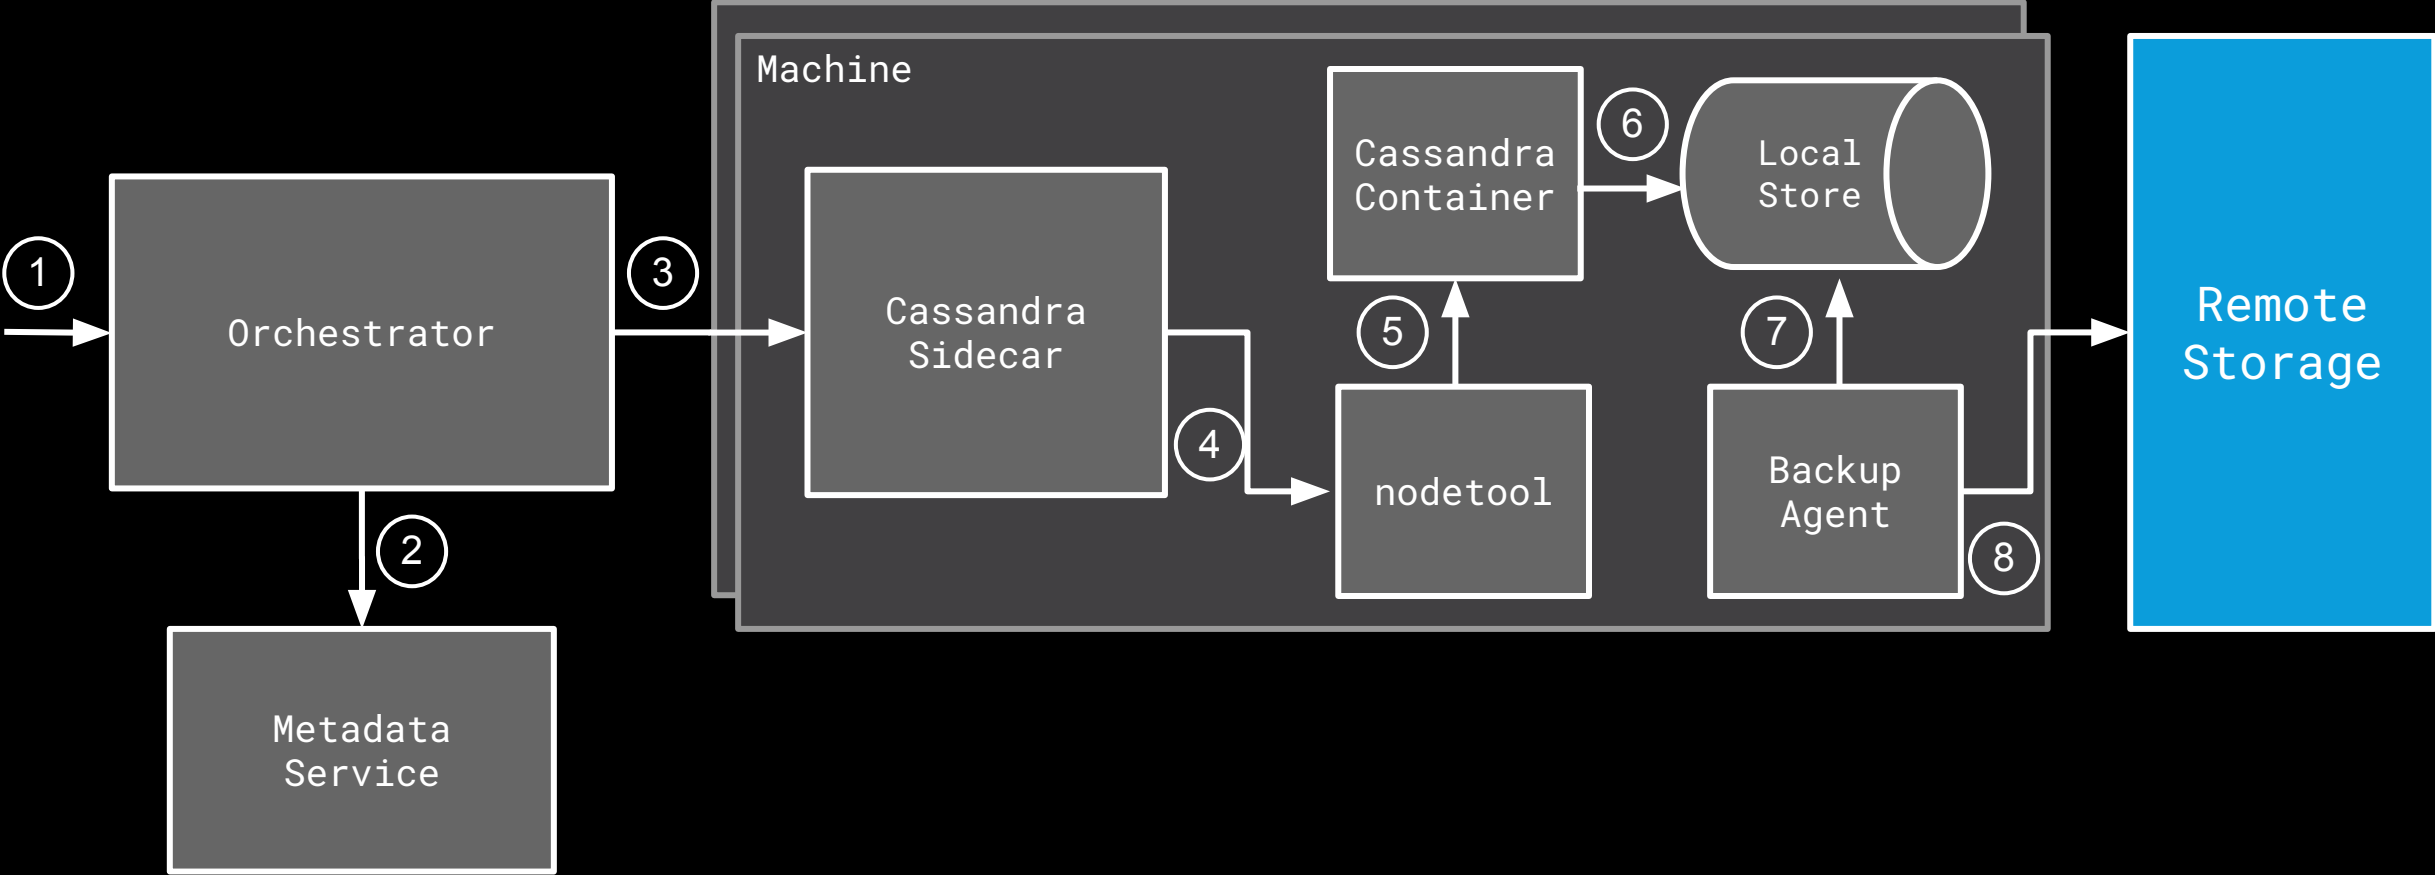
```
nodetool refresh - keyspace1 table1
```

```
sstableloader --nodes 172.17.0.2 \  
/var/lib/cassandra/loadme/keyspace1/table1/
```

APACHE CASSANDRA – PLATFORM ARCHITECTURE



PREVIOUS IMPLEMENTATIONS



BACKUP OPTIONS



UTILIZE EXISTING CASSANDRA TOOLS AND BUILD REMOTE STORAGE CONNECTION



ADD REMOTE STORAGE BACKUP FUNCTIONALITY TO CASSANDRA/CASSANDRA SIDECAR



USE MEDUSA FOR CASSANDRA BACKUPS

WHY MEDUSA?

BENEFITS

BUILT FOR CASSANDRA

SUPPORTS MULTIPLE STORAGE TYPES

SUPPORTS DIFFERENTIAL BACKUPS

HAS PURGE FUNCTIONALITY

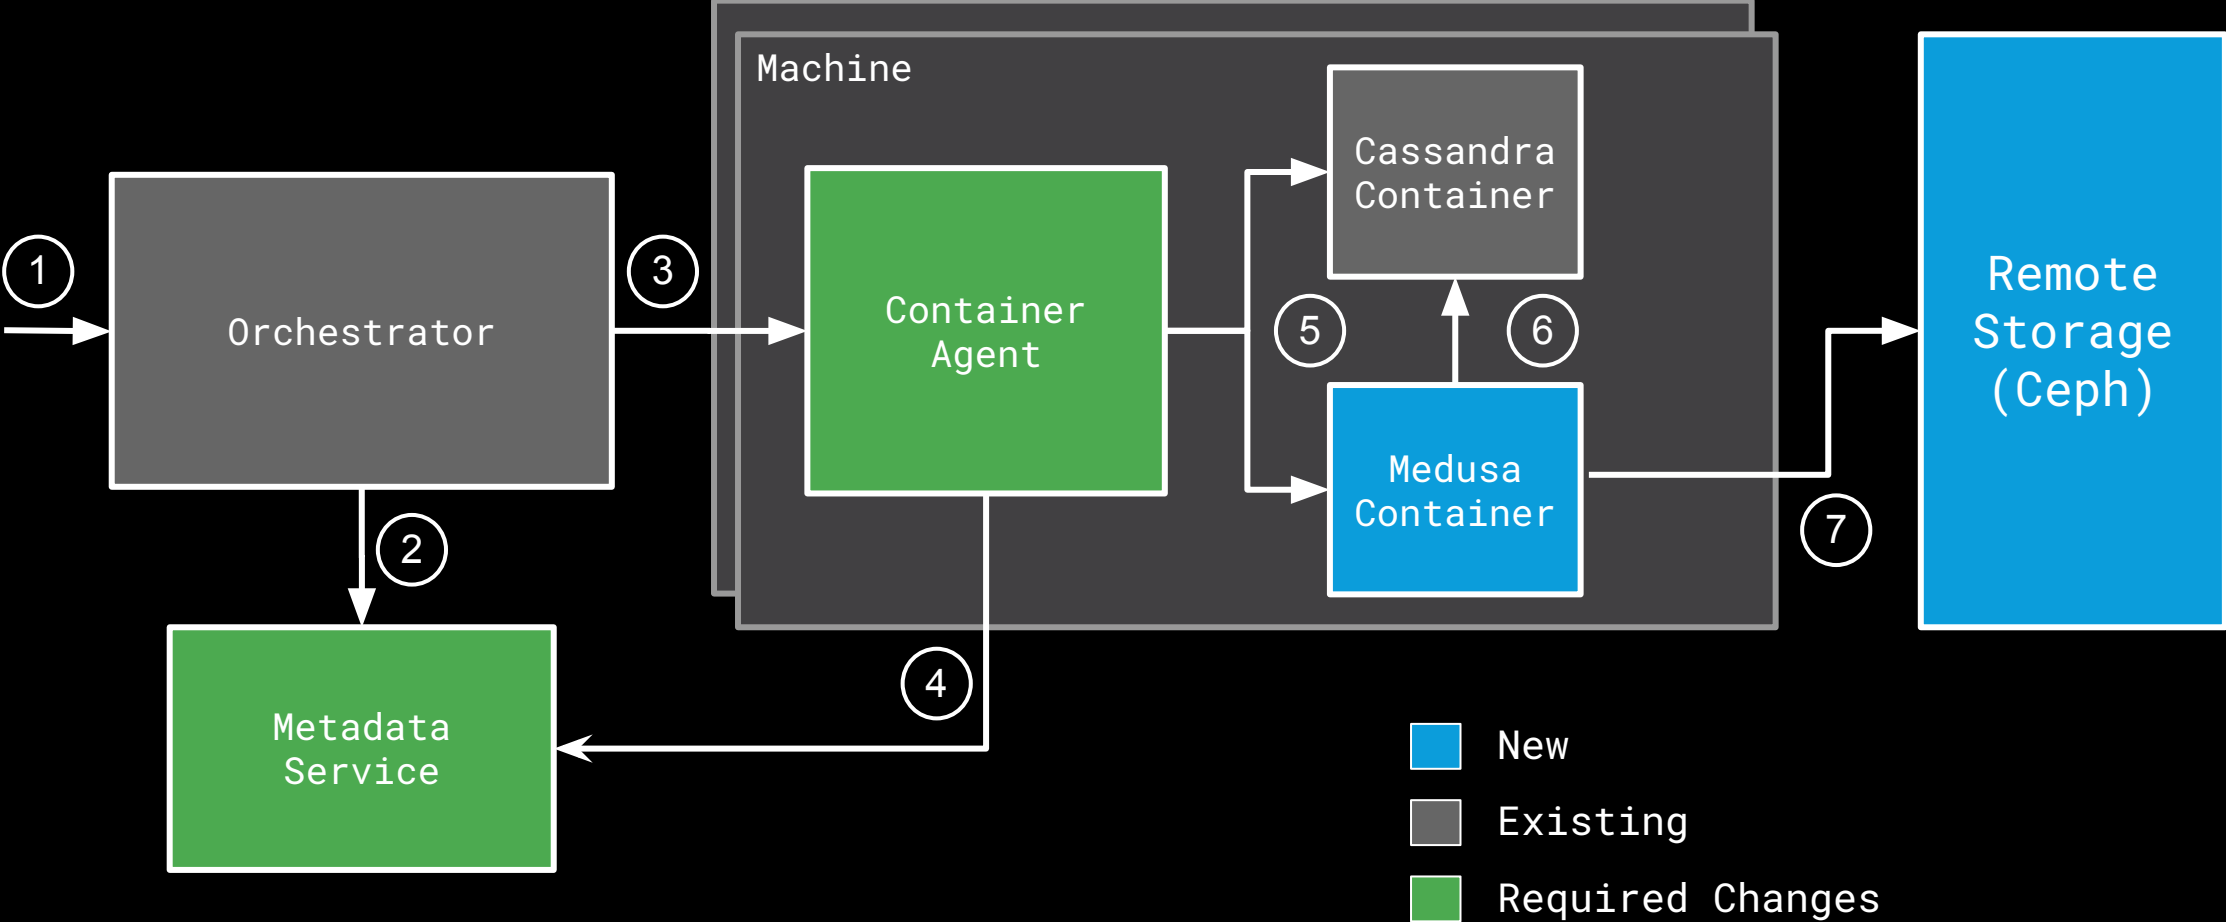
SUPPORTS NODE AND CLUSTER RESTORES

LIMITATIONS

ONLY AVAILABLE AS CLI OR GRPC SERVICE (K8)

ORCHESTRATOR CONNECTS TO NODES THROUGH PSSH

MEDUSA ARCHITECTURE AT BLOOMBERG



BACKUPS WITH MEDUSA

```
FROM python3

RUN pip install cassandra-medusa

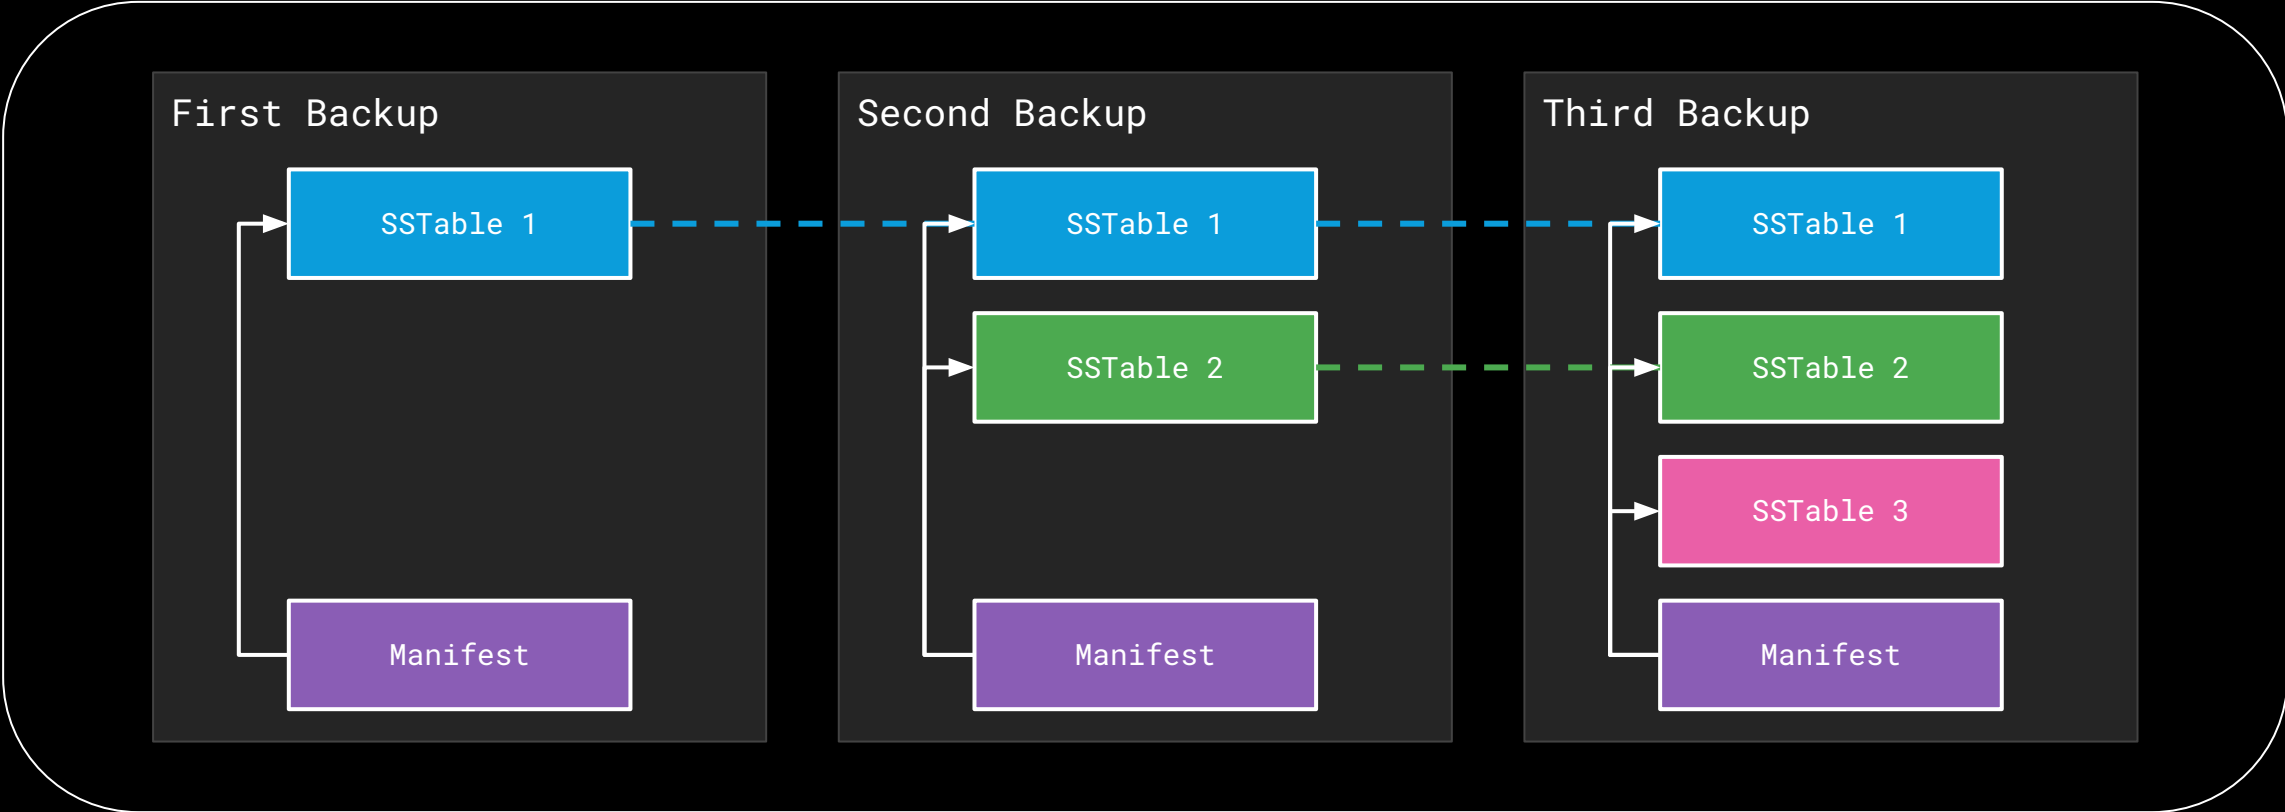
// environment variables used in config.ini file
ENV cluster_name
ENV nodetool_port

VOLUME cassandra.yaml .

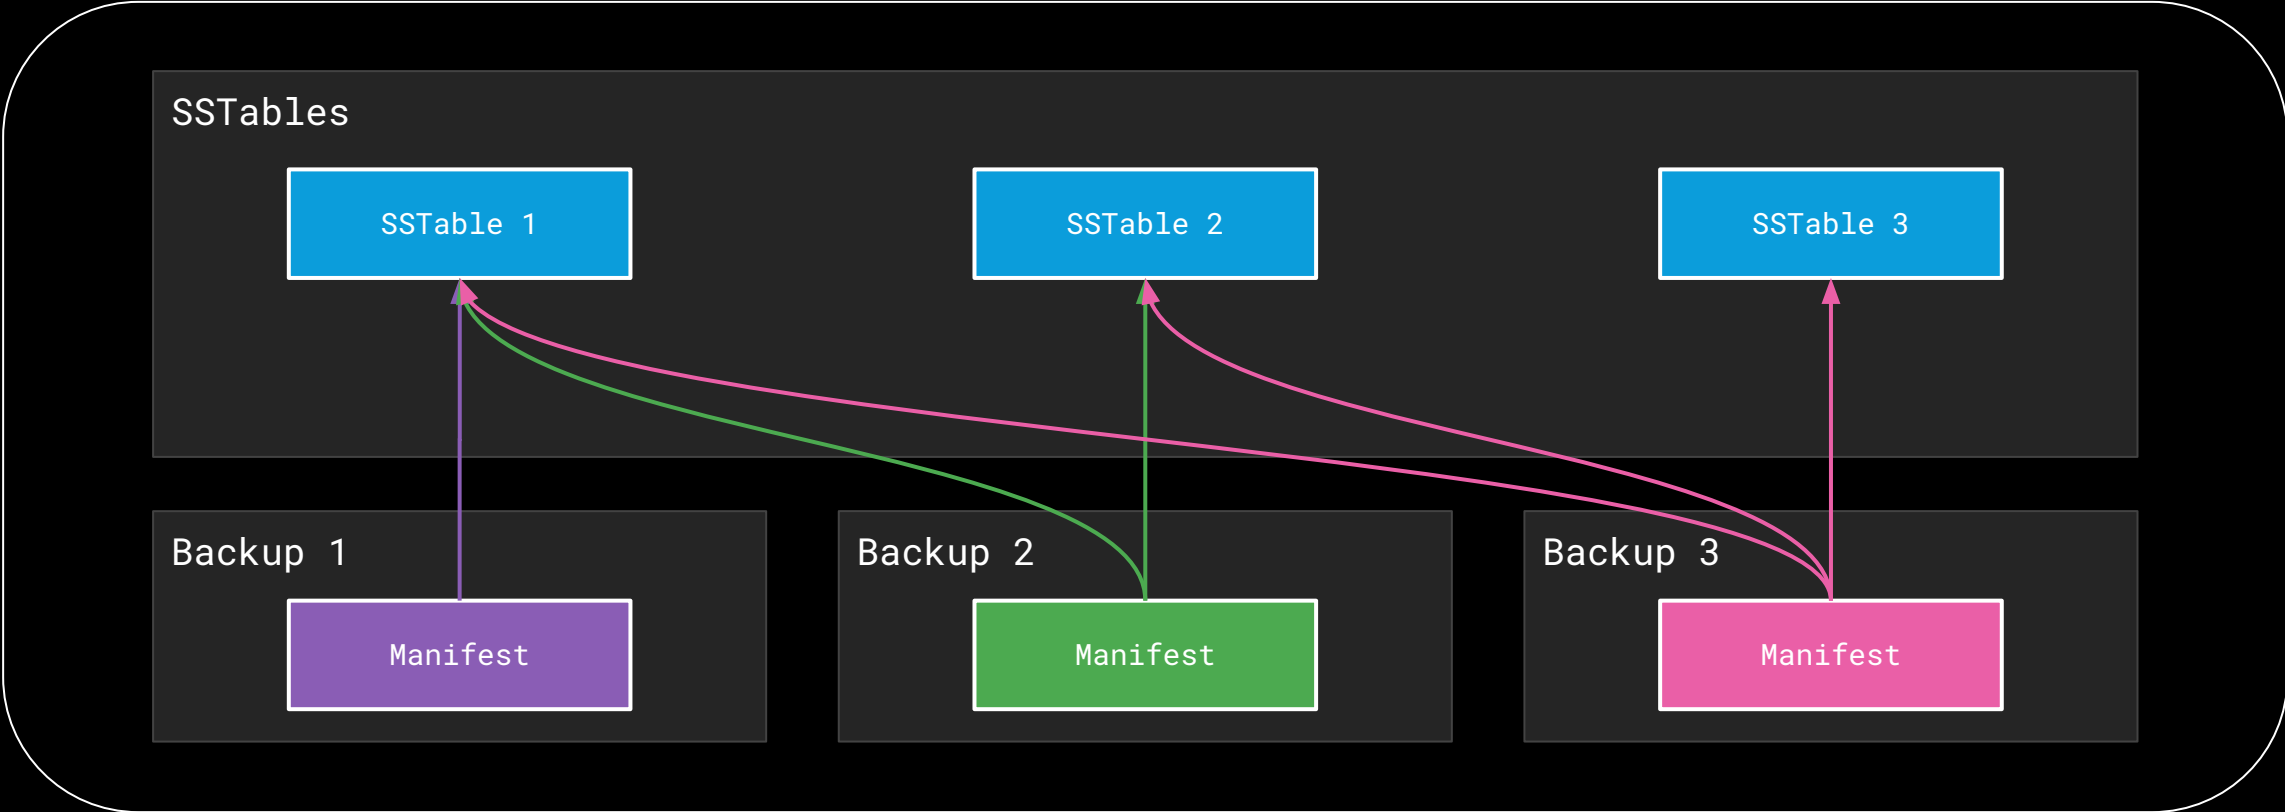
// creates config.ini file
ENTRYPOINT ./config_file_setup

CMD medusa backup --backup-name <name>
```

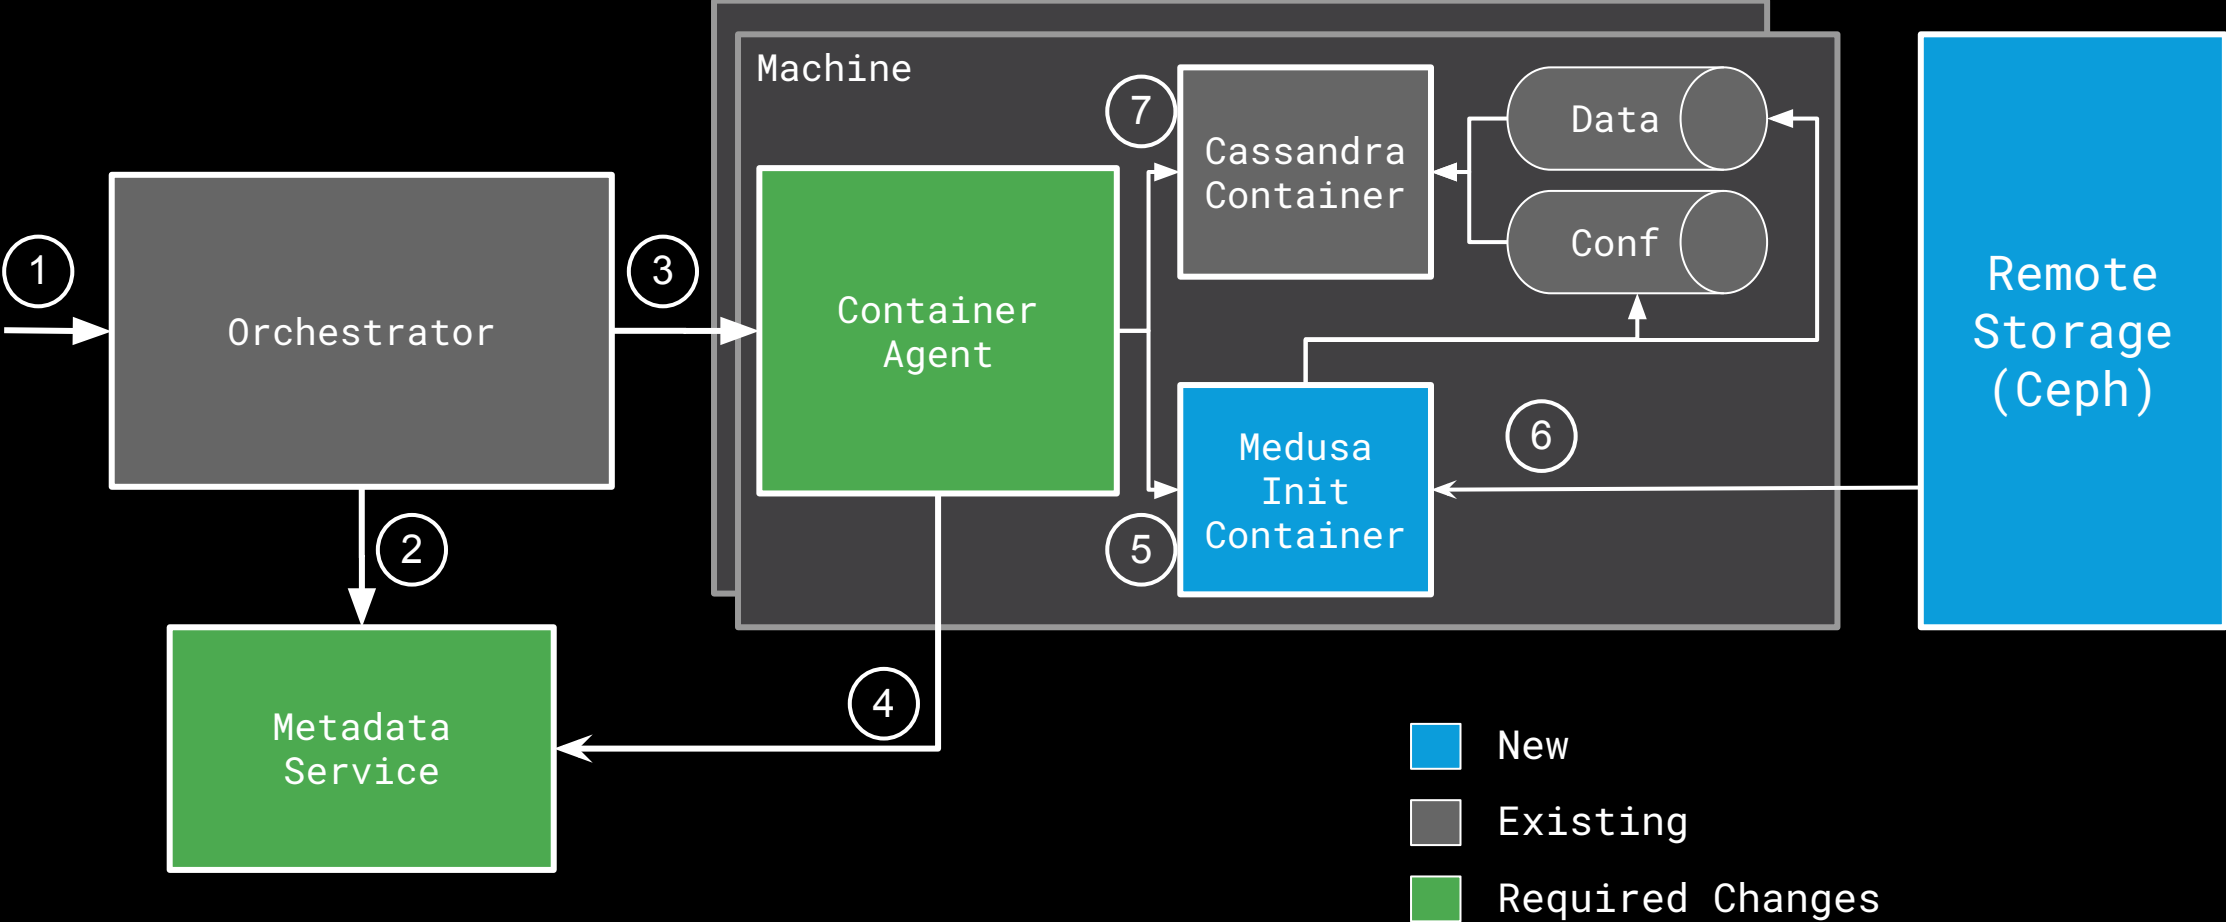
FULL BACKUPS WITH MEDUSA



DIFFERENTIAL BACKUPS WITH MEDUSA



RESTORES WITH MEDUSA



MAINTAINING BACKUPS



RETENTION POLICY

- Needed to ensure backup cost is minimized
- Cannot use bucket TTL
- Nightly execution of Medusa purge



BACKUP DATA INTEGRITY

- Ensure that you can actually restore your backups
- Today: Regularly downloading backups
- Future: Periodically performing restores

FUTURE WORK

MEDUSA IMPROVEMENTS

POINT-IN-TIME RESTORES

IMPROVE STORAGE EFFICIENCY

NATIVE CASSANDRA SOLUTION?

QUESTIONS?

