Tungsten Replicator Master Class

Basics: Tungsten Replicator – How we Move Your Data

Presented by Chris Parker

VP of Customer Success, EMEA

Topics

In this short course, we will

- Review the capabilities of Tungsten Replicator
- Take a look at the inner workings of the Replicator
- Review Topologies and Targets
- Explore Filtering and the Replicator Stages

The Replicator

Tungsten Replicator

- Core to Cluster but can also be used as a standalone Asynchronous Replicator
- Can extract/apply from/to all flavours of MySQL
  - Percona, MariaDB, Enterprise and Community
  - Amazon RDS / Aurora (Including Cross Region)
  - Google Cloud Services
  - Microsoft Azure
- Can replicate between different cloud environments, e.g. from Amazon RDS to Google Cloud
- Replicate to many heterogeneous targets
- Advanced filtering at all stages
- Advanced topologies such as Fan-In and Fan-Out

Supported Heterogeneous Targets

- JDBC Applier
  - PostgreSQL
  - Oracle
- Batch Applier
  - Amazon Redshift
  - HP Vertica
  - Hadoop
- Native Applier
  - MongoDB (Incl. Atlas)
  - Kafka

How It Works…

Tungsten Replicator for JDBC Targets

Option 1: Local Install
Extractor reads directly from the logs, even when the DBMS service is down. This is the default.

Option 2: Remote
Extractor gets log data via MySQL Replication protocols (which requires the DBMS service to be online). This is how we handle Amazon Aurora extraction tasks.

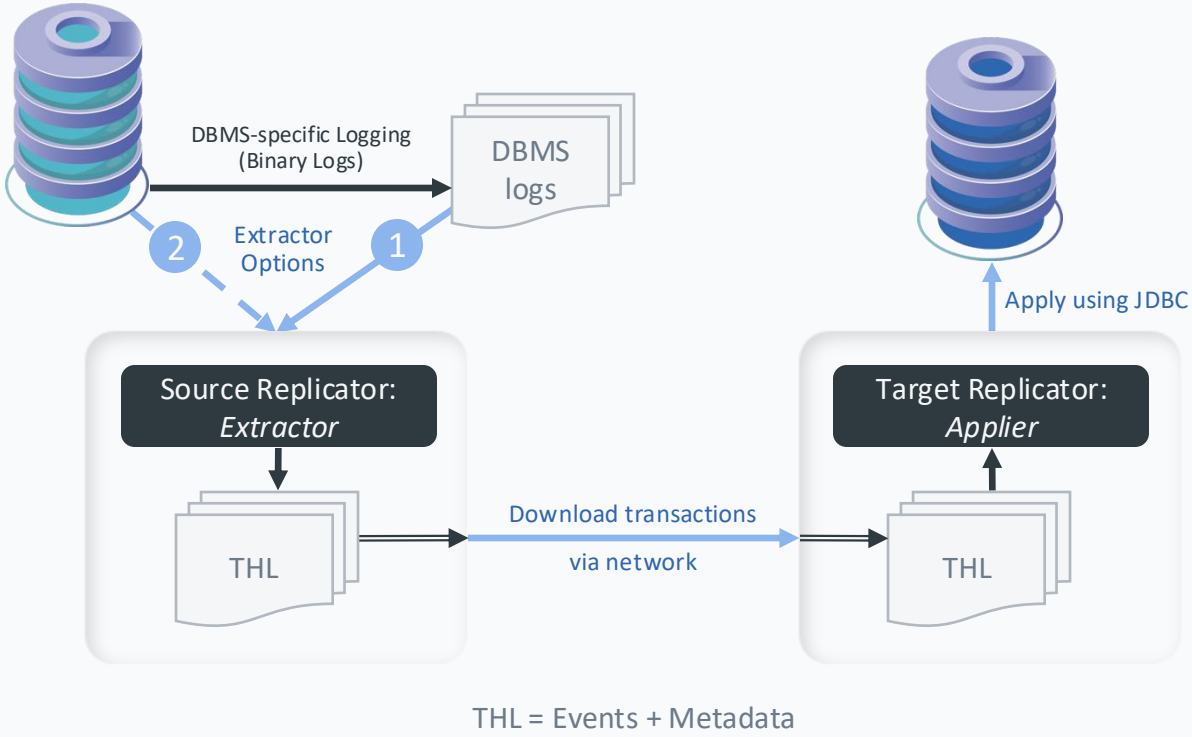

THL = Events + Metadata

Tungsten Replicator for Batch Targets

Option 1: Local Install
Extractor reads directly from the logs, even when the DBMS service is down. This is the default.

Option 2: Remote
Extractor gets log data via MySQL Replication protocols (which requires the DBMS service to be online) This is how we handle Amazon Aurora extraction tasks.

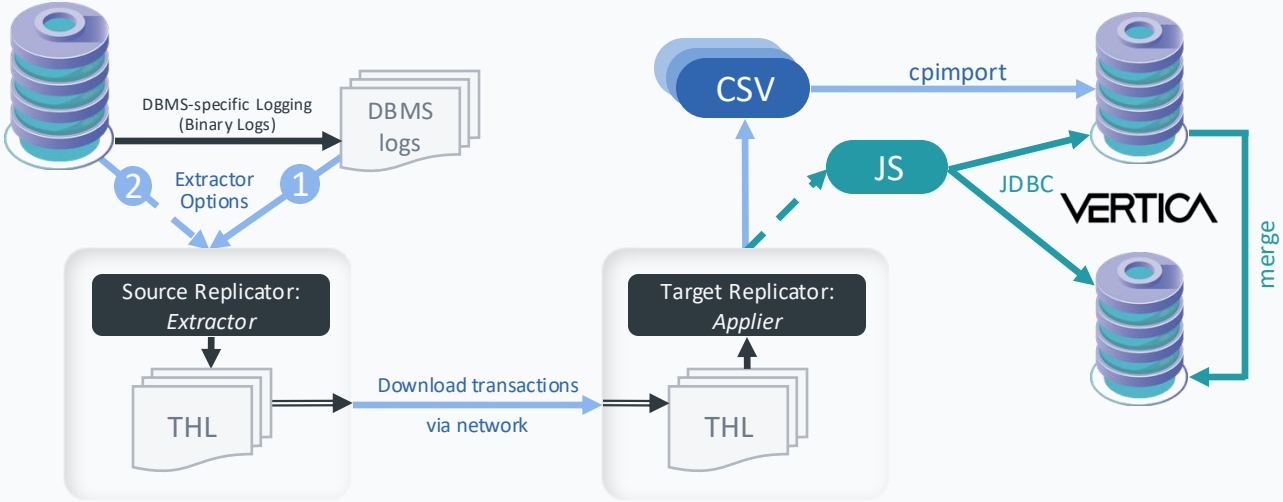

THL = Events + Metadata

Tungsten Replicator Topologies

Topologies

- Primary/Replica
- Active/Active (MySQL <> MySQL Only)
- Using Amazon Aurora/RDS as either a source and/or a target
- Fan-in
- Fan-out
- Cluster-Extractor
- Replicate into a cluster

Fan Out

- Extract Once
- Apply to many
- Targets can be a mix of any supported target
- Filtering can be used

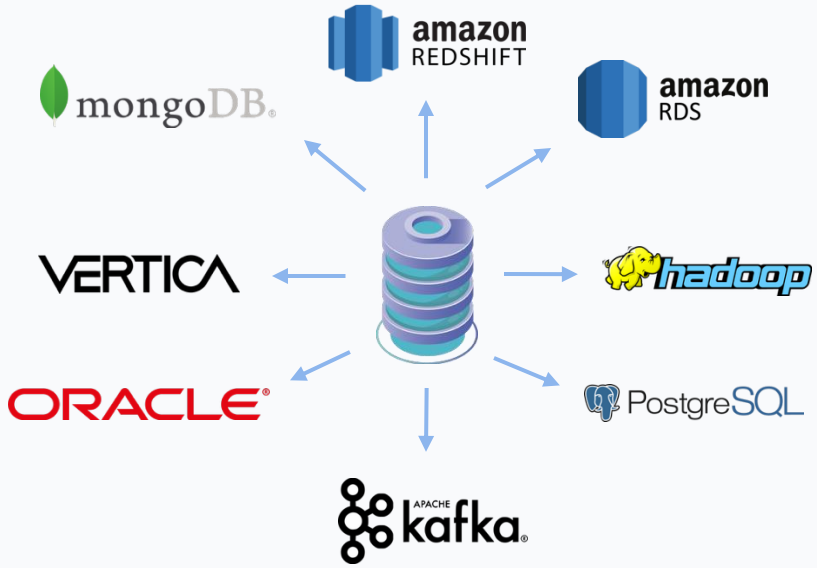

Fan In

- Many Sources into a single target
- Write to single schema or multiple
- Target can be any of the supported targets
- Mix of fan in and fan out

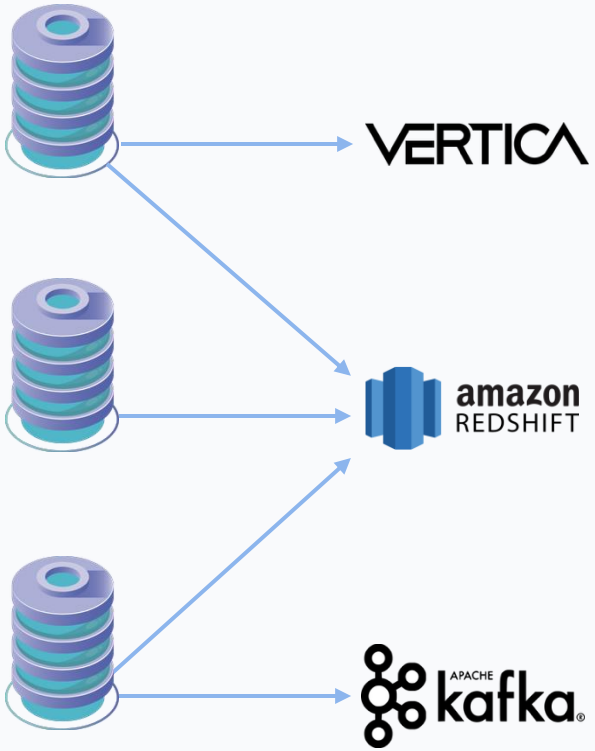

Cluster-Extractor

- Connect Applier to an existing Tungsten Cluster
  - Take advantage of THL already created by the extractor within a Cluster
- Apply to any supported target
- Continues working even after a manual Switch or an automatic Failover of the Primary node
  - No Need to reconfigure the source cluster!
- Ideal for offloading reporting functions or complementing existing backup/recovery procedures

Replicate into a Cluster

- Use replicator to setup replication into an existing Tungsten Cluster
- Extract from a single source
- Write in via Tungsten Proxy
- No need to reconfigure applier if Primary node fails within cluster

Filtering

Filtering

- Remove THL entries
  - For example, filter by schema/table name
  - Filter by data
  - Filter by type (statement, row, DDL, DML)
- Change Content
  - Remove a row or column
  - Change the contents of rows – for example string formats, or
  - Change the datatypes of rows – for example, an Int into Float, SET types into strings, or date types
  - Data Masking
- Add Data
  - Introduce entirely new transactions
  - Insert new data
  - Duplicate data into other tables

Replicator Stages

Tungsten Replicator Stages

Extractor

Each stage passes through:

`binlog-to-q`
Handles extraction of transactions from binary log and places into in-memory queue

`q-to-thl`
Takes transactions from in-memory queue and writes down to disk into thl log file

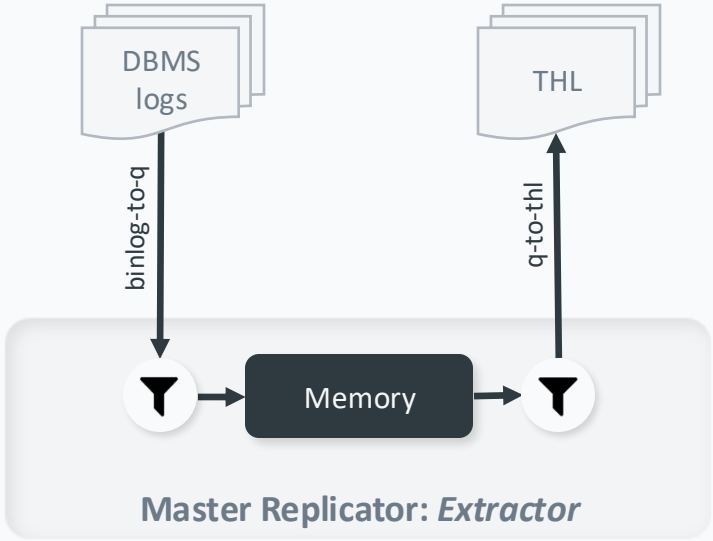

Tungsten Replicator Stages

Applier

`remote-to-thl`
Streams thl log files across network socket, from configured source. Writes out THL to local disk.

`thl-to-q`
Reads transactions from thl and places into in-memory queue

`q-to-dbms`
Reads transactions from in-memory queue and writes out to target database

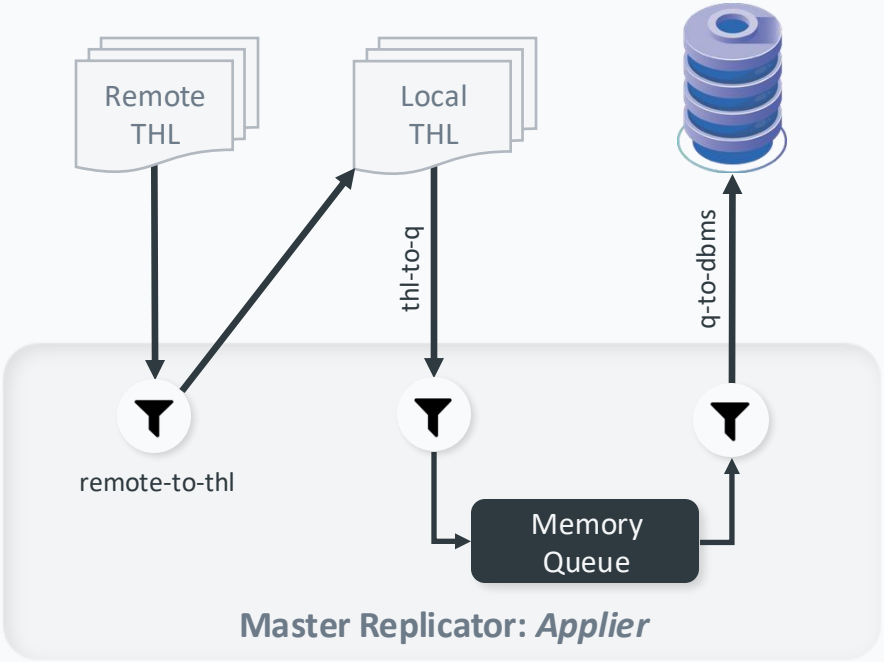

Summary

What we have learnt today

- How the Replicator Works
- Looked at possible topologies
- Discussed filtering
- Discussed the Replicator stages

Thank you for listening

continuent.com

Presented by Chris Parker

VP of Customer Success, EMEA