

Basics: Introduction to Tungsten Replication



Tungsten Replicator in a Nutshell

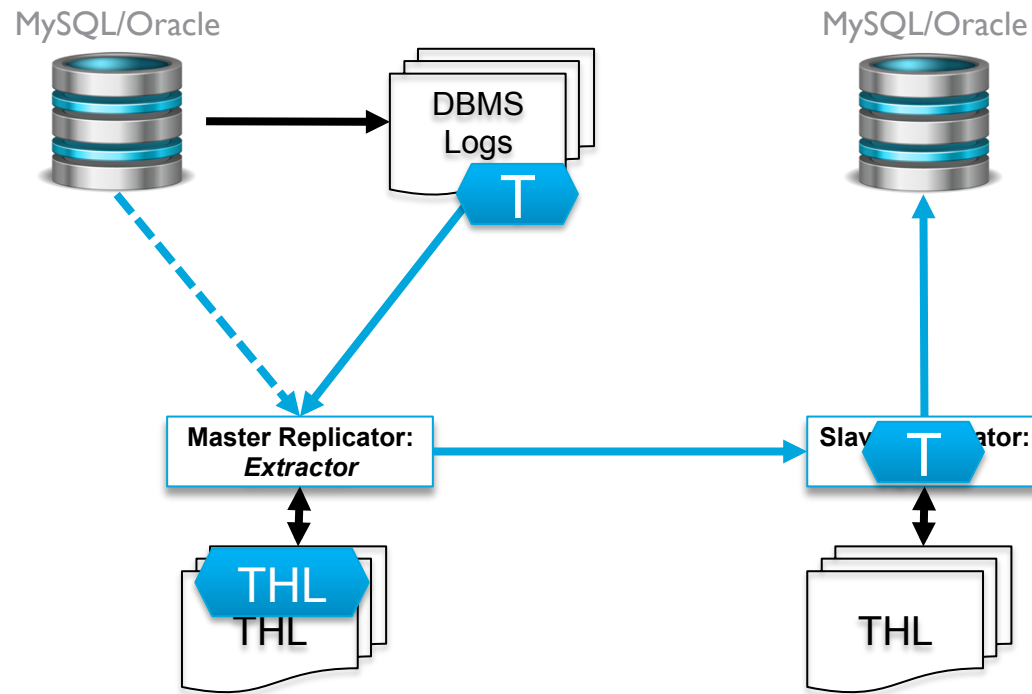


Tungsten Replicator is...

- Transactionally Consistent Asynchronous Replicator
- Supports SQL (homogeneous) and Row-based (heterogeneous)
- Database independent
- Architecture and database neutral storage format
- High performance
- Flexible Extraction
- Flexible Applying
 - JDBC
 - Native
 - Batch
- Flexible Filtering



Transactionally Consistent Asynchronous Replicator

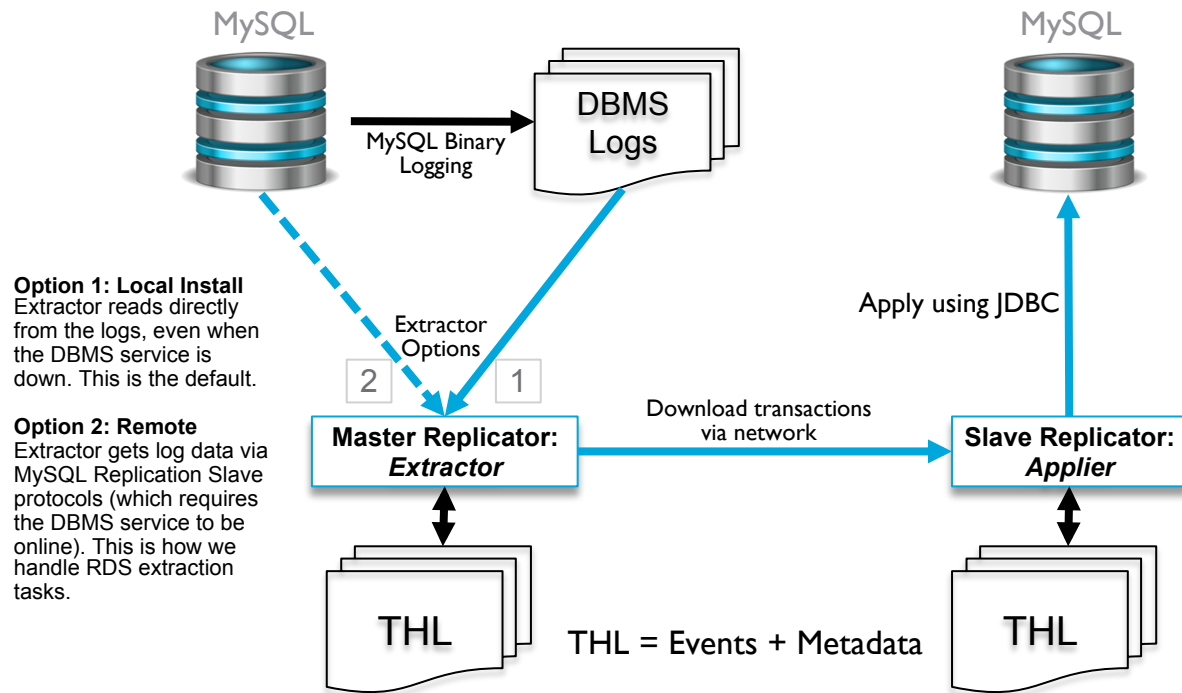


Sources

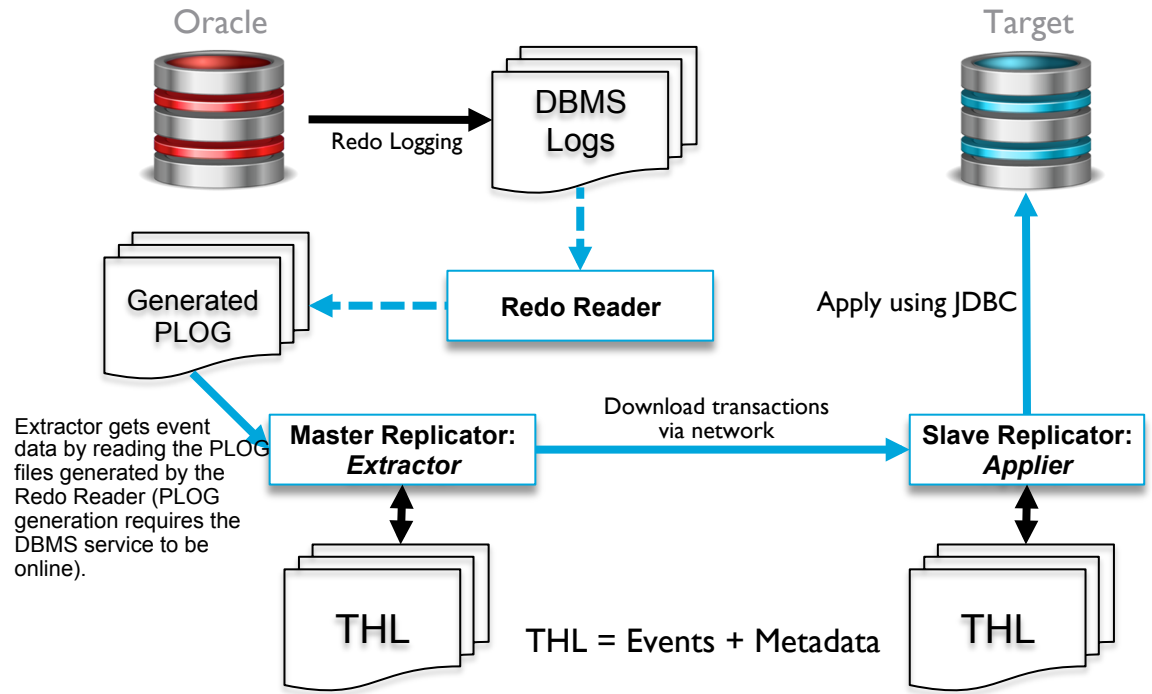
- MySQL
 - Oracle MySQL to 5.6
 - Percona MySQL to 5.6
 - MySQL 5.7 supported with limited data types (full support coming 2017)
 - MariaDB 10/10.1
- Oracle
 - Using 'Redo Reader', commercial license only
 - Supports 9i to 12c, SE1 to EE



MySQL Extraction



Oracle Extraction



Targets - JDBC

- MySQL – any version
- Oracle – any version



Targets - Native

- MongoDB
- ElasticSearch, announced
- Kafka, announced
- More coming in 2017



Targets - Batch


- Amazon RedShift
- HP Vertica
- Hadoop (native HDFS)

- Cassandra (announced)

- More coming in 2017

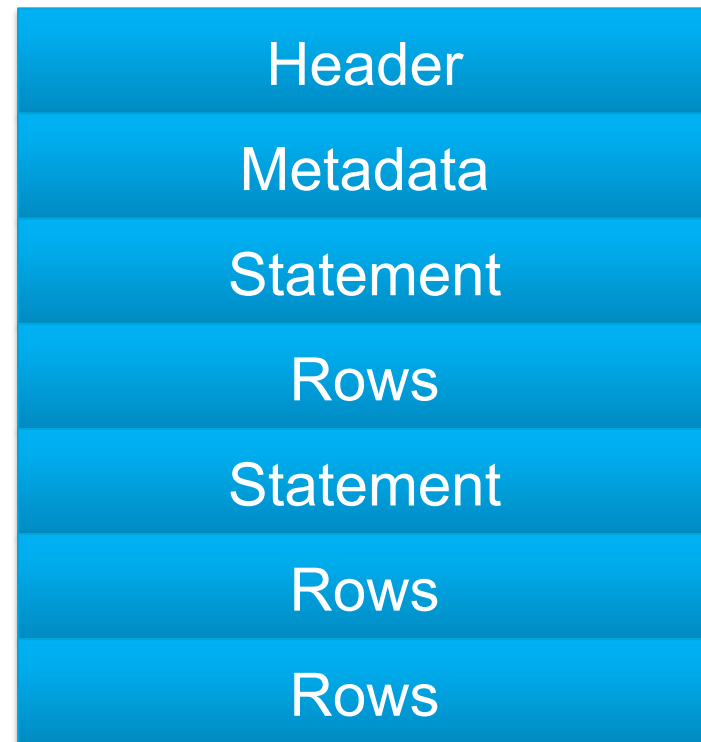


Heterogeneous




Transaction History Log (THL)

- A sequential log of each transaction
- Read from a source in the order applied
- Extracted with header, meta, and actual data
- THL is used internally by replicator
- Data extracted from source to THL
- THL exchanged between replicators
- Replicator converts single transaction into
 - JDBC statement
 - Native apply
 - Batched apply



Filtering



What Can Filtering Do?

- 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
- Add Data
 - Introduce entirely new transactions
 - Insert new data
 - Duplicate data into other tables



Tungsten Replicator is...

- Transactionally Consistent Asynchronous Replicator
- Supports SQL (Homogeneous) and Row-based (heterogeneous)
- Database independent
- Architecture and database neutral storage format
- High performance
- Flexible Extraction
- Flexible Applying
 - JDBC
 - Native
 - Batch
- Flexible Filtering



Upcoming Topics

- Basic Deployment
- How Appliers Work
- How Extraction Works
- Troubleshooting
- Filtering
- Advanced Deployments



Next Steps

- If you would like to download the clustering software so that you can the process for yourself, please visit: <https://github.com/continuent/tungsten-replicator>
- Read the documentation at <https://docs.continuent.com/tungsten-replicator-5.1/index.html>
- If you would prefer a more technical deep dive, please contact our sales team who will be able to take you through the details and setup a POC



Upcoming Training and Webinars

- **Tungsten Training Program**

- Training -- Tuesday, 27th June 2017, 09:00 PST, 17:00 BST, 30 minutes -- Basics: Introduction to Tungsten Clustering

- **Continuent Webinar Program**

- Webinar -- Wednesday, 21st June 2017, 09:00 PST, 17:00 BST, 30 minutes -- Amazon RDS vs. AWS with Continuent Tungsten Clustering



For more information, contact us:

Eero Teerikorpi
Founder, CEO
eero.teerikorpi@continuent.com
+1 (408) 431-3305

Eric Stone
COO
eric.stone@continuent.com

MC Brown
VP Products
mc.brown@continuent.com

