

Tungsten Replicator for Kafka - Real-time Data Loading from Amazon RDS into Kafka



Topics

In today's webinar, we will discuss:

- Why Kafka and Amazon RDS
- How replication into Kafka Works
- So I can replicate into Kafka, now what?
- Future direction

Why Kafka

- Kafka is a high performance message bus
- NOT a database
- Great for distributing messages and firing/triggering operations on content
- Log aggregation
- Activity/security tracking
- Metrics
- Auditing
- Data ingestion for Hadoop

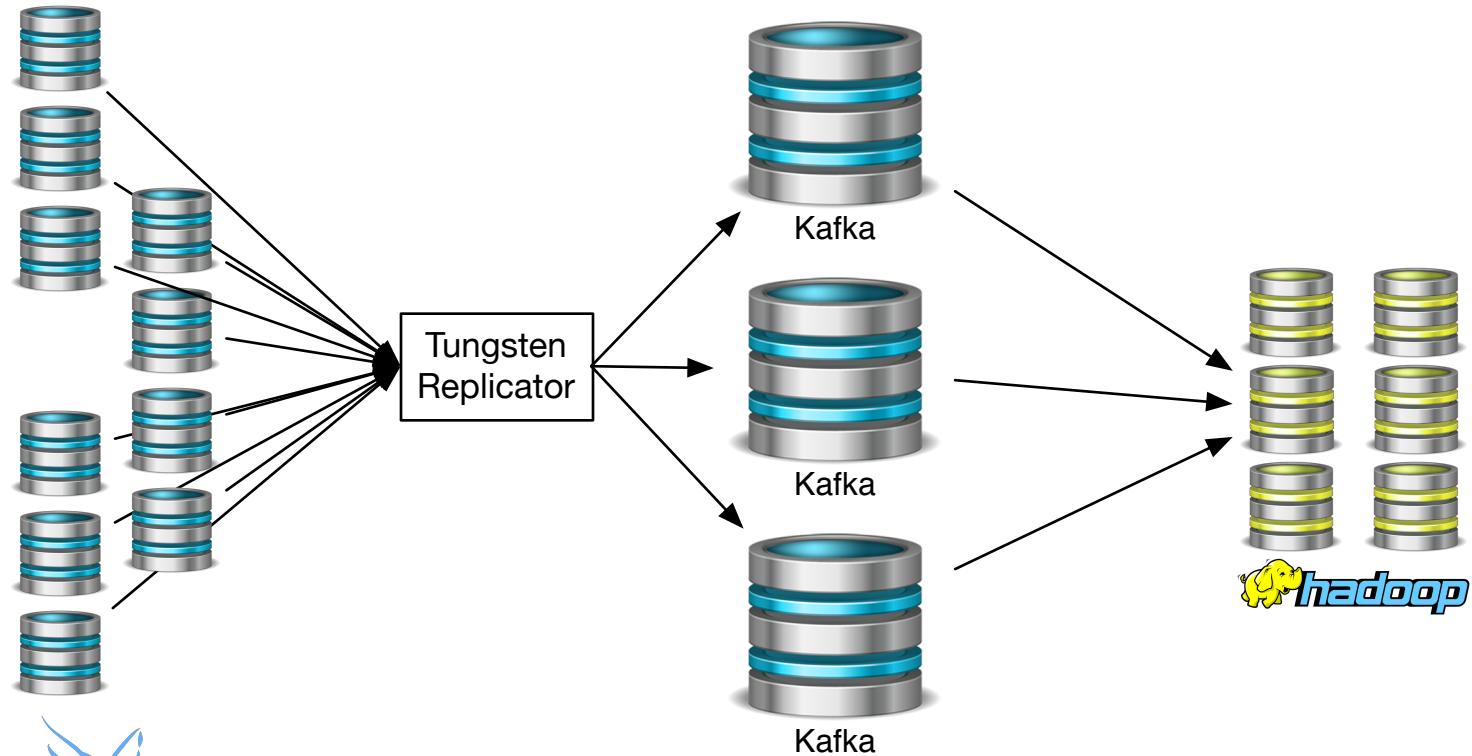


Why Amazon RDS to Kafka


- Amazon RDS allows fast MySQL deployment
- Kafka allows easy ingestion into multiple targets including Hadoop
- Cross AZ and datacenter movement
- Fully deployable with the Amazon scalable environment
- Low-latency
- Offload the responsibility of seeding your Kafka environment from your application



Mass Data Collection with Kafka

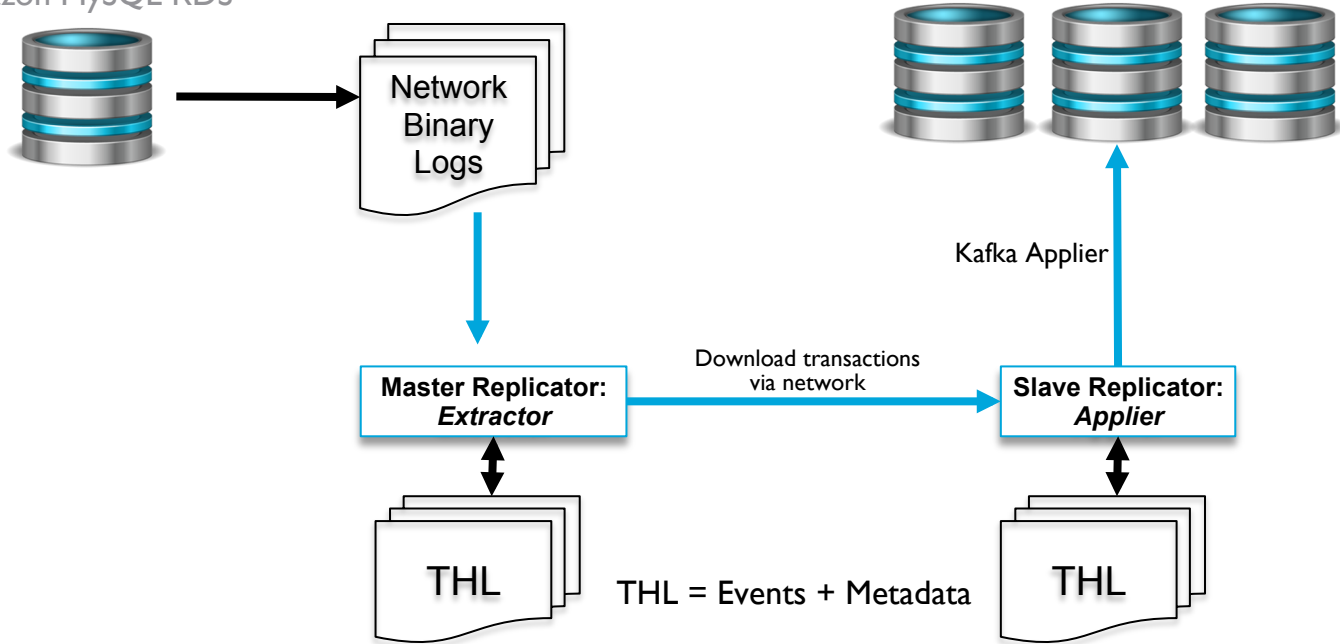


Multiple Target Distribution



How Replication Works

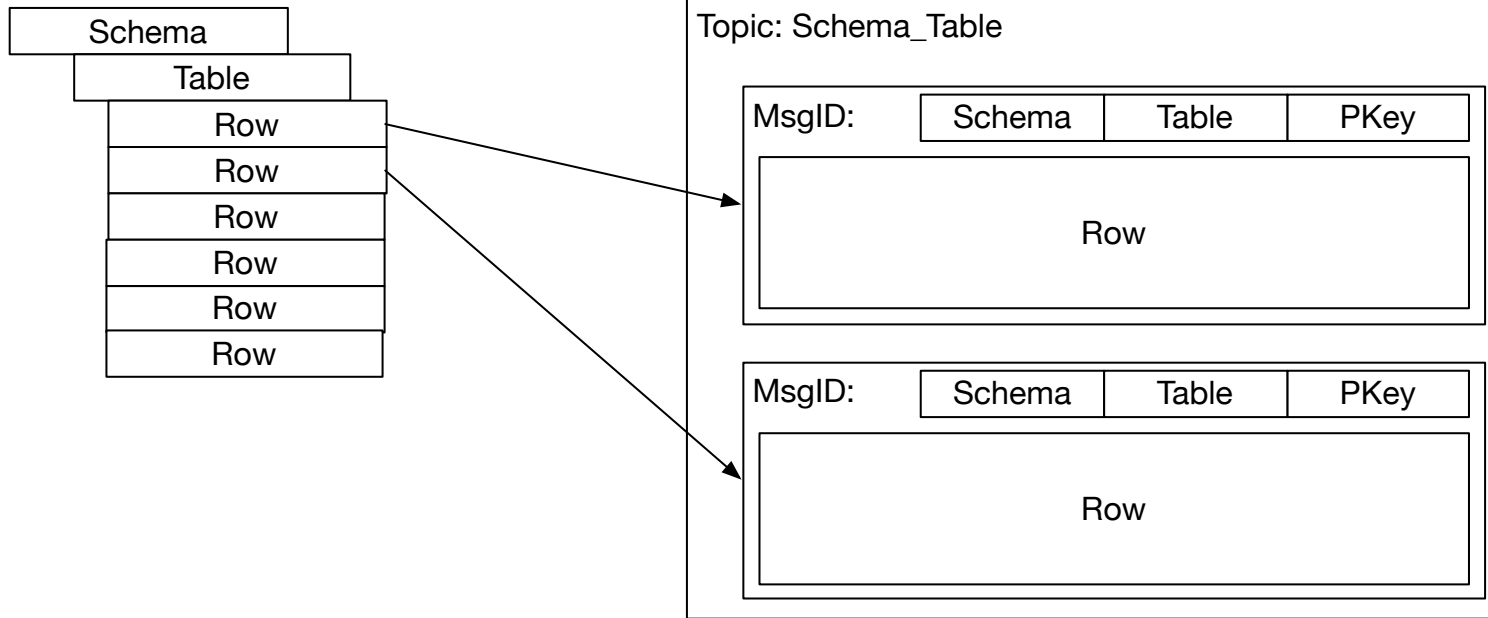
Amazon MySQL RDS



What Tungsten Replicator Does to Apply into Kafka

- Takes an incoming row and converts it to a message
- Message consists of metadata:
 - Schema name, table name
 - Sequence number
 - Commit timestamp
 - Operation Type
- Embedded Message Content

Message Structure



Customizable Elements

- Whether acknowledgements are required from Kafka
- How much distribution/replication is required before sending the message
- Format of the message key
- Whether to embed schema and table name
- Whether the commit timestamp should be embedded



Demo



Future Direction

- Full transaction support for Kafka
- New filters
 - Data obfuscation
 - Data modification
 - Row merging and basic extract augmentation
- Kafka Extraction
 - Parsing contents of Kafka message queues
 - Database updates
 - Large scale distribution of database changes
 - Filtering and re-submission



Next Steps

- If you are interested in knowing more about Tungsten Replicator and would like to try it out for yourself, please contact our sales team who will be able to take you through the details and setup a POC – sales@continuent.com
- Read the documentation at <http://docs.continuent.com/tungsten-replicator-5.2/index.html>
- Subscribe to our Tungsten University YouTube channel! <http://tinyurl.com/TungstenUni>



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

Chris Parker

Director, Professional Services EMEA & APAC

chris.parker@continuent.com