

Tungsten Replicator for Elasticsearch - Real-time data loading from Oracle and MySQL into Elasticsearch



Topics

In today's webinar, we will discuss:

- Why Elasticsearch?
- How replication into Elasticsearch works
- Customizations and configurations
- Future direction

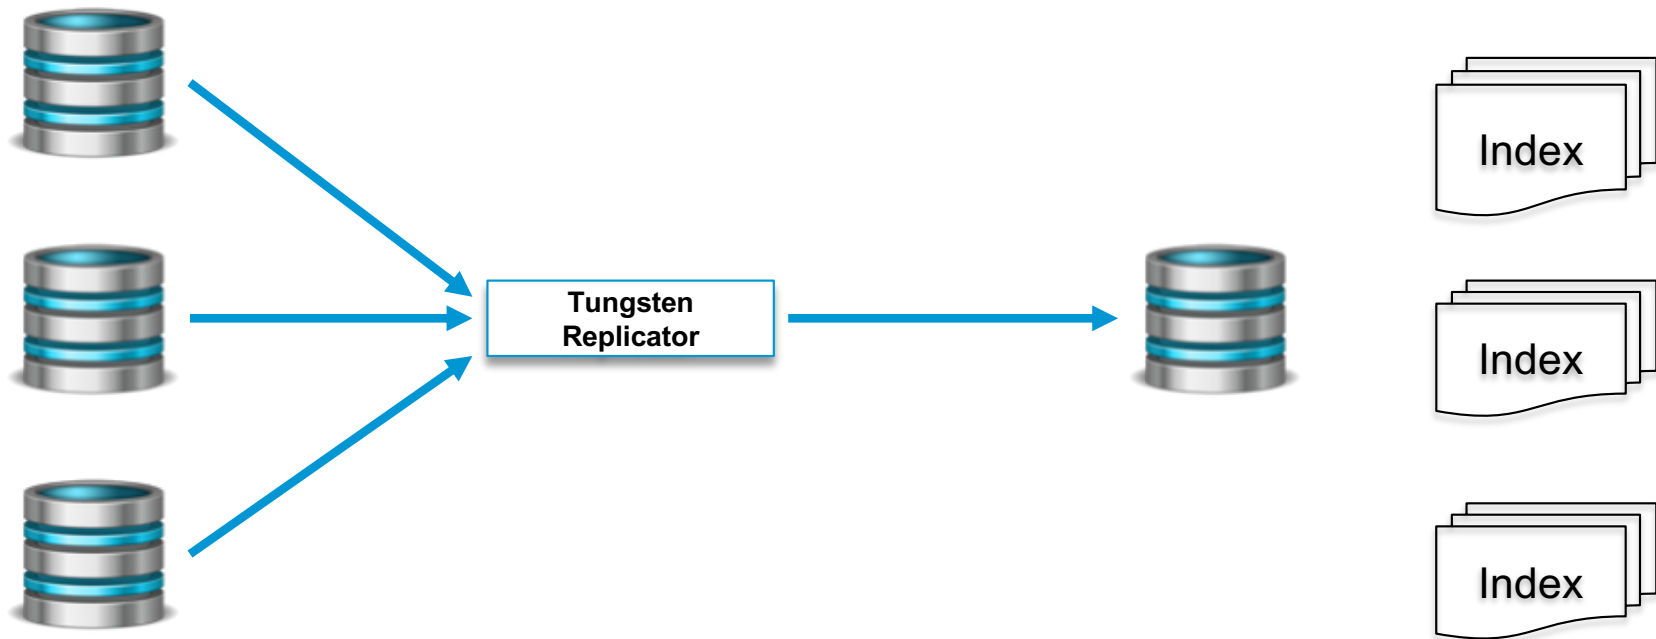


Why Elasticsearch

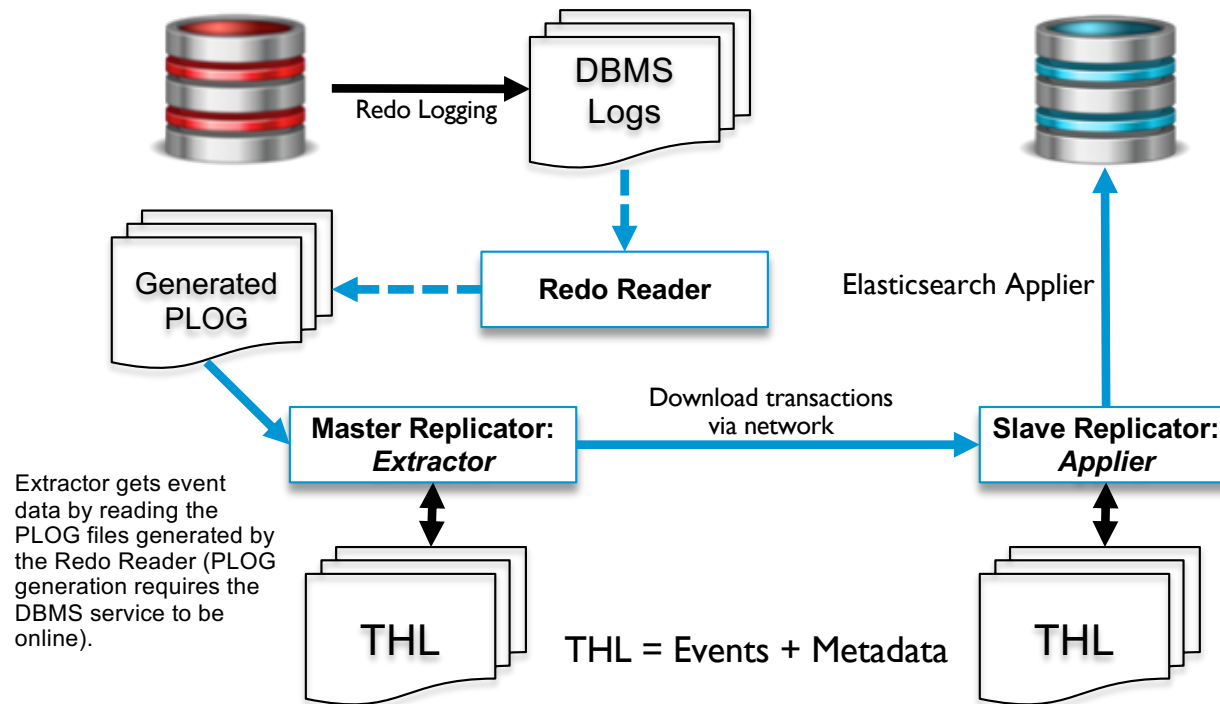
- Elasticsearch is a REST-based search and analytics engine
- Provides a quick method to store/search data and link back to source
- Very fast
- Great at doing very large scale analytics and aggregations
- Able to combine different forms and data (structured, unstructured, metrics)
- Easily scalable
- Provides different buckets and different types to allow multiple data repos



Data Collection and Organization with Elasticsearch



How Replication Works



What Tungsten Replicator Does to Apply into Elasticsearch

- Takes an incoming row and converts it to a JSON document
- By default:
 - Incoming schema/database name is used as the index name
 - Incoming table name is used as the type
- Can also be explicitly set to an index/type to concentrate data (or use filters)
- Key is taken from the primary key information
 - Can be converted/merged from multiple keys and formats
- Document is constructed with metadata and record information
- INSERTS create the data
- UPDATES update the existing data using the key information
- DELETE delete the record (but this can be disabled)



Message Structure

```
{
  "_index" : "test",
  "_type" : "messages",
  "_id" : "99999",
  "_version" : 1,
  "found" : true,
  "_source" : {
    "msg" : "Hello Elasticsearch",
    "committime" : "2017-06-23 19:02:22.0",
    "id" : "99999",
    "source_table" : "messages",
    "source_schema" : "test"
  }
}
```



Customizable Elements

- Documentation ID format
 - pkey, pkeyus, tspkey, tspkeyus
- Embedded information
 - Commit time
 - Schema/table source information
- Ignore delete or update errors
- Whether to use schema and table names as index and type
- Custom index name and type names
- Whether to use a self generated ID or one derived from primary key



Demo



Future Direction

- Further customization of the structures and stored information
- Integration with the data combination and concentration tools
- DDL integration and control of the structures
 - Currently TRUNCATE and CREATE/DROP TABLE are not replicated
- Improved data structures and record metadata



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>
- Follow us on Twitter @ continuent or www.facebook.com/Continuent
- 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

MC Brown

VP Products

mc.brown@continuent.com

Eric Stone

COO

eric.stone@continuent.com

Petri Virsunen

VP Marketing

petri.virsunen@continuent.com

+1 (408) 806-9860

