

# Advanced : Installing and Configuring Tungsten Replicator for Off-Board Oracle Extraction



# Topics

In this short course we will:

- Review Tungsten Replicator
- Review prerequisites required
- Discuss installation types (Offboard vs Onboard)
- Walkthrough an Installation (Full end to end demo)
- Recap Key Resources and Tools

## Course Prerequisite Learning

- Basics: Introduction to Tungsten Replication
- Advanced: Installing and Configuring Tungsten Replicator for Oracle Extraction
- Visit Continuent website or Tungsten University on YouTube to watch these recordings



# Tungsten Replicator

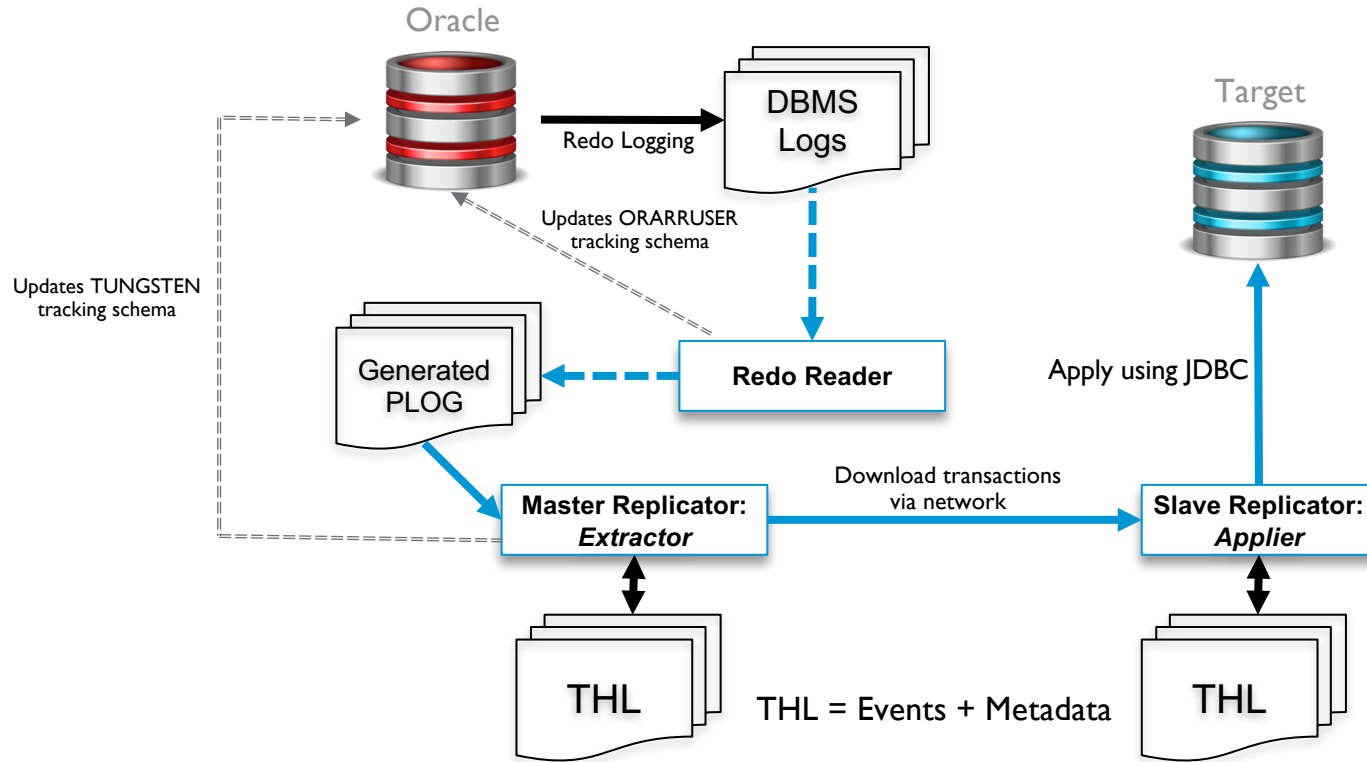


## Tungsten Replicator is...

- Transactionally Consistent, Asynchronous replication
- High performance
- Flexible Extraction
  - Oracle, MySQL and RDS
- Flexible Applying
  - Oracle, MySQL, RDS, Google Cloud SQL, Hadoop, Redshift, Vertica, Elasticsearch, MongoDB, Kafka Cassandra and many more currently in development coming soon
  - JDBC
  - Native
  - Batch
- Flexible Filtering



# Tungsten Replicator



# Prerequisites and Installation Types



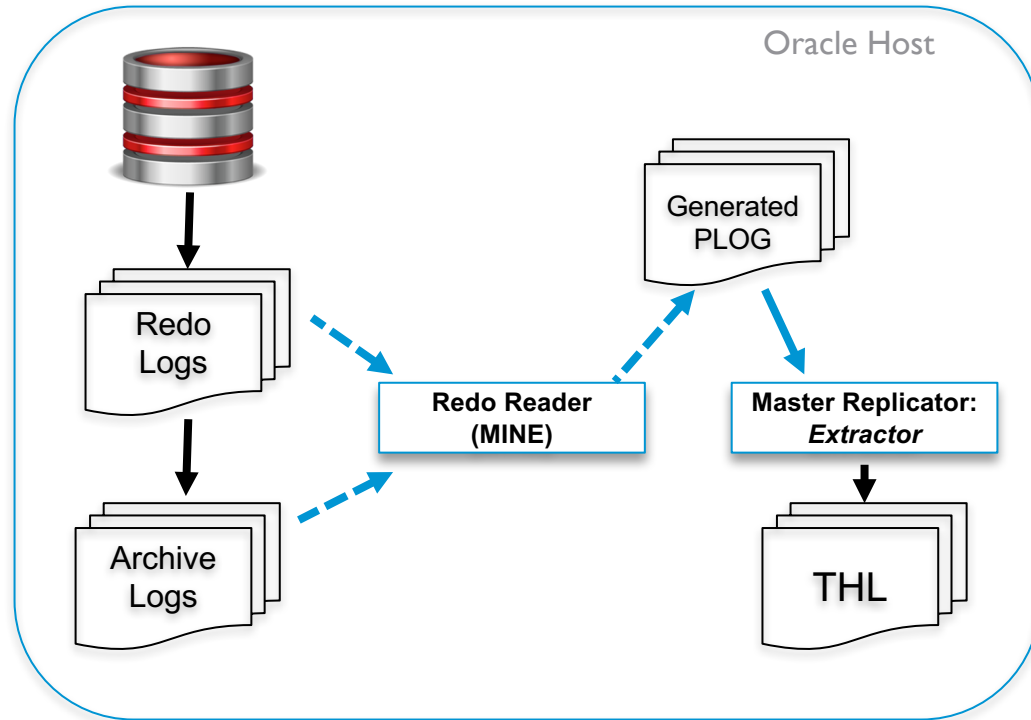
# Prerequisites

- Review docs
  - <http://docs.continuent.com/tungsten-replicator-5.2/index.html>
- Host Prerequisites
  - OS user
  - /etc/hosts
  - sudoers and ssh
  - Ruby
  - Java
  - Oracle client on Offboard host
  - NFS Mount Points configured for Offboard with NFS topology
- Network Prerequisites
  - Review port requirements
- Oracle Prerequisites
  - Archive and Supplemental Logging

tungsten user accounts

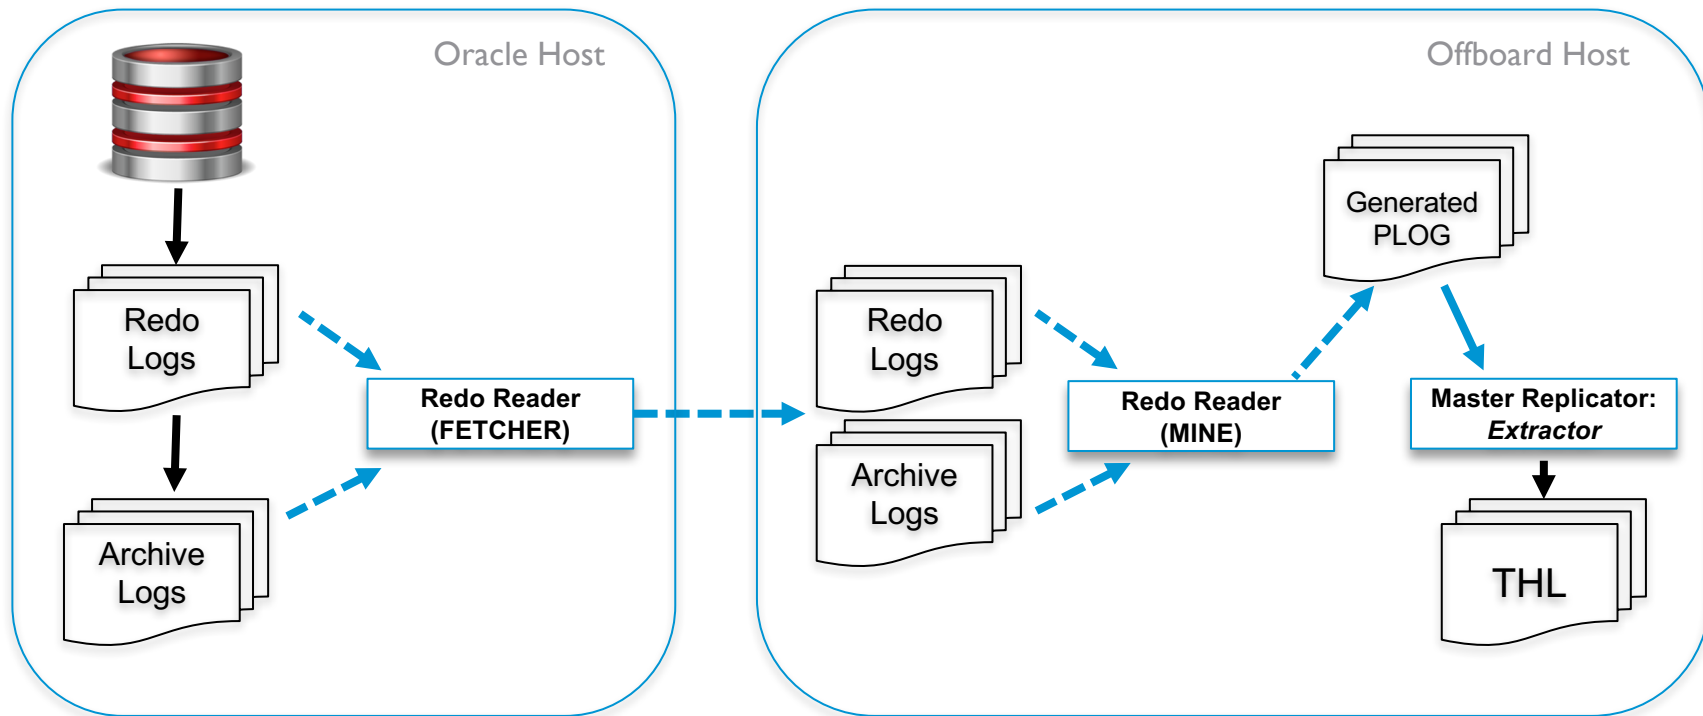


# Installation Type - Onboard

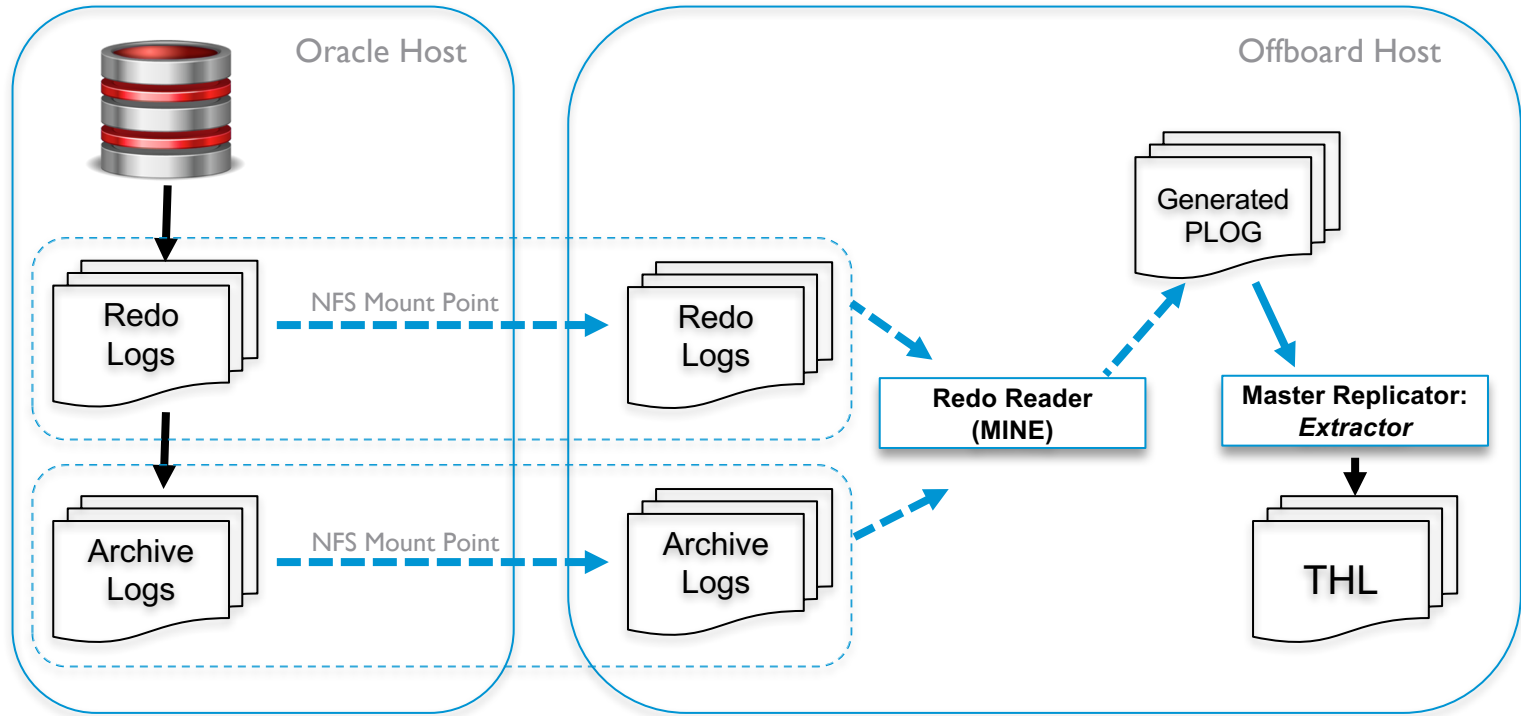




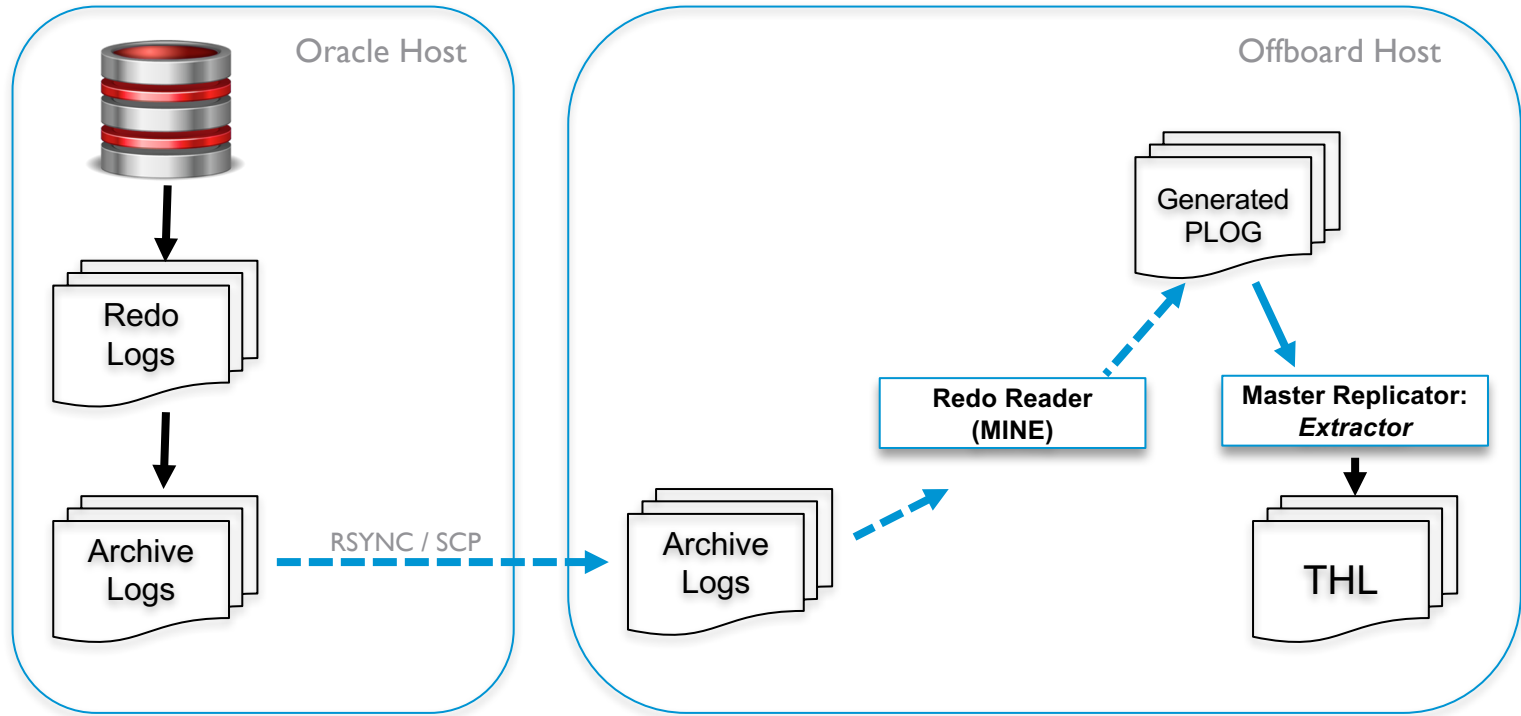
# Installation Type - Offboard with FETCHER



# Installation Type - Offboard with NFS



# Installation Type - Offboard with Log Shipping



## Installation Type - Summary

Type	Pros	Cons
Onboard	<ul style="list-style-type: none"> <li>• Faster</li> <li>• Easier to configure and manage</li> <li>• Single install for extractor</li> <li>• Less disk space required</li> </ul>	<ul style="list-style-type: none"> <li>• Currently, only supported for Oracle hosts running RHEL / CentOS / Oracle Linux. Further OS Support is planned but not yet available.</li> </ul>
Offboard with Fetcher	<ul style="list-style-type: none"> <li>• Supports more OS's</li> <li>• Separates extractor from DB host</li> <li>• Dedicated host for extraction</li> <li>• Offboard host could host multiple extractor services</li> </ul>	<ul style="list-style-type: none"> <li>• Much more disk space required</li> <li>• Multiple components to manage</li> <li>• Additional network ports</li> <li>• 2 Installations for extractor</li> </ul>
Offboard with NFS	<ul style="list-style-type: none"> <li>• As per "Offboard with Fetcher"</li> <li>• Less disk space required</li> </ul>	<ul style="list-style-type: none"> <li>• Potentially slower</li> <li>• Additional network ports</li> <li>• NFS requires managing</li> </ul>
Offboard with Log Shipping	<ul style="list-style-type: none"> <li>• As per "Offboard with Fetcher"</li> <li>• Only requires Archive Logs</li> <li>• Does not require NFS, Logs can be copied via RSYNC/SCP etc</li> </ul>	<ul style="list-style-type: none"> <li>• Always "one log behind"</li> <li>• Enough disk space required for Archive Logs</li> </ul>

# Demo



# Command Line Tools & Resources



## Tools : trepctl

- “trepctl status” displays the status of the local replicator
- “trepctl status –r 3” will show status output refreshed every 3 second until CTRL+C
- “trepctl qs” provides a quick summary overview of the local replicator
- “trepctl perf” provides deeper diagnostics of the different stages in the replicators

```
$ trepctl qs
State: east Online for 21.069s, running for 45.654s
Latency: 0.837s from DB commit time on db1 into THL
         21.839s since last database commit
Sequence: 1 last applied, 0 transactions behind (0-1 stored) estimate 0.00s before synchronization
```

## Tools : console

- The console is initiated using the start-console.sh script in the installation path
- It will connect to the MINE (or FETCHER) process depending on host run from
- “SHOW” will list all parameters
- “SUPPORT PACKAGE MINE|FECTHER|ALL” will generate zip file of diagnostics useful for support engineers when assisting with issues



## Log Files

- The `/opt/continuent/service_logs/` directory contains both text files and symbolic links.
- Symbolic links in the `service_logs` directory point to the replicator logs:
  - `/opt/continuent/tungsten/tungsten-replicator/log/`
- MINE and FETCHER processes log files to
  - `/opt/continuent/plog/<servicename>/logs`



## Next Steps

- If you are interested in knowing more or would like to try it out for yourself, please contact our sales team at [sales@continuent.com](mailto: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>
- Visit the events calendar on our website for upcoming Webinars and Training Sessions



## For more information, contact us:

**Eero Teerikorpi**

*Founder, CEO*

[eero.teerikorpi@continuent.com](mailto:eero.teerikorpi@continuent.com)

+1 (408) 431-3305

**Eric Stone**

*COO*

[eric.stone@continuent.com](mailto:eric.stone@continuent.com)

**MC Brown**

*VP Products*

[mc.brown@continuent.com](mailto:mc.brown@continuent.com)

**Chris Parker**

*Director, Professional Services, EMEA & APAC*

[chris.parker@continuent.com](mailto:chris.parker@continuent.com)

**Matthew Lang**

*Director, Professional Services, Americas*

[matthew.lang@continuent.com](mailto:matthew.lang@continuent.com)