

# Intermediate: Performing Cluster Maintenance with Zero Downtime



# Topics

In this short course we will:

- Review the Cluster Architecture
- Describe the Rolling Maintenance Process
- Explore what happens during a Master Switch
- Discuss Cluster States
- Demonstrate Rolling Maintenance
- Re-Cap Commands and resources used during the demo

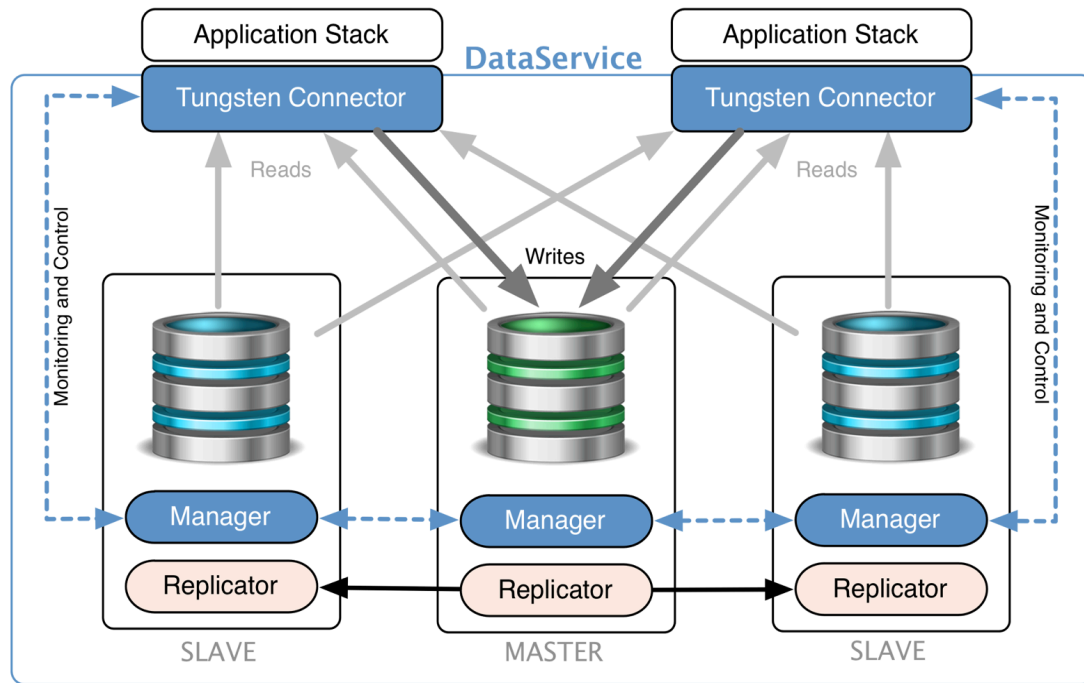
## Course Prerequisite Learning

- Basics: Introduction to Clustering
- Visit the Continuent website or Tungsten University on YouTube to watch these recordings
  - Continuent website <https://www.continuent.com/videos/>
  - Tungsten University on YouTube <https://www.youtube.com/channel/UCZ9iU-7nTIRLNnJvITFCsWA> or <http://tinyurl.com/TungstenUni>



# Tungsten Cluster Architecture

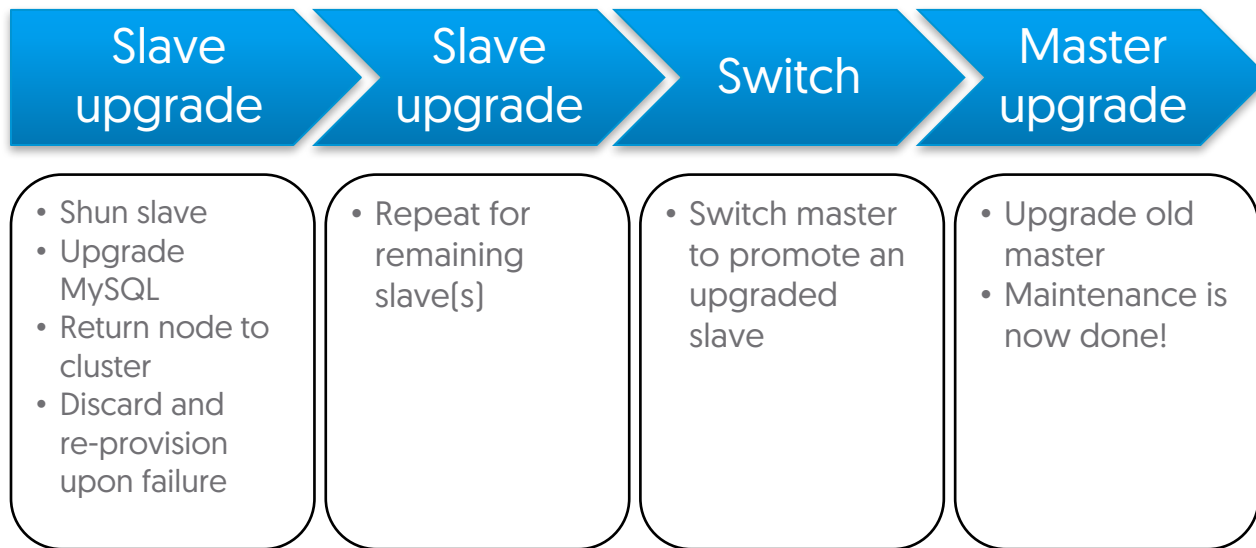
# Tungsten Cluster Architecture



# Zero-Downtime Rolling Maintenance



## Rolling maintenance proceeds node-by-node starting with the slaves and ending with the master



## Zero-Downtime Maintenance: The Switch

- The key to performing maintenance is the manual switch operation, where another node is selected to be the master.
- The switch command is invoked inside the `cctrl` command-line utility. In the below example, the Manager is allowed to select the new node:

```
cctrl> switch
```

- Instead of allowing the cluster to decide, you can switch to a specific node:

```
cctrl> switch to db3
```

- Make sure replication is 100% caught up before you switch because the cluster will let you switch to a new master that is missing transactions!

# The process of switching the Master...

```
[LOGICAL] /nyc > switch
```



Application

Reads & Writes

1. Halt  
in-coming  
connections



Tungsten  
Connector

Reads & Writes



MySQL  
**Master**

2. Master is shunned,  
wait for in flight  
queries to complete

Replication  
Traffic



MySQL  
**Slave**

Replication  
Traffic

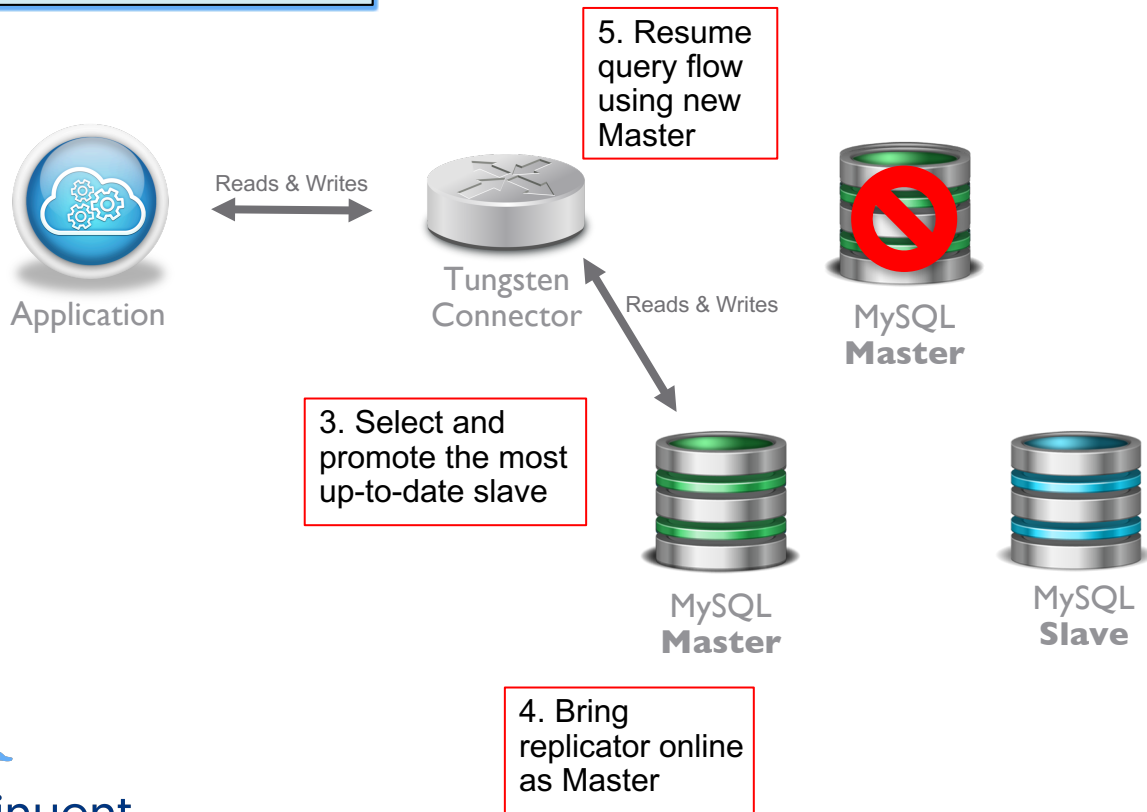


MySQL  
**Slave**



# The process of switching the Master...

```
[LOGICAL] /nyc > switch
```



# The process of switching the Master...

```
[LOGICAL] /nyc > switch
```



Application

Reads & Writes



Tungsten  
Connector

Reads & Writes



MySQL  
Slave

6. Convert old  
Master to a Slave



MySQL  
Master

Replication  
Traffic



MySQL  
Slave

7. Reconfigure  
Slave replicators  
to use the new  
Master

8. Replication  
Recovered, Switch  
Complete

## Rolling maintenance steps - Recap

- set policy maintenance
- On Slave Node:
  - `datasource <node> shun`
  - `replicator <node> offline`
    - Perform Maintenance
  - `datasource <node> recover`
- Repeat on all Slave Nodes
- Promote a Slave to Master
  - `switch`
- Repeat steps on old Master after switch complete

# Cluster States

- AUTOMATIC
  - “set policy automatic”
  - When the cluster is in Automatic Mode,
    - Automatic Recovery of services is attempted
    - Automatic Failover will happen if the Master node fails
  - Cluster should ALWAYS be in this state under normal operation
- MAINTENANCE
  - “set policy maintenance”
  - Automatic Recovery of services will NOT happen
  - Failover will not occur if Master node fails
  - Cluster should NEVER be left in this state in Normal Operation
  - This state should only be used for Maintenance Operations

## Cluster Datasource states

- ONLINE
  - “datasource <node> online”
  - The node is Online and available to Failover (If a slave) and data operations
- OFFLINE
  - “datasource <node> offline”
  - Not available for data operations.
  - Will not accept connections through the connector
  - Still receives and processes replication data
  - If cluster in AUTOMATIC Mode, cluster will attempt automatic recovery to ONLINE state
- SHUNNED
  - “datasource <node> shun”
  - As per OFFLINE but automatic recovery will NOT be attempted

# Rolling Maintenance Demo

- Tungsten Clustering 5.2
- Connectors in Port-Based Routing mode
- Amazon AWS EC2
- MySQL Community 5.7
- Java 1.8
- Ruby 2.0
- Maintenance: Change location of MySQL Error Log

# Command Line Tools & Resources

## Tools : cctrl

- “cctrl” can be run from any node within a cluster to control the local cluster and gather information
- Type “help” to get a full list of all commands available
- “ls” provides a summary overview of the entire cluster
- “datasource <node> shun”
  - Isolates node in cluster
- “replicator <node> offline”
  - Stop node receiving replication data
- “datasource <node> recover”
  - Re-Introduce node to cluster
- “switch” or “switch to <node>”

```
[LOGICAL] /nyc > ls

COORDINATOR[db1:AUTOMATIC:ONLINE]

ROUTERS:
+-----+
|connector@db1.tt-0509[25431] (ONLINE, created=0, active=0)|
|connector@db2.tt-0509[24548] (ONLINE, created=0, active=0)|
|connector@db3.tt-0509[24803] (ONLINE, created=0, active=0)|
+-----+

DATASOURCES:
+-----+
|db1(master:ONLINE, progress=0, THL latency=0.776)|
|STATUS [OK] [2017/08/17 11:00:53 AM UTC]|
+-----+
| MANAGER(state=ONLINE)|
| REPLICATOR(role=master, state=ONLINE)|
| DATASERVER(state=ONLINE)|
| CONNECTIONS(created=0, active=0)|
+-----+
. . .
```



## Tools : trepctl

- “trepctl status” can be run from any node within a cluster to view 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 offline|online”

```
$ 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
```



# Log Files

- The `/opt/continuent/service_logs/` directory contains both text files and symbolic links for cluster log files.
- Links in the `service_logs` directory go to one of three [3] subdirectories:
  - `/opt/continuent/tungsten/tungsten-connector/log/`
  - `/opt/continuent/tungsten/tungsten-manager/log/`
  - `/opt/continuent/tungsten/tungsten-replicator/log/`

```
tungsten@db1:/opt/continuent/service_logs $ ll
total 116
lrwxrwxrwx 1 tungsten tungsten    61 Jun 22 09:52 connector.log -> /opt/continuent/tungsten/tungsten-connector/log/connector.log
lrwxrwxrwx 1 tungsten tungsten    62 Jun 22 09:52 mysqldump.log -> /opt/continuent/tungsten/tungsten-replicator/log/mysqldump.log
lrwxrwxrwx 1 tungsten tungsten    55 Jun 22 09:52 tmsvc.log -> /opt/continuent/tungsten/tungsten-manager/log/tmsvc.log
lrwxrwxrwx 1 tungsten tungsten    60 Jun 22 09:52 trepsvc.log -> /opt/continuent/tungsten/tungsten-replicator/log/trepsvc.log
lrwxrwxrwx 1 tungsten tungsten    63 Jun 22 09:52 xtrabackup.log -> /opt/continuent/tungsten/tungsten-replicator/log/xtrabackup.log
```

## Tools : tpm diag

- Provides support engineers with an entire overview of the cluster state, by:
  - Gathering point-in-time status of all components in a cluster
  - Gathering log files of all components in a cluster, including database logs to provide historical information
  - Bundles everything into one easy zip file that can be attached to a support case
  - ALWAYS create a diag package when you contact support for assistance
- tungsten\_send\_diag
  - Executes “tpm diag” to generate the diagnostic package
  - Automatically uploads the package to support
  - [https://docs.continuent.com/tungsten-clustering-5.2/cmdline-tools-tungsten\\_send\\_diag.html](https://docs.continuent.com/tungsten-clustering-5.2/cmdline-tools-tungsten_send_diag.html)

```
tungsten@db1:/opt/continuent/service_logs $ tungsten_send_diag -d -c 1234
```

## Next Steps

- If you are interested in knowing more about the clustering software 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](mailto:sales@continuent.com)
- Read the documentation at <http://docs.continuent.com/tungsten-clustering-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  
<https://www.continuent.com/events/>
- Tues 3<sup>rd</sup> October : Advanced: Performing Schema Changes in a Multi-Site/Multi-Master Cluster  
<https://attendee.gotowebinar.com/register/8332481030891376387>
- Percona Live: Europe : 26th – 28th September : <https://www.percona.com/live/e17/>

## 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)