

# Basics: Backup, Recovery, and Provisioning with a Continuent Tungsten Cluster



# Topics

In this short course we will:

- Methods and Tools for taking a backup
- Verifying the backup contains the last binary log position, and the importance of this
- Restore various backups into a cluster
- Provisioning slaves from an existing datasource

## Course Prerequisite Learning

- Basics: Introduction to Tungsten Clustering
- Basics: Simple Tungsten Cluster Deployments
- 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-7nT1RLNnJvITFCsWA> or <http://tinyurl.com/TungstenUni>



# Methods and Tools for Taking Backups



# Backup Tools

- mysqldump
    - Installed with most MySQL installations
    - Easy to use for taking a backup
    - Easy to use for restoring a backup
    - Quick for small datasets
    - Very slow for large datasets
    - Locks tables
    - Stresses the MySQL database server
    - Slow to restore
  - xtrabackup
    - Open source
    - Does not lock tables
    - Although it will put load on the server and OS, it does not put load on the MySQL database server
    - Ideal for larger datasets (200– 700 GB)
- Can be difficult to use

## Additional Backup Tools

- Snapshots (lvm or other)
  - Ideal for large datasets
  - Can be extremely quick to take the snapshot
  - Requires working at the OS or product level
  - Having many snapshots can impact disk performance
- File copy
  - Good for one off backups or backup test systems
  - Fast Restore
  - Will work for large datasets
  - Requires scripting to automate
  - Must stop MySQL database server to get a consistent backup
- Tungsten\_provision\_slave
  - The easiest way to provision a failed slave!

# Taking and Restoring Backups



## Taking a backup within a Tungsten Cluster

```
[LOGICAL:EXPERT] /alpha > datasource host3 backup
```

Using the 'mysqldump' backup agent.

```
Replicator 'host3' starting backup Backup of dataSource 'host3' succeeded;  
uri=storage://file-system/store-0000000001.properties
```

- Select a slave when using this
- Backups by default will be saved into /opt/continuent/backups
- Backup command will use xtrabackup if available, otherwise it will fall back to mysqldump
- A properties file is also created, which includes backup method and checksum

## Automating backups within a Tungsten Cluster

- Use the `cluster_backup` command within cron on all database nodes
- `cluster_backup` will ONLY run on the coordinator; it will gracefully exit on other hosts
- Backup will be performed on a slave
- Will use xtrabackup if available, otherwise will fall back to mysqldump

```
shell> crontab -l  
00 00 * * * /opt/continuent/tungsten/cluster-home/bin/cluster_backup >>/opt/continuent/service_logs/cluster_backup.log 2>&1
```



## Restoring a backup within a Tungsten Cluster

```
[LOGICAL] /alpha > datasource host3 shun  
[LOGICAL] /alpha > replicator host3 offline  
[LOGICAL] /alpha > datasource host3 restore
```

- The restore process will take the latest available backup from **the host being store**
- Does not search the entire cluster for the latest backup

```
[LOGICAL] /alpha > datasource host3 shun  
[LOGICAL] /alpha > replicator host3 offline  
[LOGICAL] /alpha > datasource host3 restore storage://file-system/store-  
0000000004.properties
```

- Specify the location of the properties file to restore a specific backup

## Taking and Restoring an external backup

```
[LOGICAL] /alpha > mysqldump --opt --single-transaction --all-databases --add-drop-database --master-data=2
```

- Will lock tables, recommended to run on a slave
- Must include **--master-data=2**
  - Necessary when backing up master
  - `CHANGE MASTER TO MASTER_LOG_FILE='mysql-bin.000002', MASTER_LOG_POS=622;`

```
[LOGICAL] /alpha > datasource host3 shun  
[LOGICAL] /alpha > replicator host3 offline  
shell > th1 -service alpha purge
```

- Restore the backup (xtrabackup copy-backup, copy snapshot, etc)

```
[LOGICAL] /alpha > datasource host3 recover
```

## Provision a Slave from an Existing Node

- `tungsten_provision_slave` allows you to easily provision a node from a remote host
- Is run on the slave to be provisioned
- Runs the backup on the remote hosts specified by the `source` option
- Copies the backup to the local server and restores it
- Will recover replication services and have the newly provisioned node (re)join the cluster
- Like the Tungsten backup options, it will use `xtrabackup` if available, otherwise will use `mysqldump`

```
shell@host3> tungsten_provision_slave --source host2
```

- When run on `host3`, will provision from host `host2`

Demo: tungsten\_provision\_slave



# Restoring a Backup Taken from a Master

# The Tungsten Schema

- Tungsten creates a “tracking” schema in the MySQL database
- GTID and binlog positions are stored within this schema
- When a database is restored, Tungsten used the data in the tracking schema to set the replicator position

```
mysql> select * from tungsten_east.trep_commit_seqno;
```

task_id	seqno	fragno	last_frag	source_id	epoch_number	eventid	applied_latency	update_timestamp	shard_id	extract_timestamp
0	190	0	1	db1	0	mysql-bin.000002:000000000048972;32826				
	1			music				2017-10-17 03:20:39		2017-10-17 03:20:38



## Setting the Correct binary log position

- When backing up a busy master, the actual binary log position may be different than the values in the `trep_commit_seqno` table
- Compare the values in `trep_commit_seqno` with the values obtained from the backup (`mysqldump master-data` or `xtrabckup_info`)
- If the values are the same, no further action is needed
- If the values differ, note the sequence number in `trep_commit_seqno` table
- Find the binlog position in the THL, using the sequence number from the above step as a starting point
- Note this new sequence number obtained from the THL, and use `tungsten_set_position` on the slave node to set the replicator position with this new sequence number

## Setting the Correct binary log position (example)

```
mysql> select seqno,source_id, event_id from tungsten_east.trep_commt_seqno;  
seqno      source_id  eventid  
32033674   host1     mysql-bin.000032:0000000473860407;-1
```

Examine binary log position in backup. In our example, the position is `mysql-bin.000032:473863524`.

Now on the master:

```
Shell> thl list -service alpha -low 32033674 -headers | grep 473863524  
32033678 32030709 0 true 2014-10-17 16:58:11.0 mysql-bin.000032:0000000473863524;-1 db1-east.continuent.com
```

Back to the slave:

```
Shell> tungsten_set_position --service=alpha --source=host1 --seqno=32033678
```



# What's Next?



# Tungsten Backup

- Fully featured, long-term backup solution for any of your data sources (MySQL and Oracle, more coming)
- Full Point-in-Time recovery
- Configurable RTO/RPO: Configure how quickly you want to get running
- Restore back to source, or write your data out to Hadoop, Kafka or any other applier
- Also suitable for Point in Time Analytics!
- Tungsten Backup v1.0 will be free of charge

## 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/>



## 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)