Tungsten Cluster Master Class

Advanced: Multi Cluster Topologies

Chris Parker, Customer Success Director, EMEA & APAC



Topics

In this short course, we will

- Review Multi Cluster Use Topologies
- Understand Multi Cluster Architecture
- Explore Configuration Differences

Topologies

Tungsten Cluster+

Composite Active/Passive Cluster

- Minimum 2 clusters
 - Min 3 nodes per cluster
- 1 Cluster for read/write
- All other clusters read-only
- Can be Cross-Region
- Simple to use and implement
- Single, write-able Primary
- Managed cross-site replication
- Often used for Disaster Recovery





Tungsten Cluster+

Multi-Site/Active-Active Cluster

- Minimum 2 Clusters
 - Min 3 nodes per cluster
- All clusters are read/write
- Can be Cross-Region
- Each cluster is independent
- Un-managed cross-site replication
- Complex setup
- Dashboard will show each cluster but not cross-cluster replication
- All versions of Tungsten Clustering





Tungsten Cluster+

Composite Active/Active Cluster

- Minimum 2 Clusters
 - Min 3 nodes per cluster
- All clusters are read/write
- Cross-Region
- Vastly simplified setup and control
- Multiple write-able Primaries
- Managed cross-site replication the managers are aware of and control the cross-site Replication services.
- Full Dashboard integration
- Version 6.0+ of Tungsten Clustering





Configuration Considerations

Installation Considerations

- All usual Pre-Reqs need to be in place
- Network Latency between clusters
- Network Connectivity
 - Additional THL Ports (Active-Active)
- MUST be ROW Based (Active-Active)
- Conflict Resolution (Active-Active)
 - Auto-Increment Offsets
- Connector Behaviour

Active-Active Specifics

- Cross-Site services auto configured
 - x_from_y
- Local Cluster defaults to THL Port 2112
- Remote Cluster services are assigned THL Port 2113 onwards
 - e.g. 3 cluster configuration (east, west, north)
 - east = 2112
 - east_from_west = 2113
 - east_from_north = 2114
 - west = 2112
 - west_from_east = 2113
 - west_from_north = 2114
 - etc



Sample Configs

Composite Active/Passive

- One Cluster identified as the Active cluster
- All other clusters are Passive clusters
- Passive clusters can service reads to Connectors
- All writes, from any Connector, will go to Primary host in Active Cluster
- The Relay node in Passive Clusters becomes the Primary during Switch
- Moving the Priamry role to any Passive Cluster is a single `switch` command in cctrl
- FAILOVER between clusters is NOT Automatic
 - This avoids potential flapping if network between remote clusters unstable.
 - Avoids Split-Brain
 - Handled with one or two simple commands in cctrl

[defaults]

user=tungsten home-directory=/opt/continuent application-user=app_user application-password=secret application-port=3306 profile-script=~/.bash_profile replication-user=tungsten replication-password=secret mysql-allow-intensive-checks=true

[nyc]

topology=clustered
master=trainingdb1
members=trainingdb1,trainingdb2,trainingdb3
connectors=trainingdb1,trainingdb2,trainingdb3

[london]

topology=clustered
master=trainingdb4
members=trainingdb4,trainingdb5,trainingdb6
connectors=trainingdb4,trainingdb5,trainingdb6
relay-source=nyc

[global] composite-datasources=nyc,london

Multi-site/Active-Active

- Incredibly complex
- Requires installation of both Cluster software AND Replication Software separately
- Cross Site Replicators are totally unaware of underlying cluster
- Requires extensive configuration
- Connectors only work with their associated cluster
 - Cannot re-route sessions to any remote cluster

Multi-site/Active-Active

[defaults]

user=tungsten home-directory=/opt/continuent application-user=app_user application-password=secret application-port=3306 profile-script=~/.bash_profile replication-user=tungsten replication-password=secret mysql-allow-intensive-checks=true

[defaults.replicator] home-directory=/opt/replicator executable-prefix=mm

[nyc] topology=clustered master=trainingdb1 members=trainingdb1,trainingdb2,trainingdb3 connectors=trainingdb1,trainingdb2,trainingdb3

topology=clustered master=trainingdb4 Members=trainingdb4,trainingdb5,trainingdb6 connectors=trainingdb4,trainingdb5,trainingdb6

[nyc_london] thl-port=2114 rmi-port=10002 topology=cluster-slave master-dataservice=nyc slave-dataservice=london

[london_nyc]
thl-port=2115
rmi-port=10003
topology=cluster-slave
master-dataservice=london
slave-dataservice=nyc

Composite Active-Active

- · Both Clusters identified as Active clusters
- All writes, from any Connector, will go to Primary host in Active Cluster to which they are associated by default
- In event of a cluster failure, if the Connectors are still operational, they will re-route to another available cluster
- Completely Cluster-aware
- Everything is controlled via cctrl
- One single install no longer needs the Replicator installed separately
- Available from Version 6.0.0 onwards only
- INI Installation method only Staging-method installs are not supported



[defaults]

user=tungsten home-directory=/opt/continuent application-user=app_user application-password=secret application-port=3306 profile-script=~/.bash_profile replication-user=tungsten replication-password=secret mysql-allow-intensive-checks=true

nyc]

topology=clustered
master=trainingdb1
members=trainingdb1,trainingdb2,trainingdb3
connectors=trainingdb1,trainingdb2,trainingdb3

[london]

topology=clustered
master=trainingdb4
members=trainingdb4,trainingdb5,trainingdb6
connectors=trainingdb4,trainingdb5,trainingdb6

[global]

topology=composite-multi-master
composite-datasources=nyc,london

Summary

What we have learnt today

- Reviewed Multi Cluster Use Topologies
- Understood Multi Cluster Architecture
- Explored Configuration Differences

Next Steps

In the next session we will

- Discuss Tungsten Connector (Proxy)
 - Different Connector Configurations
 - Best Practices

THANK YOU FOR LISTENING

continuent.com

Chris Parker, Customer Success Director, EMEA & APAC

