

Tungsten Cluster Master Class

Advanced: Distributed Datasource Groups (DDG)

Chris Parker

VP of Customer Success, EMEA

Topics

In this short course, we will


- Understand what a DDG topology is
- Explore additional considerations
- Study the Configuration
- See an installation in practice

DDG Topology

Tungsten Cluster

Distributed Datasource Groups Topology

- Odd number of nodes
- Single unified cluster, split into logical site-based groups
- Each node assigned unique ID
- Single global primary node
- Local-first failover within same DDG
- Minimizes cross-site latency
- No need to “failover” if entire DDG is lost
- Improved resilience



Configuration Considerations

Installation Considerations

- All usual Pre-Reqs need to be in place
- Network Connectivity
- Network Latency between clusters
- Connector Behaviour

Sample Configs

- One Cluster with one node identified as the primary
- Same configuration as a standalone cluster
- `datasource-group-id` specified to distinguish the ID for each node
 - Same ID value must be used on all nodes within a group, e.g:
 - db1,db2 set to 100
 - db3,db4 set to 101
 - db5 set to 102

```
[defaults]
user=tungsten
install-directory=/opt/continuent
profile-script=~/.bash_profile
connector-bridge-mode=false
application-user=app_user
application-password=secret
application-port=3306
application-readonly-port=3307
replication-user=tungsten
replication-password=secret
replication-port=13306
disable-security-controls=true
rest-apis=true
rest-api-authentication=true
rest-api-ssl=true
rest-api-address=0.0.0.0
rest-api-admin-pass=secret
rest-api-admin-user=tungsten
datasource-group-id=xxx
```

```
[alpha]
topology=clustered
master=db1
members=db1,db2,db3,db4,db5
connectors=db1,db2,db3,db4,db5
```

Summary

What we have learnt today

- Reviewed what DDG is
- Understood configuration and architecture
- Seen a demonstration of it in action

Thank you for listening

continuent.com

Chris Parker

VP of Customer Success, EMEA