
OpenStack Architecture for the Enterprise



Keith Tobin
Cloud Architect

Greg Jacobs
Network Architect

Speaker Profiles

Keith Tobin



- Dell Cloud Architect
- Located in the Dell Cloud Center of Excellence in Dublin, Ireland
- Over two decades of software, platform and cloud development
- Organizer for OpenStack Ireland
- Loves all things OpenStack and Cloud

Greg Jacobs



- Dell Network Architect
- 20 years working as a network architecture and solutions specialist
- Network evangelist and packet head




Design Goals

- Design an OpenStack architecture to run both enterprise and cloud applications
- Meet enterprise performance expectations for VM's, networking and storage
- Zero downtime on the control layer
- No data loss in the event of a single node failure
- Must be durable, reliable and highly available
- Must automatically recover after a fault situation has been fixed
- Minimize management tasks



Solution Stack



Deployment With Crowbar




Deploying OpenStack with Crowbar Is Easy

- Crowbar boots and installs from ISO on a management node (server)
- Provides user with web UI for configuration and operations management
- Discovers all potential targets (servers) automatically
- Performs hardware configuration
 - BIOS configuration
 - Performs updates to BIOS and firmware
- Installs OS on targets
- Deploys OpenStack from a set of configuration modules called barclamps



Crowbar Deployment




Crowbar 2 Status




Load Balancer




Active/Passive (Normal)




Active/Passive (Fault)



Single HAProxy Bottleneck




Dynamic Load Balancing




Neutron



Neutron (Single Networking Node)



Neutron (Separate Networking Nodes)




RabbitMQ




RabbitMQ Cluster (With Mirrored Queues)

Physical View




Logical View



MySQL Database



MySQL Single Server (Normal Operation)



MySQL Single Server (Failed)

Application

DB Server 1
(Active)



MySQL Single Server (Failed)

- **Disadvantage**
 - Single point of failure

Application

DB Server 1
(Active)



Data


MySQL Active/Passive Failover



MySQL Active/Passive Failover (Primary Node Operation)

Disadvantages

- Inactive Server
 - Limits scale
 - Reduces the potential workload processing ability of the solution by half
- Async replication has potential for data inconsistencies across server nodes in fault situations




MySQL With DRDB



MySQL With DRDB (Normal Operation)

Disadvantages


- Inactive server
 - Limits scale
 - Reduces the potential workload processing ability of the solution by half




MySQL Multi-Master



MySQL Multi-Master Cluster




MySQL Multi-Master Cluster (Failed Node)




MySQL Async Replication

- No guarantee that data will be consistent across all nodes at any point in time




MySQL Single Thread replication



MySQL Multi-Master Cluster

Disadvantages

- At any point in time there is no guarantee that data is in a consistent state within the cluster
- One thread allocated for replication tasks



Percona MySQL Cluster




Percona MySQL Cluster

Capability
Synchronous Replication
Data Consistency
Parallel Applying on All Nodes
Automatic Node Provisioning
Multi-Master




Synchronous Replication



Capability	
Synchronous Replication	
Data Consistency	
Parallel Applying on All Nodes	
Automatic Node Provisioning	
Multi-Master	


Data Consistency



At any point in time data is always consistent across all nodes in the cluster


Capability	
Synchronous Replication	
Data Consistency	
Parallel Applying on All Nodes	
Automatic Node Provisioning	
Multi-Master	

Parallel Applying on All Nodes




Capability	
Synchronous Replication	
Data Consistency	
Parallel Applying on All Nodes	
Automatic Node Provisioning	
Multi-Master	

Automatic Node Provisioning



Capability	
Synchronous Replication	
Data Consistency	
Parallel Applying on All Nodes	
Automatic Node Provisioning	
Multi-Master	

Multi-Master Cluster



Capability	
Synchronous Replication	
Data Consistency	
Parallel Applying on All Nodes	
Automatic Node Provisioning	
Multi-Master	


Block Storage




Block Storage



CEPH on R720



Logical Architecture





Thank you

Chat with Keith and Greg @ the Dell booth

Keith Tobin
Cloud Architect

Greg Jacobs
Network Architect

