CentOS Cluster Server

Ryan Matteson
matty91@gmail.com
http://prefetch.net

Presentation overview

- Tonight I am going to give an overview of CentOS cluster server, and describe what is needed to build a basic HA cluster
- This presentation assumes a basic understanding of network and clustering technology, so make sure to ask questions if you aren’t sure about something

What is CentOS cluster server?

- CentOS cluster server is a suite of packages that can be used to deploy highly available services on CentOS Linux-based servers
- Based on Redhat cluster server
- Provides three main features:
  - Cluster management and service failover
  - Network load-balancing (LVS)
  - Global read-write file system (GFS)

What is required to run a cluster?

- Two or more servers that are on the HCL
- Two or more bonded NICs to send cluster heartbeat messages over (this is optional, but highly recommended!)
- Two or more bonded NICs dedicated to public network traffic
- Supported fencing solution
- Shared storage

What does a cluster consist of?

- An HA cluster typically consists of the following items:
  - Two or more nodes
  - One or more fence devices
  - Shared storage
  - Public and private network interfaces
  - One or more resources
  - One or more services
  - Quorum devices
  - Failover Domains

Quorum devices

- Quorum is used to ensure that a majority of nodes are available in the cluster
- Needed to avoid split-brain conditions
- Works by assigning one or more votes to each server and quorum device in the cluster
- To ensure quorum, a cluster needs to have 51% of the available votes to form or continue running an operational cluster
- SCSI disks that support SPR are the most common type of quorum device

Fencing devices

- Fencing devices provide a way for the cluster to remove an unresponsive node from the cluster
- Nodes are typically fenced when they are unresponsive, and fencing is done to prevent split brain configurations
- Several supported ways to fence nodes:
  - IPMI
  - Power Fencing
  - SAN fencing
  - VMWare virtual center fencing
  - Vendor specific methods (HP ILO, Dell DRAC, etc.)

Cluster resources

- Cluster resources provide the basic unit of configuration in a cluster
- Several types of resources exist by default:
  - Apache
  - GFS
  - MySQL
  - Oracle
  - Samba
  - NFS
  - Tomcat
  - Virtual machines

Cluster services

- Services are collections of resources that serve a specific purpose
- An example of this would be an HA MySQL service that contains three resources:
  - An IP address resource that is tied to the MySQL database instance
  - File system resources that contain the data and indexes needed by the database
  - A MySQL resource that starts, stops and verifies that mysql is running

Failover domains

- Failover domains allow you to define where services should go when a service faults and is migrated to another node
- Each failover domain can have a unique list of nodes, and each node can be assigned a priority to tell the cluster it is a better candidate to run the service

How do I install CCS?

- Verify your hardware meets the hardware guidelines in the CCS manuals
- Install CentOS on each node
- Install the clustering software on each node
- Create the cluster
- Add fence devices
- Add quorum devices if needed
- Create resources, services and failover domains
- Test, test and test some more!!

Installing the cluster software

- To install CentOS cluster server you can run yum groupinstall on each node in the cluster:

  $ yum groupinstall "Cluster Storage" "Clustering”

- If the software isn’t already installed on a node, the cluster will install the required packages when you add the node to the cluster

Creating a cluster

- You can create the cluster in one of three ways
  - Create the cluster configuration by hand
  - Run *system-config-cluster*
  - Use the conga web interface
- Once the cluster has been created, you can add fence devices, resources, services and failover domains using one of the methods listed above

Cluster configuration

- The cluster configuration is stored in /etc/cluster/cluster.xml on each node
- Each tag in the cluster.xml file contains a configuration entity, such as the name of a node in the cluster, the fence device to use for each node, and a list of resources, services and failover domains

Example cluster.xml

Cluster utilities

- There are a number of utilities that can be used to manage a cluster:
  - clustat – displays cluster status
  - clusvcadm – controls cluster services
  - ccs_tool – manages the cluster configuration
  - cman_tool – manages the cluster members
  - fence_tool – manages fencing operations
  - mkqdisk – manages quorum disks

Cluster processes

- There are a number of critical processes that make up the cluster suite:
  - cman – controls overall cluster operation
  - fenced – manages fencing operations
  - clurgmgrd – controls services
  - various kernel threads (visible in ps)
  - Application processes (e.g., httpd)
  - Several more not mentioned here …

Debugging cluster problems

- If your cluster is acting up, you will want to review the logging data in /var/log/* to see what is going on
- Debug stanzas can be added to each cluster facility to get additional debugging data:

```
<logger debug="on" ident=”CMAN” to_stderr="yes"/>
```

- The Redhat bugzilla archives are a great resource for finding solutions to problems, and for troubleshooting sporadic issues

Conclusion

- CentOS cluster server has a number of cool features, and won't cost you a dime to deploy (you don't get support though)
- If you decide to use CCS, make SURE you have approved hardware and fencing devices. If you don't, you are asking for trouble (and data loss!)!

Questions?

References

- Redhat cluster suite overview:

  http://www.redhat.com/docs/manuals/csgfs/
- Configuring and managing a RHEL cluster:

  http://www.redhat.com/docs/manuals/csgfs/
- Quorum devices

  http://magazine.redhat.com/2007/12/19/enhancing-cluster-quorum-with-qdisk/