OCFS2 Cluster File System Setup Guide in Linux

Article Number: 149 | Rating: Unrated | Last Updated: Sat, Jun 2, 2018 9:19 PM

OCFS2 Cluster File System Setup Guide in Linux

What is OCFS2?

- OCFS2 is a shared-disk cluster file system for Linux
- Capable of providing both high performance and high availability.
- Cluster-aware applications can make use of parallel I/O for higher performance with this FS
- OCFS2 is mostly used to host Oracle Real application clusters (RAC) database on Linux clusters.

The below are the high level steps for creating ocfs2 filesystem on top of a multipath'd SAN LUN:

1. Verify the nodes that will be part of your cluster.
2. Make sure your LUNs on the SAN end are accessible on all the nodes of the cluster.
3. If you need multipathing, configure multipath and the multipathing policy based on solution whatever you have.
4. The following example configuration (**/etc/ocfs2/cluster.conf**) shows you a sample

configuration of a 2 node cluster pool.

5. If you have heartbeat IP configured on these cluster nodes, use the heartbeat IP for ocfs2 cluster communication and specify the hostname without FQDN.
6. Copy the same file to all the hosts in the cluster.

Sample Config file:

```
[root@rac-cluster ~]# cat /etc/ocfs2/cluster.conf
node:
  ip_port = 7777
  ip_address = 192.168.1.2
  number = 0
  name = rac-cluster
  cluster = ocfs2
node:
  ip_port = 7777
  ip_address = 192.168.1.3
  number = 1
  name = rac-cluster
  cluster = ocfs2
```

On each node check the status of OCFS2 cluster service and stop "o2cb" if the service is already running.

 # service o2cb status # service o2cb stop

On each node, load the OCFS2 module.

 # service o2cb load

Make the OCFS2 service online on all the nodes.

 # service o2cb online

Now your OCFS2 cluster is ready.

Format the SAN lun device from any one of the cluster node.

 # mkfs.ocfs2 -b 4k -C 32k -L oraclerac /dev/emcpowera

Where,

-b : Block size
-C : Cluster size
-L : Label

Update /etc/fstab on all the nodes in the cluster with the mount point.

/dev/emcpowera /u01 ocfs2 _netdev 0 0

Mount the /u01

Enable ocfs and o2b service on boot

 # chkconfig --level 345 o2cb on# chkconfig --level 345 ocfs2 on

The /u01 repository setup on a SAN Lun is done.

You can now configure Oracle RAC on this file System.

Posted - Sat, Jun 2, 2018 9:19 PM. This article has been viewed 7474 times.

Online URL: http://kb.ictbanking.net/article.php?id=149