Role STORE - DRBD File System (Redundant with LVM)

Create LVM Physical Volume, Volume Group and Logical Volume (Both Nodes)

The commands below assume that **/dev/sdb** will be used for the DRBD on top of LVM configuration, and that the disks are **EXACTLY** the same size.

```
# Create the phsyical volume - this is based on sdb assuming it is the second drive on the system
lvm pvcreate /dev/sdb

# Create the volume group
lvm vgcreate "vg_drbd_jtelshared" /dev/sdb

# Create the logical volume
lvm lvcreate -l +100%FREE vg_drbd_jtelshared -n lv_drbd_jtelshared
```

Configure Firewall (Both Nodes)

```
# Prepare the firewall
firewall-cmd --zone=public --add-port=7788-7799/tcp --permanent
firewall-cmd --reload
```

Install elrepo to yum (Both Nodes)

```
# Install elrepo
rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
rpm -Uvh http://www.elrepo.org/elrepo-release-7.0-3.el7.elrepo.noarch.rpm
```

Install and enable DRBD (Both Nodes)

```
# Install DRBD
yum -y install drbd84-utils kmod-drbd84 lsof
# Enable drbd at boot, and in this session
systemctl enable drbd
systemctl start drbd
```

Configure DRBD (Both Nodes)

NOTE: The following commands requires the hostname of both machines and the IP Address. These are obtained as follows:

```
ip addr
hostname
```

Create a DRBD config file for jtelshared on /dev/sdb

/etc/drbd.d/jtelshared.res

```
# Configure DRBD
cat <<EOFF > /etc/drbd.d/jtelshared.res
resource jtelshared {
    protocol C;
    meta-disk internal;
    device /dev/drbd0;
    syncer {
        verify-alg sha1;
    }
    net {
        allow-two-primaries;
    }
    on acd-store1.jtel.local {
        disk   /dev/vg_drbd_jtelshared/lv_drbd_jtelshared;
        address 10.42.14.98:7789;
    }
    on acd-store2.jtel.local {
        disk   /dev/vg_drbd_jtelshared/lv_drbd_jtelshared;
        address 10.42.14.198:7789;
    }
    startup {
        become-primary-on both;
    }
}
EOFF
```

Note: it has been observed, that the fully qualified host name is required in the configuration file.

Create Metadata and start (Both Nodes)

```
# Create metadata and start DRBD
drbdadm create-md jtelshared
drbdadm up jtelshared
```

Make one node primary (First Node)

```
drbdadm primary jtelshared --force
```

Tune the transfer (Second Node)

```
drbdadm disk-options --c-plan-ahead=0 --resync-rate=110M jtelshared
```

Wait for initial sync to complete (Either Node)

```
cat /proc/drbd

-->

# When not yet done:

version: 8.4.10-1 (api:1/proto:86-101)
GIT-hash: a4d5de01fffd7e4cde48a080e2c686f9e8cebf4c build by mockbuild@, 2017-09-15 14:23:22

 1: cs:SyncTarget ro:Secondary/Primary ds:Inconsistent/UpToDate C r-----
    ns:0 nr:3955712 dw:3950592 dr:0 al:8 bm:0 lo:5 pe:0 ua:5 ap:0 ep:1 wo:f oos:264474588
        [>...................] sync'ed:  1.5% (258272/262132)M
        finish: 2:08:08 speed: 34,388 (25,652) want: 112,640 K/sec


-->

# When done:

version: 8.4.10-1 (api:1/proto:86-101)
GIT-hash: a4d5de01fffd7e4cde48a080e2c686f9e8cebf4c build by mockbuild@, 2017-09-15 14:23:22

 1: cs:Connected ro:Secondary/Primary ds:UpToDate/UpToDate C r-----
    ns:0 nr:15626582 dw:284051762 dr:0 al:8 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:0
```

Untune the transfer and make primary (Second Node)

```
drbdadm adjust jtelshared
drbdadm primary jtelshared
```

Create filesystem (First Node)

```
mkfs.xfs -L data /dev/drbd/by-res/jtelshared/0
```

Create fstab entry for file system (Both Nodes)

Add the following line to **/etc/fstab**

```
/dev/drbd/by-res/jtelshared/0  /srv/jtel/shared         xfs noauto,noatime,nodiratime  0   0
```

Mount the file system (First Node)

```
mkdir /srv/jtel
mkdir /srv/jtel/shared
chown -R jtel:jtel /srv/jtel
mount /srv/jtel/shared
```

Create a test file and Unmount (First Node)

```
cat <<EOFF > /srv/jtel/shared/test.txt
test 123
EOFF
umount /srv/jtel/shared
```

Mount the file system and check the test file (Second Node)

```
mkdir /srv/jtel
mkdir /srv/jtel/shared
chown -R jtel:jtel /srv/jtel
mount /srv/jtel/shared
cat /srv/jtel/shared/test.txt

# Check contents of file before proceeding

rm /srv/jtel/shared/test.txt
umount /srv/jtel/shared
```

Install Samba and lsof (Both Nodes)

Install SAMBA

```
yum -y install samba samba-client lsof
```

Configure Samba (Both Nodes)

Configure SAMBA

```
cat <<EOFF > /etc/samba/smb.conf
[global]
        workgroup = SAMBA
        security = user
        passdb backend = tdbsam
        printing = cups
        printcap name = cups
        load printers = yes
        cups options = raw
        min protocol = NT1
        ntlm auth = yes

[homes]
        comment = Home Directories
        valid users = %S, %D%w%S
        browseable = No
        read only = No
        inherit acls = Yes

[printers]
        comment = All Printers
        path = /var/tmp
        printable = Yes
        create mask = 0600
        browseable = No

[print$]
        comment = Printer Drivers
        path = /var/lib/samba/drivers
        write list = root
        create mask = 0664
        directory mask = 0775

[shared]
    comment = jtel ACD Shared Directory
    read only = no
    public = yes
    writable = yes
    locking = yes
    path = /srv/jtel/shared
    guest ok = yes
    create mask = 0644
    directory mask = 0755
    force user = jtel
    force group = jtel
    acl allow execute always = True

EOFF
```

Setup SeLinux, jtel User access and Firewall for Samba (Both Nodes)

Replace <password> with the actual password for the jtel user:

SeLinux, jtel User, Firewall

```
setsebool -P samba_enable_home_dirs=on samba_export_all_rw=on use_samba_home_dirs=on use_nfs_home_dirs=on
printf '<password>\n<password>\n' | smbpasswd -a -s jtel
firewall-cmd --zone=public --add-port=445/tcp --add-port=139/tcp --add-port=138/udp --add-port=137/udp --permanent
firewall-cmd --reload
```

If necessary, add further users to samba:

More SAMBA users

```
useradd -m Administrator
printf 'F1r3B²11\nF1r3B²11\n' | smbpasswd -a -s Administrator
```

Test SAMBA (Both Nodes)

This test should be performed on the node which currently has /srv/jtel/shared mounted:

Test SAMBA

```
mount /srv/jtel/shared
service nmb start
service smb start


# Now check access to the SMB share via (for example) one of the windows machines.

service smb stop
umount /srv/jtel/shared

# do same again on other node
```

Unmount and disable SAMBA (Both Nodes)

Unmount

```
service smb stop
umount /srv/jtel/shared
systemctl disable smb
```

Install PCS Services (Both Nodes)

See Redundancy - Installing PCS Cluster.

Setup virtual IP (One Node Only!)

Change the following to set the virtual IP which should be shared between the nodes.

Set virtual IP

```
KE_VIP=10.4.8.22
```

Configure PCS Resources (One Node Only!)

Configure the PCS resources with the following commands:

Configure PCS Resources

```
pcs resource create ClusterDataJTELSharedMount ocf:heartbeat:Filesystem device="/dev/drbd/by-res/jtelshared/0" directory="/srv/jtel/shared" fstype="xfs" --
group=jtel_portal_group
pcs resource create ClusterIP ocf:heartbeat:IPaddr2 ip=${KE_VIP} cidr_netmask=32 op monitor interval=30s --group=jtel_portal_group
pcs resource create samba systemd:smb op monitor interval=30s --group=jtel_portal_group
pcs constraint order start ClusterDataJTELSharedMount then ClusterIP
pcs constraint order start ClusterIP then samba
```

Test

Test as follows:

Test pcs status

```
pcs status

--> shows the status of the newly created resources on both nodes, one node should be active.

Cluster name: portal
Stack: corosync
Current DC: uk-acd-store2 (version 1.1.16-12.el7_4.8-94ff4df) - partition with quorum
Last updated: Mon Mar 19 15:40:24 2018
Last change: Mon Mar 19 15:40:16 2018 by root via cibadmin on uk-acd-store1

2 nodes configured
3 resources configured

Online: [ uk-acd-store1 uk-acd-store2 ]

Full list of resources:

 Resource Group: jtel_portal_group
     ClusterDataJTELSharedMount (ocf::heartbeat:Filesystem):    Started uk-acd-store1
     ClusterIP  (ocf::heartbeat:IPaddr2):       Started uk-acd-store1
     samba      (systemd:smb):  Started uk-acd-store1

Daemon Status:
  corosync: active/enabled
  pacemaker: active/enabled
  pcsd: active/enabled
```

Test the file mount:

Test file mount

```
# From the windows machines:

dir \\uk-acd-store\shared
```

Test manual failover:

Test file mount

```
# Failover to node 2
pcs cluster standby uk-acd-store1

# ... (wait)

pcs status

# Then test the availability of the files from the windows machines.
# Create a new file before failing back (to make sure DRBD working ok).

# Fail back to node 1
pcs cluster unstandby uk-acd-store1
pcs cluster standby uk-acd-store2

# ... (wait)

pcs status

# Then test the availability of the files from the windows machines.
# Check that the new file created above is available.

# Unstandby node 2

pcs cluster unstandby uk-acd-store2
```

Manually link /home/jtel/shared (Both Nodes)

link /home/jtel/shared

```
ln -s /srv/jtel/shared /home/jtel/shared
```