

Role STORE - Redundant (Debian/Win2019)

Notes

We no longer support DRBD on the base file system. We always install with LVM for maintenance purposes.

We use the whole disk, so increasing size is done using LVM by adding new disks.

It is recommended to create the STORE machine without the disk for the storage being mounted by the installation routines.

Disk and DRBD Setup

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 physical 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 for DRBD (**Both Nodes**)

```
# Prepare the firewall
ufw allow 7788:7799/tcp
```

Prepare Mount Point (**Both Nodes**)

The data should be mounted to the directory `/srv/jtel/shared`.

The following commands prepare for this:

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

Install DRBD (Both Nodes)

We now install DRBD. The kernel package is included in debian, but the tools must be installed.

```
apt-get -y install drbd-utils
```

Configure DRBD (Both Nodes)

DRBD must be configured with static ip addresses and correct hostnames.

The IP addresses below must be modified:

CAUTION IP ADDRESSES **CAUTION HOSTNAMES**

```
# 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-lb1.jtel.local {
        disk   /dev/vg_drbd_jtelshared/lv_drbd_jtelshared;
        address 10.1.1.1:7789;
    }
    on acd-lb2.jtel.local {
        disk   /dev/vg_drbd_jtelshared/lv_drbd_jtelshared;
        address 10.1.1.2:7789;
    }
}
EOFF
```

Global configuration (note, the c-max-rate is good for a 1Gbit network. You might want to change this.

```
cp /etc/drbd.d/global_common.conf /etc/drbd.d/global_common.conf.orig
cat << EOFF > /etc/drbd.d/global_common.conf
global {
    usage-count no;
    udev-always-use-vnr;
}

common {
    handlers {
    }

    startup {
    }

    options {
    }

    disk {
        c-plan-ahead 10;
        c-fill-target 24M;
        c-min-rate 10M;
        c-max-rate 100M;
    }
}

net {
    max-buffers 36k;
    sndbuf-size 1024k;
    rcvbuf-size 2048k;
}
}
EOFF
```

Start and Enable Kernel Module (**Both Nodes**)

```
modprobe drbd
systemctl enable drbd
systemctl start drbd
```

Create Metadata and Start (**Both Nodes**)

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

Make ONE Node Primary

```
# Make ONE node primary  
drbdadm primary jtelsshared --force
```

Wait for Sync

DRBD will now sync. This might take some time.

Note: with DRBD9 we currently have no options to tune the transfer.

You can watch the initial sync with the following command:

```
drbdadm status jtelsshared
```

You will see output like this:

```
jtelsshared role:Primary  
disk:UpToDate  
acd-store2 role:Secondary  
replication:SyncSource peer-disk:Inconsistent done:7.19
```

This means the following:

- The local machine is primary.
- The disk in the local machine is up to date.
- acd-store2 is secondary
- It is inconsistent and syncing, 7.19% done

Do not continue until this step is complete.

Make Second Node Primary (**SECOND NODE ONLY**)

```
drbdadm primary jtelsshared
```

Create the Filesystem (**FIRST NODE ONLY**)

```
mkfs.ext4 /dev/drbd/by-res/jtelsshared/0
```

Create fstab entry for file system (**Both Nodes**)

This command adds a line to **/etc/fstab**

```
cat << EOF >> /etc/fstab
/dev/drbd/by-res/jtelshared/0  /srv/jtel/shared          ext4 noauto,noatime,nodiratime 0 0
EOF
```

Test DRBD

Now, we can test the DRBD setup.

Mount the file system (**FIRST Node**)

```
mount /srv/jtel/shared
```

Create a test file and Unmount (**FIRST Node**)

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

Mount the file system and check test file (**SECOND Node**)

```
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
```

Do not proceed unless you can see the contents of the test file.

Comment Mount out in fstab (**BOTH nodes**) and disable DRBD

```
sed -i '/jtelshared/s/^/#/' /etc/fstab
systemctl disable drbd
umount /srv/jtel/shared
```

Install PCS Cluster (**BOTH NODES**)

If you have not installed Pacemaker / Corosync on both LB machines, do this now - see here: [Redundancy - Installing PCS Cluster](#)

Install and Configure Samba

Installation (**BOTH NODES**)

These commands install the samba server and client and lsof.

```
apt-get -y install samba samba-client lsof
```

Next disable smbd (this will be managed by the pcs cluster):

```
systemctl stop smbd
systemctl disable smbd
```

Configure Samba (**BOTH NODES**)

The following creates a samba configuration file with a minimum configuration.

```
# SMB Conf
cat <<EOFF > /etc/samba/smb.conf
[global]
    workgroup = JTEL
    security = user
    passdb backend = tdbsam
    min protocol = SMB2
    reset on zero vc = yes
[shared]
    comment = jtel ACD Shared Directory
    public = no
    read only = no
    writable = yes
    locking = yes
    path = /srv/jtel/shared
    create mask = 0644
    directory mask = 0755
    force user = jtel
    force group = jtel
    acl allow execute always = True
EOFF
```

Setup the Firewall (**BOTH NODES**)

The following command sets up the firewall:

```
ufw allow 445/tcp
```

Link /home/jtel/shared (**BOTH NODES**)

Link the /home/jtel/shared folder.

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

Setup Access to the Samba Server

jtel User Access (**BOTH NODES**)

The following command creates the smb credentials for the jtel user.

CAUTION PASSWORD

```
printf '<password>\n<password>\n' | smbpasswd -a -s jtel
```

Further User Access (**BOTH NODES**)

If necessary, add further users to samba - replacing password with the actual password for the user. Here, for example, the windows administrator user:

CAUTION PASSWORD

```
useradd -m Administrator
printf '<password>\n<password>\n' | smbpasswd -a -s Administrator
```

Configure Cluster Resources

Now all resources will be configured in the pacemaker cluster.

Setup virtual IP (One Node Only!)

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

CAUTION IP ADDRESS

```
JT_VIP=10.1.1.100
```

Configure PCS Resources for DRBD Mount, DRBD Primary / Secondary, Samba and Virtual IP Address (**One Node Only!**)

Configure the PCS resources with the following commands:

```
# Configure using a file jtel_cluster_config
cd
pcs cluster cib jtel_cluster_config
# DRBD Primary Secondary
pcs -f jtel_cluster_config resource create DRBDClusterMount ocf:linbit:drbd drbd_resource=jtelshared op monitor interval=60s
pcs -f jtel_cluster_config resource promotable DRBDClusterMount promoted-max=1 promoted-node-max=1 clone-max=2 clone-node-max=1 notify=true
# DRBD File System Mount
pcs -f jtel_cluster_config resource create DRBDClusterFilesystem ocf:heartbeat:Filesystem device="/dev/drbd/by-res/jtelshared/0" directory="/srv/jtel/shared"
fstype="ext4"
# Colocation of File System Mount with Primary DRBD instance
pcs -f jtel_cluster_config constraint colocation add DRBDClusterFilesystem with DRBDClusterMount-clone INFINITY with-rsc-role=Master
# Promote first, then start filesystem
pcs -f jtel_cluster_config constraint order promote DRBDClusterMount-clone then start DRBDClusterFilesystem
# Resource for Samba
pcs -f jtel_cluster_config resource create Samba systemd:smbd op monitor interval=30s
# Resource for virtual IP
pcs -f jtel_cluster_config resource create ClusterIP ocf:heartbeat:IPaddr2 ip=${JT_VIP} cidr_netmask=32 op monitor interval=30s
# Samba must be with active DRBD filesystem
pcs -f jtel_cluster_config constraint colocation add Samba with DRBDClusterFilesystem INFINITY
# Cluster IP must be with Samba
pcs -f jtel_cluster_config constraint colocation add ClusterIP with Samba INFINITY
# Start DRBD File system then start Samba
pcs -f jtel_cluster_config constraint order DRBDClusterFilesystem then Samba
# Start Samba then start Cluster IP
pcs -f jtel_cluster_config constraint order Samba then ClusterIP
```

Check the configuration:

```
# Check the config file
pcs -f jtel_cluster_config config
```

Push the configuration to the cluster:

```
# Push the config to the cluster  
pcs cluster cib-push jtel_cluster_config --config
```

Ensure ownership of jtel directory:

```
chown -R jtel:jtel /srv/jtel
```

Tests

Test pcs status

First of all, we test the cluster status:

```
pcs status
```

You should see output similar to this:

```
Cluster name: jtel_cluster
Cluster Summary:
  * Stack: corosync
  * Current DC: acd-lb1 (version 2.0.3-5.el8_2.1-4b1f869f0f) - partition with quorum
  * Last updated: Sat Oct  3 12:59:34 2020
  * Last change: Sat Oct  3 12:31:22 2020 by root via cibadmin on acd-lb2
  * 2 nodes configured
  * 5 resource instances configured

Node List:
  * Online: [ acd-lb1 acd-lb2 ]

Full List of Resources:
  * Clone Set: DRBDClusterMount-clone [DRBDClusterMount] (promotable):
    * Masters: [ acd-lb1 ]
    * Stopped: [ acd-lb2 ]
  * DRBDClusterFilesystem      (ocf::heartbeat:Filesystem):     Started acd-lb1
  * Samba        (systemd:smb):     Started acd-lb1
  * ClusterIP   (ocf::heartbeat:IPAddr2):     Started acd-lb1

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

Make sure all of the resources are started and both nodes are online.

Test File Mount

You should now be able to access **\acd-store\shared** from the windows machines for example.

If you want to test from linux, you will need to mount STORE as described here: [Mounting STORE - All Linux except for STORE \(CentOS8/Win2019\)](#)

Test Failover and Fallback

You can test failover and fallback with any of the following commands:

Standby and Unstandby

Caution: standby and unstandby have been observed to not failover the resources correctly.

Use with caution.

```
pcs node standby acd-lb1
# TEST
pcs node unstandby acd-lb1
# TEST
pcs node standby acd-lb2
# TEST
pcs node unstandby acd-lb2
# TEST
```

Stop Cluster Node

```
pcs cluster stop acd-lb1
# TEST
pcs cluster start acd-lb1
# TEST
pcs cluster stop acd-lb2
# TEST
pcs cluster start acd-lb2
# TEST
```

Reboot

Rebooting is also a good way to test.

Power Off

This is the best way to test, but be aware, you may cause split brain on DRBD and need to repair it.

Debian 11 - Possible ruby problem

It has been observed, that the following file contains errors, even if the cluster "appears" to be working properly.

```
less /var/log/pcsd/pcsd.log

-->

E, [2022-04-03T00:06:25.007 #43472]     ERROR -- : Unable to connect to node acd-store4, the node is not known
E, [2022-04-03T00:06:25.007 #43472]     ERROR -- : Unable to connect to node acd-store3, the node is not known
```

This can be fixed as follows:

```
# ON BOTH NODES

# Install missing ruby library
gem install orderedhash

# Unauthorize the cluster
pcs pcsc deauth

# ON ONE NODE

# Authorize the cluster
pcs cluster auth -u hacluster -p <password>

# CHECKS
less /var/log/pcsd/pcsd.log

##### Log should look "normal"

pcs status

##### Cluster should look „normal“
```