

Cloning the XNAT Relays

- Cloning
 - Configure BIOS
 - Boot a Cento 7 Live DVD
 - Create /etc/zfs/vdev_id.conf
 - Activate vdevs
 - Partition the SSDs
 - Create rpool
 - Populate rpool from image
 - Remount in proper order
 - Prepare rpool to boot on new host
 - Set the hostname
 - Prepare the relaypool
 - Populate the relaypool
 - Clean Up
 - Leave chroot and Copy vdev_id.conf
 - Reboot via hard reset
- TODO:

Cloning

Configure BIOS

Use ["Building the XNAT Relays"](#) as a reference

Boot a Cento 7 Live DVD

Install ZFS

```
yum -y install http://download.zfsonlinux.org/epel/zfs-release.el7_3.noarch.rpm
sed -i -e '/\[zfs\]/,/^\[s/enabled=1/enabled=0/' /etc/yum.repos.d/zfs.repo
sed -i -e '/\[zfs-kmod\]/,/^\[s/enabled=0/enabled=1/' /etc/yum.repos.d/zfs.repo
yum -y install zfs zfs-dracut pigz
modprobe zfs
```

Create /etc/zfs/vdev_id.conf

with the following contents:

vdev_id.conf for Supermicro 5028A-TN4

```
alias disk1 pci-0000:00:17.0-ata-1.0
alias disk2 pci-0000:00:17.0-ata-2.0
alias disk3 pci-0000:00:17.0-ata-3.0
alias disk4 pci-0000:00:17.0-ata-4.0
alias ssd1 pci-0000:00:18.0-ata-1.0
alias ssd2 pci-0000:00:18.0-ata-2.0
```

vdev_id.conf for Intel NUC NUC6i3SYK

```
alias ssd1 pci-0000:00:17.0-ata-3.0
```

Alternate method to populate vdev_id.conf

```
wget -qO- "ftp://ftp.nrg.wustl.edu/pub/XNAT_relay/vdev_id.conf.5028A-TN4" > /etc/zfs/vdev_id.conf
udevadm trigger
ls -l /dev/disk/by-vdev/
```

Activate vdevs

```
udevadm trigger
```

Partition the SSDs

Partition for Supermicro 5028A-TN4

```
sgdisk -al -n2:34:2047 -t2:EF02 /dev/disk/by-vdev/ssd1
sgdisk -n9:-8M:0 -t9:BF07 /dev/disk/by-vdev/ssd1
# rpool
sgdisk -n1:0:40G -t1:BF01 /dev/disk/by-vdev/ssd1
# cache
sgdisk -n3:0:+40G -t3:BF01 /dev/disk/by-vdev/ssd1
# swap
sgdisk -n4:0:+8G -t4:8200 /dev/disk/by-vdev/ssd1

sgdisk -al -n2:34:2047 -t2:EF02 /dev/disk/by-vdev/ssd2
sgdisk -n9:-8M:0 -t9:BF07 /dev/disk/by-vdev/ssd2
# rpool
sgdisk -n1:0:40G -t1:BF01 /dev/disk/by-vdev/ssd2
# cache
sgdisk -n3:0:+40G -t3:BF01 /dev/disk/by-vdev/ssd2
# swap
sgdisk -n4:0:+8G -t4:8200 /dev/disk/by-vdev/ssd2

mkswap /dev/disk/by-vdev/ssd1-part4
mkswap /dev/disk/by-vdev/ssd2-part4
```

Partition for single 500GB SSD

```
sgdisk -al -n2:34:2047 -t2:EF02 /dev/disk/by-vdev/ssd1
sgdisk -n9:-8M:0 -t9:BF07 /dev/disk/by-vdev/ssd1
# rpool
sgdisk -n1:0:40G -t1:BF01 /dev/disk/by-vdev/ssd1
# relaypool
sgdisk -n3:0:+400G -t3:BF01 /dev/disk/by-vdev/ssd1
# swap
sgdisk -n4:0:+8G -t4:8200 /dev/disk/by-vdev/ssd1

mkswap /dev/disk/by-vdev/ssd1-part4
```

Create rpool

Create rpool

```
zpool create -f -d \
    -o feature@async_destroy=enabled \
    -o feature@empty_bpobj=enabled \
-o feature@lz4_compress=enabled \
    -o ashift=12 \
    -O compression=lz4 \
    -O canmount=off \
-O mountpoint=/ \
-R /mnt/rpool \
    rpool \
mirror /dev/disk/by-vdev/ssd1-part1 /dev/disk/by-vdev/ssd2-part1
```

Populate rpool from image

Substitute URL as necessary:

```
URL="ftp://ftp.nrg.wustl.edu/pub/XNAT_relay/relay-server_rpool-to-clone_2017-02-14.zfs.gz"
wget -qO- "$URL" | unpigz | zfs receive -vF rpool
```

Remount in proper order

```
zfs umount -a && rm -rf /mnt/rpool
zfs mount rpool/ROOT/centos
zfs mount -a
```

Prepare rpool to boot on new host

```
for dir in proc sys dev;do mount --rbind /$dir /mnt/rpool/$dir;done

chroot /mnt/rpool /bin/bash --login

cd /dev;ln -s /dev/disk/by-vdev/* .;cd

rm -f /etc/zfs/zpool.cache
zpool set cachefile=/etc/zfs/zpool.cache rpool

version=`rpm -qa kernel|sort|tail -1|cut -c 8-`
dracut -f -v /boot/initramfs-${version}.img ${version}

grub2-mkconfig -o /boot/grub2/grub.cfg
grep ROOT /boot/grub2/grub.cfg
grub2-set-default 0
grub2-install /dev/disk/by-vdev/ssd1
grub2-install /dev/disk/by-vdev/ssd2
```

Set the hostname

Adjust as necessary

```
hostname='site-relay1'
echo $hostname > /etc/hostname
```

Prepare the relaypool

```
zpool create -f relaypool raidz2 disk1 disk2 disk3 disk4
zfs set compression=lz4 relaypool
zfs set sync=disabled relaypool
zfs set quota=7T relaypool
zfs set edu.wustl.nrg:quotareports=2 relaypool
zfs set edu.wustl.nrg:quotareport:1="2T|warning|relay-ops@nrg.wustl.edu|2d" relaypool
zfs set edu.wustl.nrg:quotareport:2="512G|critical|relay-ops@nrg.wustl.edu|12h" relaypool
zpool add relaypool cache ssd1-part3
zpool add relaypool cache ssd2-part3
```

Populate the relaypool

Substitute URL as necessary:

```
URL="ftp://ftp.nrg.wustl.edu/pub/XNAT_relay/relay-server_relaypool-to-clone_2017-02-14.zfs.gz"
wget -qO- "$URL" | unpigz | zfs receive -vF relaypool
```

Clean Up

```
snaps=`zfs list -H -o name -t snapshot -r|grep clone`;for snap in $snaps ;do echo $snap; zfs destroy $snap;done
rm -rf /tmp/*
```

Leave chroot and Copy vdev_id.conf

```
exit
cp /etc/zfs/vdev_id.conf /mnt/rpool/etc/zfs/vdev_id.conf
```

Reboot via hard reset

A proper reboot is not possible with mapped /dev. ZFS does not care.

System will get IP via DHCP

[Continue to Post Clone Configuration](#)

TODO:

Setup Nagios Monitoring:

<http://www.admin-magazine.com/Archive/2014/22/Nagios-Passive-Checks>