

Building the XNAT Relays

- [Server Setup](#)
 - [Hardware](#)
 - [BIOS Settings](#)
 - [Operating System](#)
 - [CentOS 7 - Graphical Server install](#)
 - [ZFS](#)
 - [ZFS root](#)
 - [Enter BIOS and change boot to ZFS root disk](#)
 - [Confirm zfs root is running properly](#)
 - [Setup mirror to original SSD](#)
 - [Addition server configuration](#)

Update procedure:

- [TODO:](#)
- [References](#)

This is a verbose documentation of the build process used to assemble and build the DICOM and RAW data relays used for the Lifespan studies.


Server Setup

Hardware

1. Set boot disk(s) to the SSDs and disable spinning disks from booting
2. Label disks using [XNAT DICOM/RAW Relay Build](#)
3. Label network ports [XNAT DICOM/RAW Relay Build](#)

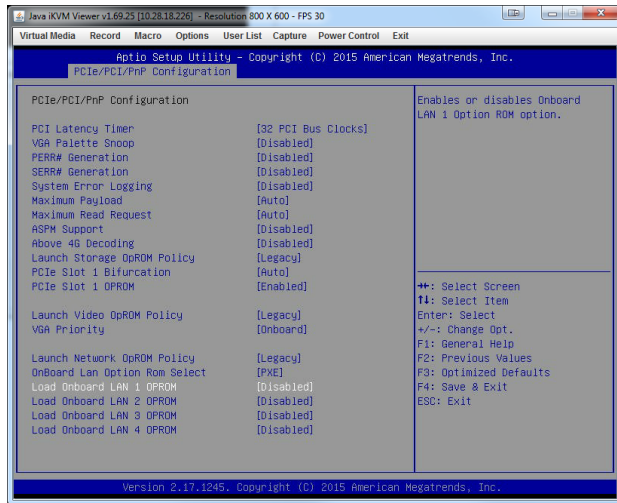
BIOS Settings

Boot Features




PCIe/PCI/PnP Configuration

- Turn off LAN OPROM on all ports

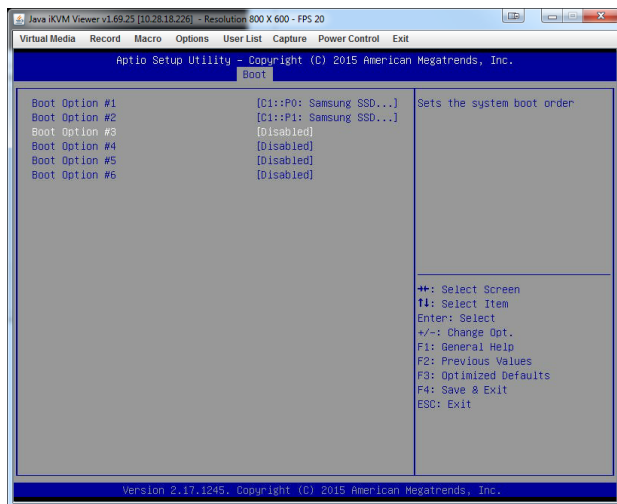


Set the boot order



Set Hard Drive BBS Priorities

Set boot to the C1::P0: and C1::P1: attached SSD and disable booting from the data disks.



Operating System

1. CentOS 7 - Graphical Server install

- a. Use manual partitioning
 - i. Select SSD and auto generate legacy layout
 - ii. Delete the /home partition (This space will be used for ZFS cache)
- b. Set timezone to US/Central (Chicago)
 - i. Turn on network time
- c. Create hcpadmin user
- d. Set root password (needed if runlevel 1 is necessary)
- e. Install additional packages

Packages

```
yum install httpd tomcat mutt pgadmin3 httpd-devel gcc gcc-c++ make libtool apxs mod_ssl xinetd tree
```

2. ZFS

- a. Install ZFS

Install ZFS

```
yum -y install epel-release
yum -y install http://download.zfsonlinux.org/epel/zfs-release.el7_3.noarch.rpm

# Disable DKMS and enable KMOD
vim /etc/yum.repos.d/zfs.repo

yum -y install zfs zfs-dracut
yum update
```

- b. Reboot to apply updated kernel
- c. Create /etc/zfs/vdev_id.conf with the following contents:

vdev_id.conf

```
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
```

- d. Reboot to populate /dev/disk/by-vdev
- e. Create the zpool and set it up.

```

Create zpool

zpool create -f relaypool raidz2 disk1 disk2 disk3 disk4
zfs set compression=lz4 relaypool
zfs set sync=disabled relaypool
zfs create relaypool/zfs_tools
zfs create -o mountpoint=/raw relaypool/raw
zfs create relaypool/raw/data
zfs create relaypool/raw/scripts
zfs create -o mountpoint=/xnat relaypool/xnat
zfs create relaypool/xnat/archive
zfs create relaypool/xnat/prearchive
zfs create relaypool/xnat/tmp
zfs create relaypool/xnat/build
zfs create relaypool/xnat/cache
zfs create relaypool/xnat/logs
zfs create relaypool/xnat/home
zfs create -p relaypool/xnat/pgsql/9.4
zfs create relaypool/rpool_backup
zfs set reservation=100G relaypool/zfs_tools

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
# Make pool auto import and mount on boot
systemctl enable zfs-import-cache.service
systemctl enable zfs-mount.service
# Enable ZFS event daemon
systemctl enable zfs-zed.service

```

- f. Reboot the system so the kernel is familiar with the changes.
- g. Limit ARC to 1/2 the system ram. (Assuming 16GB ram installed)

```

Limit ARC

echo "options zfs zfs_arc_max=8589934592" >> /etc/modprobe.d/zfs.conf
echo "8589934592" > /sys/module/zfs/parameters/zfs_arc_max

```

- h. Install OZMT and create snapshot jobs

```

Install OZMT

yum install mercurial
cd /opt
hg clone https://bitbucket.org/ozmt/ozmt
cd ozmt
./install-ozmt.sh
ozmt-snapjobs-add.sh relaypool/raw/scripts daily/7 weekly/4 monthly/3
ozmt-snapjobs-add.sh relaypool/raw/data hourly/6 daily/1
ozmt-snapjobs-add.sh relaypool/zfs_tools hourly/12 daily/7 weekly/4 monthly/3
ozmt-snapjobs-add.sh relaypool/xnat/archive hourly/6 daily/1
ozmt-snapjobs-add.sh relaypool/xnat/prearchive hourly/6 daily/1
ozmt-snapjobs-add.sh relaypool/xnat/home daily/7 weekly/4
ozmt-snapjobs-add.sh relaypool/xnat/logs daily/7 weekly/4 monthly/6
ozmt-snapjobs-add.sh relaypool/xnat/pgsql/9.4 daily/7 weekly/4 monthly/6

```

- i. Scrub the pool daily
Add to root's crontab:

Scrub zpool

```
# Scrub the zpool
0 10 * * * /sbin/zpool scrub relaypool
1 10 * * * /sbin/zpool scrub rpool
```

- j. Comment out the replication jobs in root's crontab:

Disable ZFS Replication

```
## * * * * * /opt/ozmt/replication/schedule-replication.sh
#1,11,21,31,41,51 * * * * /opt/ozmt/replication/replication-job-runner.sh
#5,15,25,35,45,55 * * * * /opt/ozmt/replication/replication-job-cleaner.sh
```

3. ZFS root

- a. Setup second SSD

Setup rpool

```
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

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 /dev/disk/by-vdev/ssd2-part1

zfs create -o canmount=off -o mountpoint=none rpool/ROOT
zfs create -o canmount=noauto -o mountpoint=/ rpool/ROOT/centos
zfs mount rpool/ROOT/centos
zfs create -o setuid=off rpool/home
zfs create -o mountpoint=/root rpool/home/root
zfs create -o canmount=off -o setuid=off -o exec=off rpool/var
zfs create -o com.sun:auto-snapshot=false rpool/var/cache
zfs create rpool/var/log
zfs create rpool/var/spool
zfs create -o com.sun:auto-snapshot=false -o exec=on rpool/var/tmp
zfs create -o mountpoint=/owncloud rpool/owncloud

mkdir /mnt/tmp
mount --bind / /mnt/tmp

rsync -avhX --stats /mnt/tmp/. /mnt/rpool/.
rsync -avhX --stats /boot/. /mnt/rpool/boot/.

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 /etc/zfs/zpool.cache
```

- b. Comment out all mounts in /etc/fstab

/etc/fstab

```
#
# /etc/fstab
# Created by anaconda on Mon Oct 31 10:11:52 2016
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
#UUID=0078e373-f435-4b66-b08d-df5cf1300554 / xfs defaults 0 0
#UUID=c952f347-1ebe-49eb-be3d-213f8f69f37d /boot xfs defaults 0 0
#UUID=fa96a497-bbd0-42ef-9507-1c6b6a27586a swap swap defaults 0 0
```

- c. Add two lines to /etc/default/grub

/etc/default/grub

```
GRUB_CMDLINE_LINUX="crashkernel=auto boot=zfs rpool=rpool bootfs=rpool/ROOT/centos zfsforce=1"
GRUB_PRELOAD_MODULES="part_gpt zfs"
```

You might need to comment-out GRUB_HIDDEN_TIMEOUT so you get grub menu during boot. This is needed to be able to select other boot entries.

```
#GRUB_HIDDEN_TIMEOUT=0
```

- d. Generate new grub config, and verify it has the correct root entry

grub config

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

- e. Enable ZFS services

```
systemctl enable zfs-import-cache
systemctl enable zfs-import-scan
systemctl enable zfs-mount
systemctl enable zfs-share
systemctl enable zfs-zed
systemctl enable zfs.target
```

- f. Add zfs to the list of modules dracut should include by default to /etc/dracut.conf

/etc/dracut.conf

```
# PUT YOUR CONFIG HERE OR IN separate files named *.conf
# in /etc/dracut.conf.d
# SEE man dracut.conf(5)
# Sample dracut config file
#logfile=/var/log/dracut.log
#fileloglvl=6
# Exact list of dracut modules to use.  Modules not listed here are not going
# to be included.  If you only want to add some optional modules use
# add_dracutmodules option instead.
#dracutmodules+=" "
# dracut modules to omit
#omit_dracutmodules+=" "
# dracut modules to add to the default
add_dracutmodules+="zfs"
# additional kernel modules to the default
#add_drivers+=" "
# list of kernel filesystem modules to be included in the generic initramfs
#filesystems+=" "
# build initrd only to boot current hardware
#hostonly="yes"
#
# install local /etc/mdadm.conf
#mdadmconf="no"
# install local /etc/lvm/lvm.conf
#lvmconf="no"
# A list of fsck tools to install.  If it's not specified, module's hardcoded
# default is used, currently: "umount mount /sbin/fsck* xfs_db xfs_check
# xfs_repair e2fsck jfs_fsck reiserfsck btrfsck".  The installation is
# opportunistic, so non-existing tools are just ignored.
#fscks=""
# inhibit installation of any fsck tools
#nofscks="yes"
# mount / and /usr read-only by default
#ro_mnt="no"
# set the directory for temporary files
# default: /var/tmp
#tmpdir=/tmp
```

g. Rebuild initramfs

Rebuild initramfs

```
dracut -f -v /boot/initramfs-$(uname -r).img $(uname -r)
```

h. Setup links in /dev on each boot

```
chmod +x /etc/rc.local
vim /etc/rc.local
```

/etc/rc.local

```
#!/bin/bash
# THIS FILE IS ADDED FOR COMPATIBILITY PURPOSES
#
# It is highly advisable to create own systemd services or udev rules
# to run scripts during boot instead of using this file.
#
# In contrast to previous versions due to parallel execution during boot
# this script will NOT be run after all other services.
#
# Please note that you must run 'chmod +x /etc/rc.d/rc.local' to ensure
# that this script will be executed during boot.
touch /var/lock/subsys/local

cd /dev;ln -s /dev/disk/by-id/* .
```

i. Setup ZED

Edit `/etc/zfs/zed.d/zed.rc`:

ZED

```
##
# zed.rc
#
# This file should be owned by root and permissioned 0600.
##
##
# Absolute path to the debug output file.
#
#ZED_DEBUG_LOG="/tmp/zed.debug.log"
##
# Email address of the zpool administrator for receipt of notifications;
# multiple addresses can be specified if they are delimited by whitespace.
# Email will only be sent if ZED_EMAIL_ADDR is defined.
# Disabled by default; uncomment to enable.
#
ZED_EMAIL_ADDR="root nrg-admin@nrg.wustl.edu"
##
# Name or path of executable responsible for sending notifications via email;
# the mail program must be capable of reading a message body from stdin.
# Email will only be sent if ZED_EMAIL_ADDR is defined.
#
#ZED_EMAIL_PROG="mail"
##
# Command-line options for ZED_EMAIL_PROG.
# The string @ADDRESS@ will be replaced with the recipient email address(es).
# The string @SUBJECT@ will be replaced with the notification subject;
# this should be protected with quotes to prevent word-splitting.
# Email will only be sent if ZED_EMAIL_ADDR is defined.
#
ZED_EMAIL_OPTS="-s '@SUBJECT@' @ADDRESS@"
##
# Default directory for zed lock files.
#
#ZED_LOCKDIR="/var/lock"
##
# Minimum number of seconds between notifications for a similar event.
#
ZED_NOTIFY_INTERVAL_SECS=3600
##
# Notification verbosity.
# If set to 0, suppress notification if the pool is healthy.
# If set to 1, send notification regardless of pool health.
#
ZED_NOTIFY_VERBOSE=1
##
```



```

# Pushbullet access token.
# This grants full access to your account -- protect it accordingly!
# <https://www.pushbullet.com/get-started>
# <https://www.pushbullet.com/account>
# Disabled by default; uncomment to enable.
#
#ZED_PUSHBULLET_ACCESS_TOKEN=""
##
# Pushbullet channel tag for push notification feeds that can be subscribed to.
# <https://www.pushbullet.com/my-channel>
# If not defined, push notifications will instead be sent to all devices
# associated with the account specified by the access token.
# Disabled by default; uncomment to enable.
#
#ZED_PUSHBULLET_CHANNEL_TAG=""
##
# Default directory for zed state files.
#
#ZED_RUNDIR="/var/run"
##
# Replace a device with a hot spare after N checksum errors are detected.
# Disabled by default; uncomment to enable.
#
#ZED_SPARE_ON_CHECKSUM_ERRORS=10
##
# Replace a device with a hot spare after N I/O errors are detected.
# Disabled by default; uncomment to enable.
#
#ZED_SPARE_ON_IO_ERRORS=1
##
# The syslog priority (e.g., specified as a "facility.level" pair).
#
#ZED_SYSLOG_PRIORITY="daemon.notice"
##
# The syslog tag for marking zed events.
#
#ZED_SYSLOG_TAG="zed"

```

j. Create swap space

Setup swap space

```

mkswap /dev/disk/by-vdev/ssd2-part4
echo "/dev/disk/by-vdev/ssd2-part4 swap swap defaults 0 0" >> /etc
fstab

```

k. Reboot

4. Enter BIOS and change boot to ZFS root disk

5. Confirm zfs root is running properly

a. List zpool status and zfs

Check ZFS

```
# zpool status
pool: rpool
state: ONLINE
status: Some supported features are not enabled on the pool. The pool can
still be used, but some features are unavailable.
action: Enable all features using 'zpool upgrade'. Once this is done,
the pool may no longer be accessible by software that does not support
the features. See zpool-features(5) for details.
scan: none requested
config:
NAME                                STATE      READ WRITE CKSUM
rpool                                ONLINE    0   0   0
  ata-Samsung_SSD_750_EVO_120GB_S33MNB0H911742A-part1 ONLINE    0   0   0
errors: No known data errors
```

```
# zfs list
NAME                                USED  AVAIL  REFER  MOUNTPOINT
rpool                                4.96G 33.5G  136K  /
rpool/ROOT                            4.53G 33.5G  136K  none
rpool/ROOT/centos                     4.53G 33.5G  4.39G  /
rpool/home                             4.97M 33.5G  156K  /home
rpool/home/hcpadmin                   4.04M 33.5G  3.85M  /home/hcpadmin
rpool/home/root                       712K 33.5G  540K  /root
rpool/var                              433M 33.5G  136K  /var
rpool/var/cache                       413M 33.5G  412M  /var/cache
rpool/var/log                         17.4M 33.5G  16.8M  /var/log
rpool/var/spool                       1.70M 33.5G  1.19M  /var/spool
rpool/var/tmp                         876K 33.5G  604K  /var/tmp
```

```
# zfs list -o name,mounted,mountpoint
NAME                                MOUNTED  MOUNTPOINT
rpool                                no       /
rpool/ROOT                            no       none
rpool/ROOT/centos                     yes      /
rpool/home                             yes      /home
rpool/home/hcpadmin                   yes      /home/hcpadmin
rpool/home/root                       yes      /root
rpool/var                              no       /var
rpool/var/cache                       yes      /var/cache
rpool/var/log                         yes      /var/log
rpool/var/spool                       yes      /var/spool
rpool/var/tmp                         yes      /var/tmp
```

```
# Confirm swap is mounted
```

```
# free
```

	total	used	free	shared	buff/cache	available
Mem:	16405364	905724	14940016	9152	559624	14992148
Swap:	8388604	0	8388604			

If any zfs folders are not mounted besides rpool, rpool/ROOT and rpool/home these need to be corrected before proceeding.

6. Setup mirror to original SSD

a. Re-import relaypool

zpool import

```
zpool import relaypool
```

b. Delete all partitions from original SSD

```
fdisk

fdisk /dev/disk/by-vdev/ssd1
```

c. Setup partition

```
Partition SSD

sgdisk -g -a1 -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
```

d. Mirror rpool.

First determine the device name of the existing disk.

```
zpool status

# zpool status
pool: rpool
state: ONLINE
status: Some supported features are not enabled on the pool. The pool can
still be used, but some features are unavailable.
action: Enable all features using 'zpool upgrade'. Once this is done,
the pool may no longer be accessible by software that does not support
the features. See zpool-features(5) for details.
scan: none requested
config:
NAME                                STATE      READ WRITE CKSUM
rpool                                ONLINE    0     0     0
  ata-Samsung_SSD_750_EVO_120GB_S33MNB0H911742A-part1  ONLINE    0     0     0
errors: No known data errors
```

Attach the other SSD. (Substitute the device name returned from zpool status).

```
zpool attach

# zpool attach rpool ata-Samsung_SSD_750_EVO_120GB_S33MNB0H911742A-part1 ssd1-part1
```

Wait for the rpool to resilver.

zpool status

```
# zpool status
pool: rpool
state: ONLINE
status: One or more devices is currently being resilvered.  The pool will
        continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
        scan: resilver in progress since Wed Nov 30 11:30:23 2016
              3.03G scanned out of 4.97G at 73.9M/s, 0h0m to go
              3.03G resilvered, 61.06% done
config:
      NAME                                STATE     READ WRITE CKSUM
      rpool                                ONLINE    0    0    0
        mirror-0
          ata-Samsung_SSD_750_EVO_120GB_S33MNB0H911742A-part1  ONLINE    0    0    0
          ssd1-part1                                ONLINE    0    0    0
(resilvering)
errors: No known data errors
```

- e. Reboot
- f. Re-apply grub setup

grub config

```
cd /dev;ln -s /dev/disk/by-id/*
grub2-mkconfig -o /boot/grub2/grub.cfg
grep ROOT /boot/grub2/grub.cfg
grub2-install /dev/disk/by-vdev/ssd1
grub2-install /dev/disk/by-vdev/ssd2
```

- g. Create snapshot policies

Snapshots

```
zfs create -o mountpoint=/rpool/zfs_tools rpool/zfs_tools
ozmt-snapjobs-mod.sh rpool/ROOT/centos daily/7 weekly/4 monthly/6
ozmt-snapjobs-mod.sh rpool/home/root daily/7 weekly/4 monthly/6
```

- h. Add additional zfs cache

Additional ZFS cache

```
zpool add relaypool cache ssd1-part3
```

- i. Add additional swap space

Setup swap space

```
mkswap /dev/disk/by-vdev/ssd1-part4
echo "/dev/disk/by-vdev/ssd1-part4 swap swap defaults 0 0" >> /etc
fstabswapon -afree
```

7. Addition server configuration

- a. Setup email
Preserve [main.cf](#)

Setup email

```
cd /etc/postfix
mv main.cf main.dist
```

Create new [main.cf](#). Update as necessary per site.

main.cf

```
# See /usr/share/postfix/main.cf.dist for a commented, more complete version
smtpd_banner = $myhostname ESMTP $mail_name
biff = no
append_dot_mydomain = no
# this will add 'POP.yourdomain.tld' to the domain
#sender_canonical_maps = regexp:/etc/postfix/sender_regexp
# Uncomment the next line to generate "delayed mail" warnings
#delay_warning_time = 4h
alias_maps = hash:/etc/aliases
alias_database = hash:/etc/aliases
myorigin = nrg.wustl.edu
relayhost = mail.nrg.wustl.edu
recipient_delimiter = +
inet_interfaces = 127.0.0.1
local_transport = error:local delivery is disabled
```

Add email alias for root

Setup email

```
echo "root:          nrg-admin@nrg.wustl.edu" >> /etc/aliases
newaliases
```

Create /etc/ozmt/reporting.muttrc Adjust as necessary.

reporting.muttrc

```
set realname="lifespan-relay1"
set hostname=lifespan-relay1.nrg.mir
set from="lifespan-relay1-no-reply@myrealdomain.com"
set use_envelope_from=yes
set ssl_use_sslv3=no
set ssl_use_tlsv1=no
```

Set email_to in /etc/ozmt/config

/etc/ozmt/config

```
# Address to send reports to
email_to="nrg-admin@nrg.wustl.edu"
```

Setup rpool backups

rpool backup

```
echo "
```

b. Setup Tomcat for XNAT

- i. Add xnat user

XNAT user

```
useradd -c "XNAT system user" -d /xnat/home -s /bin/false xnat
```

- ii. Fix up /xnat directories

/xnat

```
cd /xnat
mkdir -p logs/{catalina,xnat,tomcat} tmp/{catalina,xnat,tomcat} home/{config,plugins,work}
cd home
ln -s ../logs/xnat logs
ln -s ../tmp/xnat tmp
cd ..
chmod -R g-rx,o-rx .
chown -R xnat .
```

- iii. Fix up the symlinks in /usr/share/tomcat

/usr/share/tomcat

```
cd /usr/share/tomcat
rm -f tmp logs
ln -s /xnat/tmp/tomcat tmp
ln -s /xnat/logs/tomcat logs
```

- iv. Create XNAT Tomcat service

systemctl

```
systemctl enable tomcat@xnat.service
```

- v. Fix up /etc/tomcat/server.xml
Make the Host block match

server.xml

```
<Host name="localhost" appBase="webapps"
      unpackWARs="true" autoDeploy="true"
      xmlValidation="false" xmlNamespaceAware="false">
  <!-- SingleSignOn valve, share authentication between web applications
       Documentation at: /docs/config/valve.html -->
  <!--
  <Valve className="org.apache.catalina.authenticator.SingleSignOn" />
  -->
  <!-- Access log processes all example.
       Documentation at: /docs/config/valve.html
       Note: The pattern used is equivalent to using pattern="common" -->
  <!--
  <Valve className="org.apache.catalina.valves.AccessLogValve" directory="logs"
        prefix="localhost_access_log." suffix=".txt"
        pattern="%h %l %u %t &quot;%r&quot; %s %b" />
  -->
</Host>
```

- vi. Modify /etc/sysconfig/tomcat

/etc/sysconfig/tomcat

```
# Service-specific configuration file for tomcat. This will be sourced by
# the SysV init script after the global configuration file
# /etc/tomcat/tomcat.conf, thus allowing values to be overridden in
# a per-service manner.
#
# NEVER change the init script itself. To change values for all services make
# your changes in /etc/tomcat/tomcat.conf
#
# To change values for a specific service make your edits here.
# To create a new service create a link from /etc/init.d/<your new service> to
# /etc/init.d/tomcat (do not copy the init script) and make a copy of the
# /etc/sysconfig/tomcat file to /etc/sysconfig/<your new service> and change
# the property values so the two services won't conflict. Register the new
# service in the system as usual (see chkconfig and similars).
#
# Where your java installation lives
#JAVA_HOME="/usr/lib/jvm/java"
# Where your tomcat installation lives
#CATALINA_BASE="/usr/share/tomcat"
CATALINA_HOME="/usr/share/tomcat"
#JASPER_HOME="/usr/share/tomcat"
CATALINA_TMPDIR="/xnat/tmp/catalina"
# You can pass some parameters to java here if you wish to
#JAVA_OPTS="-Xminf0.1 -Xmaxf0.3"
# Use JAVA_OPTS to set java.library.path for libtcnative.so
#JAVA_OPTS="-Djava.library.path=/usr/lib64"
JAVA_OPTS="-Xms872m -Xmx2620m -Xmn524m -XX:-OmitStackTraceInFastThrow -XX:MaxPermSize=256m -
Dsun.net.inetaddr.ttl=30 -Dxnat.home=/xnat/home"
# What user should run tomcat
TOMCAT_USER="xnat"
# You can change your tomcat locale here
#LANG="en_US"
# Run tomcat under the Java Security Manager
SECURITY_MANAGER="false"
# Time to wait in seconds, before killing process
SHUTDOWN_WAIT="30"
# Whether to annoy the user with "attempting to shut down" messages or not
#SHUTDOWN_VERBOSE="false"
# Connector port is 8080 for this tomcat instance
#CONNECTOR_PORT="8080"
# If you wish to further customize your tomcat environment,
# put your own definitions here
# (i.e. LD_LIBRARY_PATH for some jdbc drivers)
TOMCAT_LOG="/xnat/logs/tomcat/catalina.out"
CATALINA_OPTS="-Xdebug -Xrunjdwp:transport=dt_socket,server=y,suspend=n,address=8000 -
Dcatalina.ext.dirs=/usr/share/tomcat/shared/lib:/usr/share/tomcat/common/lib -Dcom.sun.
management.jmxremote -Dcom.sun.management.jmxremote.port=9004 -Dcom.sun.management.
jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false"
```

vii. Edit Tomcat systemd unit. (Change the User line and disable OOM kill)

```
systemctl edit --full tomcat.service
```

tomcat@xnat.service

```
# Systemd unit file for tomcat instances.
#
# To create clones of this service:
# 0. systemctl enable tomcat@name.service
# 1. create catalina.base directory structure in
#    /var/lib/tomcats/name
# 2. profit.

[Unit]
Description=Apache Tomcat Web Application Container
After=syslog.target network.target

[Service]
Type=simple

EnvironmentFile=/etc/tomcat/tomcat.conf
Environment="NAME=%I"
EnvironmentFile=-/etc/sysconfig/tomcat@%I

ExecStart=/usr/libexec/tomcat/server start
ExecStop=/usr/libexec/tomcat/server stop
SuccessExitStatus=143
User=xnat
Group=tomcat

# Disable OOM kill on tomcat
OOMScoreAdjust=-1000

[Install]
WantedBy=multi-user.target
```

c. Setup PostgreSQL

i. Install PostgreSQL

PostgreSQL 9.4

```
yum install -y https://download.postgresql.org/pub/repos/yum/9.4/redhat/rhel-7-x86_64/pgdg-centos94-9.4-3.noarch.rpm
yum install -y postgresql94-server
```

ii. Edit postgresql-9.4.service

```
systemctl edit --full postgresql-9.4.service
```


postgresql-9.4.service

```
# It's not recommended to modify this file in-place, because it will be
# overwritten during package upgrades.  If you want to customize, the
# best way is to create a file "/etc/systemd/system/postgresql-9.4.service",
# containing
#     .include /lib/systemd/system/postgresql-9.4.service
#     ...make your changes here...
# For more info about custom unit files, see
# http://fedoraproject.org/wiki/Systemd#How\_do\_I\_customize\_a\_unit\_file.
2F_add_a_custom_unit_file.3F

# Note: changing PGDATA will typically require adjusting SELinux
# configuration as well.

# Note: do not use a PGDATA pathname containing spaces, or you will
# break postgresql-setup.

[Unit]
Description=PostgreSQL 9.4 database server
After=syslog.target
After=network.target

[Service]
Type=forking

User=postgres
Group=postgres

# Note: avoid inserting whitespace in these Environment= lines, or you may
# break postgresql-setup.

# Location of database directory
Environment=PGDATA=/xnat/pgsql/9.4
Environment=PGLOG=/xnat/logs/pgsql/pgstartup-9.4.log

# Where to send early-startup messages from the server (before the logging
# options of postgresql.conf take effect)
# This is normally controlled by the global default set by systemd
# StandardOutput=syslog

# Disable OOM kill on the postmaster
OOMScoreAdjust=-1000

ExecStartPre=/usr/pgsql-9.4/bin/postgresql94-check-db-dir ${PGDATA}
ExecStart=/usr/pgsql-9.4/bin/pg_ctl start -D ${PGDATA} -s -w -t 300
ExecStop=/usr/pgsql-9.4/bin/pg_ctl stop -D ${PGDATA} -s -m fast
ExecReload=/usr/pgsql-9.4/bin/pg_ctl reload -D ${PGDATA} -s

# Give a reasonable amount of time for the server to start up/shut down
TimeoutSec=300

[Install]
WantedBy=multi-user.target
```

iii. Start Postgres for the first time

Start Postgres

```
cd /xnat/pgsql
chown postgres:postgres 9.4
chmod 755 /xnat

cd /xnat/logs
mkdir pgsq
chown postgres pgsq

cd
/usr/pgsql-9.4/bin/postgresql94-setup initdb

systemctl start postgresql-9.4.service
```

iv. Setup postgresql.conf

```
cd /xnat/pgsql/9.4
mv postgresql.conf postgresql.conf.dist
```

postgresql.conf

```
max_connections = 100
shared_buffers = 612MB
work_mem = 50MB
maintenance_work_mem = 256MB
dynamic_shared_memory_type = posix
archive_mode = off
effective_cache_size = 1225MB
log_destination = 'stderr'
logging_collector = on
log_directory = 'pg_log'
log_filename = 'postgresql-%a.log'
log_truncate_on_rotation = on
log_rotation_age = 1d
log_rotation_size = 0
log_min_duration_statement = 30000
log_line_prefix = '< %m >'
log_timezone = 'US/Central'
datestyle = 'iso, mdy'
timezone = 'US/Central'
lc_messages = 'en_US.UTF-8'
lc_monetary = 'en_US.UTF-8'
lc_numeric = 'en_US.UTF-8'
lc_time = 'en_US.UTF-8'
default_text_search_config = 'pg_catalog.english'
standard_conforming_strings = off
```

v. Give 'postgres' user a password

postgres user password

```
su - postgres
psql -c "ALTER USER postgres with password '<put your postgres password here>'"
exit
```

vi. Setup pg_hba.conf

Update procedure:

Check the ZFS on Linux wiki for new information before proceeding:

<https://github.com/zfsonlinux/zfs/wiki/RHEL-and-CentOS>

Remove and re-install zfs for the latest Centos release

```
yum remove zfs zfs-kmod spl spl-kmod libzfs2 libnvpair1 libuutil1 libzpool2 zfs-release
yum install http://download.zfsonlinux.org/epel/zfs-release.el7_4.noarch.rpm
find /lib/modules/ \( -name "splat.ko" -or -name "zcommon.ko" \
    -or -name "zpios.ko" -or -name "spl.ko" -or -name "zavl.ko" -or \
    -name "zfs.ko" -or -name "znpvpair.ko" -or -name "zunicode.ko" \) \
-exec /bin/rm {} \;
```

Disable [zfs] and enable [zfs-kmod] in the zfs repo:

vim /etc/yum.repos.d/zfs.repo

```
[zfs]
name=ZFS on Linux for EL7 - dkms
baseurl=http://download.zfsonlinux.org/epel/7.4/$basearch/
enabled=0
metadata_expire=7d
ggpgcheck=1
ggpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-zfsonlinux

[zfs-kmod]
name=ZFS on Linux for EL7 - kmod
baseurl=http://download.zfsonlinux.org/epel/7.4/kmod/$basearch/
enabled=1
metadata_expire=7d
ggpgcheck=1
ggpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-zfsonlinux

[zfs-source]
name=ZFS on Linux for EL7 - Source
baseurl=http://download.zfsonlinux.org/epel/7.4/SRPMS/
enabled=0
ggpgcheck=1
ggpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-zfsonlinux

[zfs-testing]
name=ZFS on Linux for EL7 - dkms - Testing
baseurl=http://download.zfsonlinux.org/epel-testing/7.4/$basearch/
enabled=0
metadata_expire=7d
ggpgcheck=1
ggpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-zfsonlinux

[zfs-testing-kmod]
name=ZFS on Linux for EL7 - kmod - Testing
baseurl=http://download.zfsonlinux.org/epel-testing/7.4/kmod/$basearch/
enabled=0
metadata_expire=7d
ggpgcheck=1
ggpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-zfsonlinux

[zfs-testing-source]
name=ZFS on Linux for EL7 - Testing Source
baseurl=http://download.zfsonlinux.org/epel-testing/7.4/SRPMS/
enabled=0
ggpgcheck=1
ggpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-zfsonlinux
```

Reinstall ZFS:

```
yum autoremove
yum install zfs zfs-dracut
systemctl preset zfs-import-cache zfs-import-scan zfs-mount zfs-share zfs-zed zfs.target

# Do the system update
yum update

# Re-enable services
systemctl enable zfs-import-cache
systemctl enable zfs-import-scan
systemctl enable zfs-mount
systemctl enable zfs-share
systemctl enable zfs-zed
systemctl enable zfs.target
```

Make the latest kernel the default in grub:

```
grep '^menuentry' /boot/grub2/grub.cfg
vim /etc/default/grub
grub2-mkconfig -o /boot/grub2/grub.cfg
```

TODO:

Setup Nagios Monitoring:

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

References

<https://github.com/zfsonlinux/zfs/wiki/RHEL-%26-CentOS>

[https://github.com/zfsonlinux/pkg-zfs/wiki/HOWTO-install-EL7-\(CentOS-RHEL\)-to-a-Native-ZFS-Root-Filesystem](https://github.com/zfsonlinux/pkg-zfs/wiki/HOWTO-install-EL7-(CentOS-RHEL)-to-a-Native-ZFS-Root-Filesystem)

<https://github.com/zfsonlinux/zfs/wiki/Ubuntu%2016.04%20Root%20on%20ZFS>