

Configuring: How do I integrate ZFS with Bright?

[How do I integrate ZFS on Bright?](#)

Here's a recipe to follow:

Installing and Configuring A ZFS filesystem on top of a Bright Cluster

CENTOS 6

On the head node

1. Install the required repositories and packages:

```
root@b70-c6 ~]# yum localinstall --nogpgcheck http://archive.zfslinux.org/epel/zfs-release.el6.noarch.rpm
```

optional step to save the repository RPM for future/local usage

```
root@b70-c6 ~]# wget -c http://archive.zfslinux.org/epel/zfs-release.el6.noarch.rpm
```

```
root@b70-c6 ~]# yum install zfs
```

2. Create a 10GB file to host the ZFS filesystem for testing

```
root@b70-c6 ~]# dd if=/dev/zero of=/opt/zfs.img bs=1073741824 count=10
```

```
10+0 records in
```

```
10+0 records out
```

```
10737418240 bytes (11 GB) copied, 238.08 s, 45.1 MB/s
```

3. Create a ZFS pool

```
root@b70-c6 ~]# zpool create localzfs /opt/zfs.img
```

Configuring: How do I integrate ZFS with Bright?

4. Create a ZFS filesystem

```
root@b70-c6 ~]# zfs create localzfs/data
```

```
root@b70-c6 ~]# df -hT
```

```
Filesystem      Type      Size  Used Avail Use% Mounted on
```

```
dev/vda3        ext3      38G   23G   14G   64% /
```

```
tmpfs           tmpfs     2.0G   0 2.0G   0% /dev/shm
```

```
dev/vda1        ext2      504M   24M  456M   5% /boot
```

```
localzfs        zfs       9.8G  128K  9.8G   1% /localzfs
```

```
localzfs/data   zfs       9.8G  128K  9.8G   1% /localzfs/data
```

5. On the head node, the ZFS daemon must be loaded at startup so that it can automatically mount the ZFS pools which are defined in `/etc/zfs/zpool.cache` without adding an entry in `/etc/fstab`

```
root@b70-c6 ~]# chkconfig zfs on
```

```
root@b70-c6 ~]# /etc/init.d/zfs status
```

```
pool: localzfs
```

```
state: ONLINE
```

```
scan: none requested
```

```
config:
```

```
NAME      STATE  READ WRITE CKSUM
```

```
localzfs  ONLINE  0   0   0
```

```
/opt/zfs.img  ONLINE  0   0   0
```

```
errors: No known data errors
```

```
NAME      USED AVAIL REFER MOUNTPOINT
```

Configuring: How do I integrate ZFS with Bright?

```
localzfs 148K 9.78G 31K /localzfs
```

```
localzfs/data 30K 9.78G 30K /localzfs/data
```

Zfs is up and running on the head at this point.

On the compute nodes

1. Install the required repositories and packages inside the software image:

```
root@b70-c7 ~]# yum localinstall --nogpgcheck http://archive.zfslinux.org/epel/zfs-release.el7.noarch.rpm --installroot=/cm/images/default-image
```

```
root@b70-c7 ~]# yum install zfs --nogpgcheck --installroot=/cm/images/default-image
```

```
root@b70-c6 ~]# yum install zfs --installroot=/cm/images/default-image
```

2. Make sure that the zfs service is disabled on boot in the software image

```
root@b70-c6 ~]# chroot /cm/images/default-image
```

```
root@b70-c6 /]# chkconfig --list | grep zfs
```

```
zfs          0:off  1:off  2:off  3:off  4:off  5:off  6:off
```

3. Create a 5G file to host the ZFS filesystem for testing on the node:

```
root@node001 ~]# dd if=/dev/zero of=/opt/zfs.img bs=1073741824 count=5
```

```
5+0 records in
```

```
5+0 records out
```

```
5368709120 bytes (5.4 GB) copied, 73.3398 s, 73.2 MB/s
```

Configuring: How do I integrate ZFS with Bright?

4. Configure the `excludelistsyncinstall` and `excludelistupdate` `excludelists`:

```
root@b70-c6 ~# cmt
```

```
b70-c6]% category use default
```

```
b70-c6->category[default]% set excludelistsyncinstall
```

```
...
```

```
~/opt/zfs.img
```

```
~/etc/zfs/
```

```
...
```

```
no-new-files: - /opt/zfs.img
```

```
no-new-files: - /etc/zfs/
```

```
...
```

```
b70-c6->category[default]% commit
```

```
b70-c6->category[default]% set excludelistupdate
```

```
...
```

```
~/opt/zfs.img
```

```
~/etc/zfs/
```

```
...
```

```
no-new-files: - /opt/zfs.img
```

```
no-new-files: - /etc/zfs/
```

```
...
```

```
b70-c6->category[default]% commit
```

5. Create a ZFS pool:

Configuring: How do I integrate ZFS with Bright?

```
root@node001 ~# zpool create localnfs /opt/zfs.img
```

```
root@node001 ~# zfs create localnfs/data
```

```
root@node001 ~# /etc/init.d/zfs status
```

```
pool: localnfs
```

```
state: ONLINE
```

```
scan: none requested
```

```
config:
```

```
NAME      STATE  READ WRITE CKSUM
```

```
localnfs  ONLINE  0   0   0
```

```
/opt/zfs.img  ONLINE  0   0   0
```

```
errors: No known data errors
```

```
NAME      USED AVAIL REFER MOUNTPOINT
```

```
localnfs  148K 4.89G 31K /localnfs
```

```
localnfs/data 30K 4.89G 30K /localnfs/data
```

6. Modify `/etc/rc.local` to start ZFS automatically on startup after the modules are generated

```
root@b70-c6 ~# cat /cm/images/default-image/etc/rc.local
```

```
[[
```

```
/etc/init.d/zfs restart
```

7. Mark the node as a datanode:

```
root@b70-c6 ~# cmsf
```

Configuring: How do I integrate ZFS with Bright?

```
b70-c6]% device use node001
```

```
b70-c6->device[node001]]% set d
```

```
datanode deviceheight deviceposition disksetup
```

```
b70-c6->device[node001]]% set datanode yes
```

```
b70-c6->device*[node001]]% commi
```

No restart needed for the regular nodes.

That's it, ZFS is running on the cluster.

Unique solution ID: #1256

Author: Frank Furter

Last update: 2015-01-27 17:57