

iSCSI + ZFS

Easy SAN

Quick Setup

- First time setup/startup of iSCSI target
 - Svcadm enable iscsitgt (/etc/iscsi config directory)
- Use ZFS as the backing store
 - `zfs create pool/iscsi_luns`
 - `zfs create -s -V 1T pool/iscsi_luns/vol001`
 - `zfs set shareiscsi=on pool/iscsi_luns/vol001`

Repeat for any additional volumes needed, but each should be mounted on client before continuing .
- ACL iSCSI volume to client
 - `iscsitadm modify target -acl CLIENT_IQN pool/iscsi_luns/vol001`
- CHAP Authentication for enhanced security

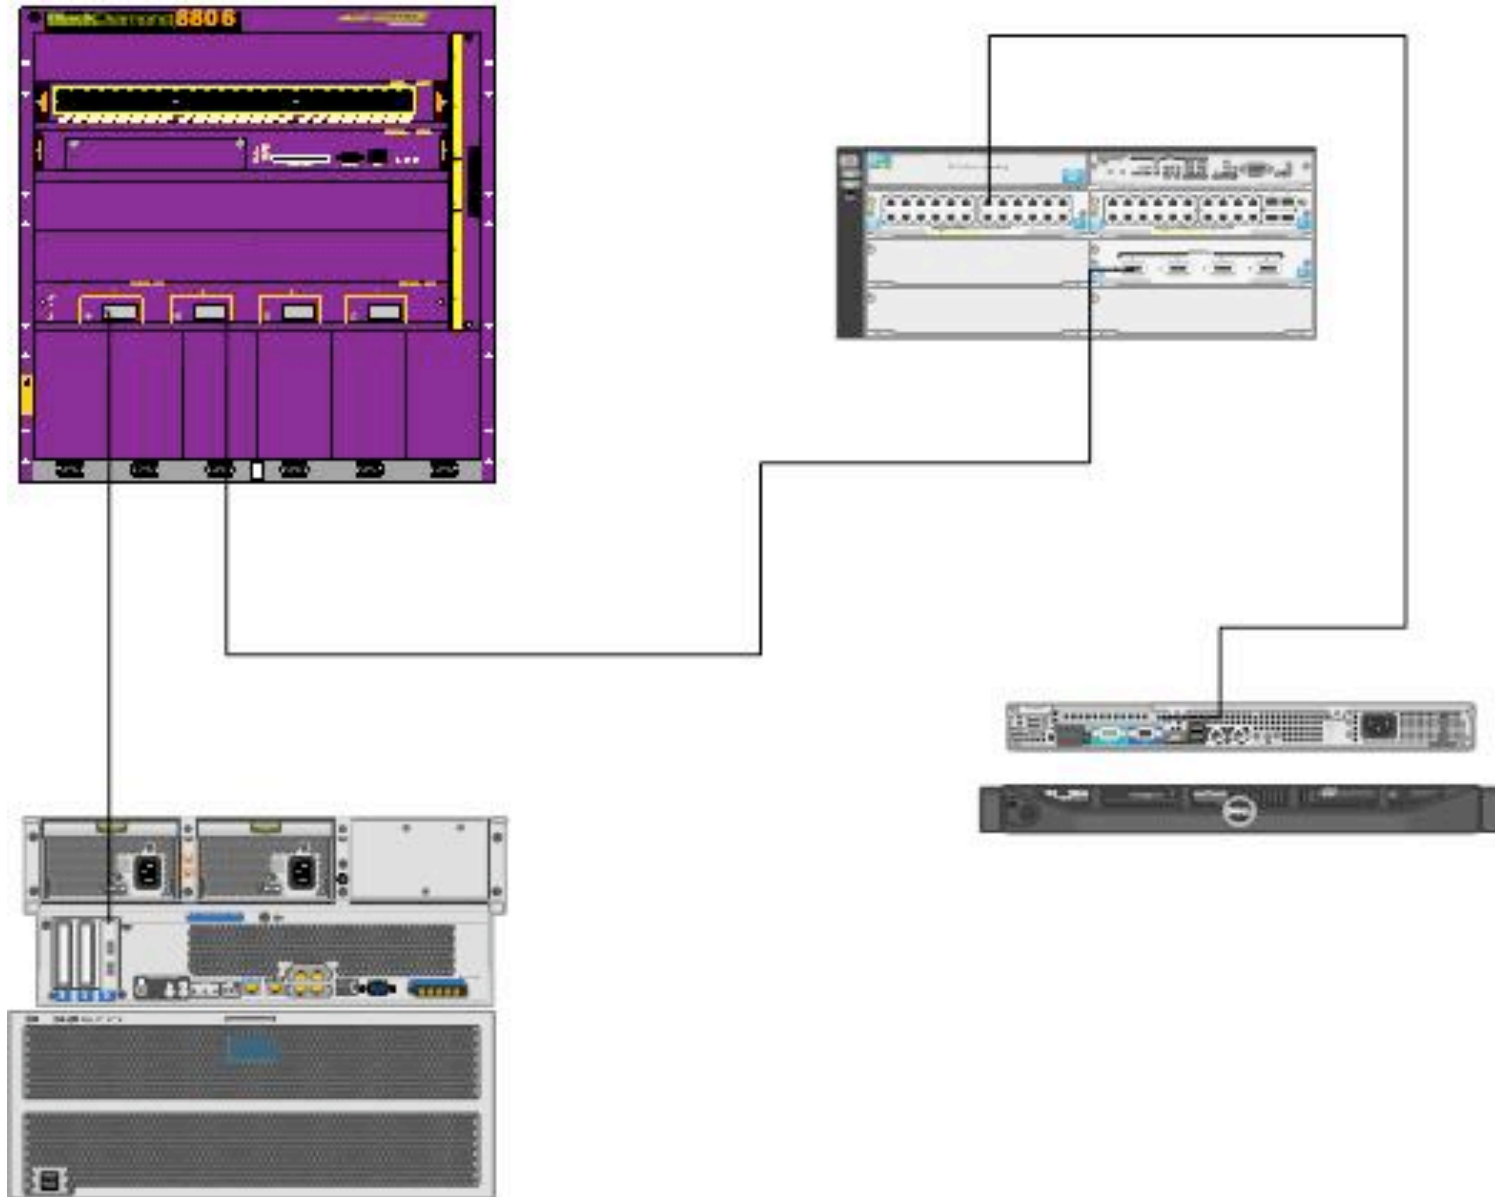
CentOS client

- iscsi-initiator-utils package
 - iSCSI Qualified Name (IQN) from /etc/iscsi/initiatorname.iscsi
- IQN should be configured on Server as ACL
- Query iSCSI server
 - `iscsiadm -mode discovery -type sendtargets -portal SERVER_IP_ADDRESS`
- Login to iSCSI volume
 - `iscsiadm -mode node -targetname SERVER_IQN -portal SERVER_IP_ADDRESS -login`
- Raw disk should be detected by Kernel and normal fdisk/mke2fs/mount/etc procedures can follow.

Example

```
# iscsitadm list target -v
Target: pool2/iscsi_luns/LI_GAINESTOWN_1
  iSCSI Name: ign.1986-03.com.sun:02:641cced6-0f1f-4239-8317-eea96c90a18f
  Alias: pool2/iscsi_luns/LI_GAINESTOWN_1
  Connections: 1
    Initiator:
      iSCSI Name: ign.1994-05.com.redhat:a4af318e6ba
      Alias: gainestown.dynopt.ece.udel.edu
  ACL list:
    Initiator: gainestown
  TPGT list:
  LUN information:
    LUN: 0
      GUID: 600144f04ae6e41100001b211cb9a500
      VID: SUN
      PID: SOLARIS
      Type: disk
      Size: 100G
      Backing store: /dev/zvol/rdisk/pool2/iscsi_luns/LI_GAINESTOWN_1
      Status: online
```

Test Platform Configuration



Testing platform

- Server: Sun x4550, Dual 2.3GHz Barcelona, 32GB Ram, Solaris snv_111, 10GB Ethernet.
- Client: Dell PowerEdge, 2.26GHz nehalem, 12GB Ram, CentOS 2.6.18-128.el5 64-bit, 1GB Ethernet.

Extreme Networks L3 switch connected to 10GB NIC and HP5400zl 10GB L2 switch.

100GB iSCSI lun(s) carved from 9TB pool, on a 20TB 4x10 disk RAIDZ2 backing store.

Bonnie++ v1.96 and DD testing on 30GB file size tests.

Results

- Baseline 100GB local disk partition
 - Bonnie-> 60/66MBs W/R
 - DD-> 72MBs Write
- iSCSI 100GB EXT2 format LUN
 - Bonnie-> 24/35MBs W/R
 - DD-> 24MBs Write
- NFS 100GB mount
 - Bonnie-> 34/35MBs W/R
 - DD-> 25MBs Write

Extended Results

- Software RAID0 of 4x 100GB iSCSI LUNs
- iSCSI
 - Bonnie-> 32/22MBs W/R
 - DD-> 36MBs Write

Intrinsic Issues

- 10GbE <-> 1GbE
- Layer3 switching
- Nagle/TCP tuning
- Jumbo Frames
- Client CPU overhead
- ZIL (writes) / ARC & L2ARC (reads)