



## SOP – Install PC File Server

---

### 1 Preamble

These are installation instructions for installing CentOS 7 x64 on PC file servers running Supermicro hardware equipped with ARC-1882i or ARC-1883i hardware RAID controllers.

### 2 Network configuration for Common PC file server service

- Put server on Vlan “BMC-hall-server” and management on Vlan “BMC-hall-IPMI”.
- Start server and enter BIOS and change IPMI interface from failover to **dedicated** and **activate DHCP**.
- Use hostnames bmc-pcfsX.bmc.uu.se and bmc-pcfsX-ipmi.bmc.uu.se
- The second 10 Gbit/s interface is used for connecting to the backup server. Use static configuration there. Use a CAT7 cable for connecting.

IPMI MAC-address    \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_

IPMI IP-address     192.168.234.\_\_\_\_\_/27

IPMI hostname      bmc-pcfs\_\_\_\_\_-ipmi.uu.se

Server MAC 1st     \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_

Server IP-address 1st 130.238.54.\_\_\_\_\_/27

Server hostname    bmc-pcfs\_\_\_\_\_.bmc.uu.se

Server MAC 2nd     \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_ - \_\_\_\_\_

Server IP 2nd    192.168.0. \_\_\_\_\_/24

- Enter server in IPAM with DHCP reserved for the above.
- Open in router filter the ports 138, 139, 140, 445 for TCP and UDP from UpUnet by mailing Netsupport.
- Optionally open port 548/tcp for Time Machine.
- Check that the server and IPMI get IP over DHCP automatically from Bluecat.
- Create static IP configuration for the second Ethernet interface enabling 10 Gbit/s between the primary and backup server.



- Old servers may have to be configured via ipmitool. Do this after installation to set IMPI to get IP via DHCP:

```
ipmitool lan set 1 ipsrc dhcp
```

- You want to set the IPMI password:

```
ipmitool user set password 2 secretpassword
```

- If you try to run the Supermicro Java Applet for the remote control you may have to add this line to allow connecting to the IP in the file java.policy


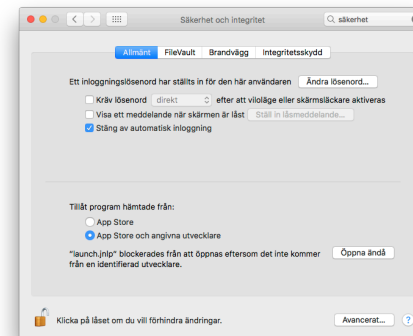
```
permission java.net.SocketPermission "A.B.C.D:*", "connect, accept ,resolve, listen";
```


### 3 Network configuration for the HPC file server

- Put the IPMI interface on the Vlan "ORG-IPMI" network, the first Ethernet interface on the Vlan "ORG" network and the second Ethernet interface on the Vlan "ORG-cluster"

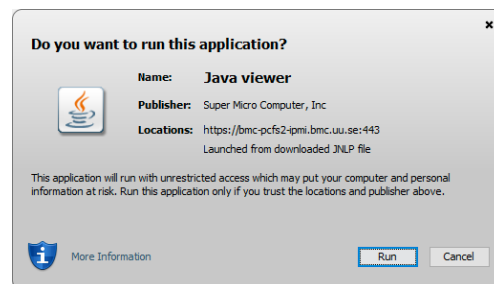
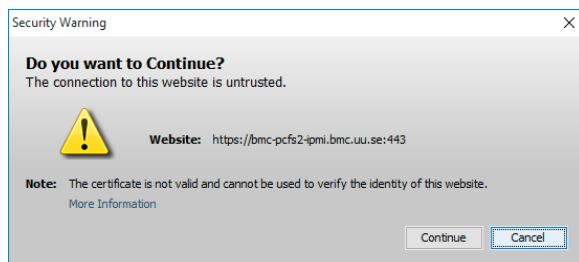
### 4 Open the Management console

- Go to <https://bmc-pcfsX-ipmi.bmc.uu.se> or <https://ORG-name.ORG.uu.se> and login with default management account **ADMIN** and password **ADMIN**.
- Open the Remote Console and the Java Web Start.
- In Mac OS X that requires opening the security preferences and manually allowing the application to start.





- In Windows that requires just to **Continue** and then **Run**.



For Windows and Linux it is also possible to download the IPMIView20 from Supermicro which is supposed to be better than the Java application, but I did not get that one to connect to the servers.

- If you really need to use the keyboard, special characters like the pipe symbol and larger and smaller than signs work better on the Java-console in Windows with English keyboard than Swedish in Mac. Temporarily change keyboard in server with: **loadkeys us**

## 5 Begin installation of host OS

- Start network boot by hammering **F12** during the American Megatrends BIOS startup. It's just after the RAID controller has been initialized. If there is no OS installed network boot is the default.
- Start **CentOS 7.2 RAID 1 kickstart** by typing **c72kr** at the PXE-menu prompt. This will overwrite /dev/sda and /dev/sdb so do not do it if the hardware RAID is configured with devices.
- Otherwise start the manual **CentOS 7.3 netinstall**. Remember to manually configure *Boot loader* for the two SSDs and use software RAID on them.

## 6 Configuration of OS

Some of this is done automatically during the kickstart installation. But these are all the steps.



- Get some standard tools:

```
yum -y install kexec-tools wget zip unzip emacs-nox screen dstat sysstat  
nmap nano ntpdate ntp net-tools yum-cron deltarpm dstat bind-utils rsync bc  
pciutils procmail lvm2 tcpdump smartmontools psmisc elinks dnsmasq iotop  
ipmitool telnet
```

- Turn off selinux

```
sed -i 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
```

- Set local hostname so it will be active before network is up.

```
hostname bmc-pcfsX.bmc.uu.se  
hostnamectl set-hostname $(hostname)  
echo HOSTNAME=$(hostname) >>/etc/sysconfig/network
```

- Login with SSH and change root password from the default to something better.

This is the new password \_\_\_\_\_

- Check that the hostname is correct (otherwise change in /etc/hostname and run command hostname)
- Download and install the command line tool for configuring the Areca hardware RAID controller:

```
wget  
http://www.areca.us/support/s_linux/driver/cli/linuxcli_V1.50.0_170105.zip  
unzip linuxcli_V1.50.0_170105  
cp -av linuxcli_V1.50.0_170105/x86_64/cli64 /usr/local/bin/  
chmod +x /usr/local/bin/cli64
```

- Download and install the latest driver for the Areca 1882 / 1883 hardware RAID controller. Reboot at some time to make it active.

```
mkdir 7_3  
cd 7_3  
wget http://www.areca.us/support/s_linux/driver/rhel/7_3.zip  
unzip 7_3.zip  
rpm -Uvh kmod-arcmsr-1.30.0X.27_rhel7.3-1.x86_64.rpm
```

- Only let root in.

```
echo AllowUsers root >>/etc/ssh/sshd_config  
systemctl enable sshd
```

- Remove NetworkManager and mlocate. Mlocate will be very confused with all the snapshots. Use iptables instead of firewall.

```
yum -y remove NetworkManager NetworkManager-libnm mlocate
```

```
yum -y install iptables-services
```



```
systemctl disable firewalld
systemctl enable iptables
iptables -F
service iptables save
```

- Apply automatic updates:

```
sed -i 's/apply_updates = no/apply_updates = yes/' /etc/yum/yum-cron.conf
```

- Activate rc.local

```
chmod +x /etc/rc.d/rc.local
systemctl enable rc-local
```

- Activate dnsmasq but still use DHCP

```
echo 'resolv-file=/etc/resolv.dnsmasq' > /etc/dnsmasq.d/resolv.file
echo 'DNS=127.0.0.1' >>/etc/sysconfig/network
host resolver.uu.se | grep -v IPv6 | awk '{print "nameserver " $4}'
>/etc/resolv.dnsmasq
echo 'nameserver 8.8.8.8' >>/etc/resolv.dnsmasq
echo 'nameserver 8.8.4.4' >>/etc/resolv.dnsmasq
systemctl start dnsmasq
systemctl enable dnsmasq
```

- Increase timeout for block devices to make the Areca controller happy.

```
echo 'for i in /sys/block/sd?/device/timeout ; do echo 900 >$i ; done'
>>/etc/rc.local
```

## 7 Make volumes and file systems

### 7.1 Check that all disk are there

```
mkdir /data
cli64 disk info
cli64 set password=0000
cli64 sys changept p=3
cli64 sys showcfg
```

The changept speeds up background tasks. The drives are numbered 9-44 for a total of 36 drives.

```
cli64 sys ncq p=1
cli64 adsys timeout p=8
cli64 adsys tler p=1
cli64 hddpwr spoweron p=4
```

Alternate setting to max timeout for ARC-I883:

```
cli64 adsys timeout p=17
```

Disable NCQ (Native Command Queuing) support, disable TLER (Time-Limited Error Recovery), set hard disk device timeout to 8 (22 seconds), set stagger power on control to 4 (2.0).



## 7.2 Put stickers on the drives

Find the little stickers with numbers in the box and put them on the drives in the right order. Stop the lights by identifying drive 0.

```
while true ; do cli64 set password=0000 ; for i in {1..60} ; do echo $i ; cli64 disk
identify drv=$i ; sleep 1 ; done ; done
cli64 disk identify drv=0
```

## 7.3 Create raidsets

```
[root@bmc-pcfs1 ~]# cli64 set password=0000
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 disk info
...
[root@bmc-pcfs1 ~]# cli64 rsf create drv=9~12 name=IMB-GenomicsCJR
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 rsf create drv=13~18 name=FBV-MSImaging
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 rsf create drv=19~24 name=IMB-GenomicsLA1
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 rsf create drv=25~30 name=IMB-GenomicsLA2
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 rsf create drv=31~36 name=IMB-GenomicsLA3
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 rsf create drv=37~42 name=MOL-EXTBMC
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 rsf info
# Name Disks TotalCap FreeCap MinDiskCap State
=====
1 IMB-GenomicsCJR 4 32000.0GB 32000.0GB 8000.0GB Normal
2 FBV-MSImaging 6 48000.0GB 48000.0GB 8000.0GB Normal
3 IMB-GenomicsLA1 6 48000.0GB 48000.0GB 8000.0GB Normal
4 IMB-GenomicsLA2 6 48000.0GB 48000.0GB 8000.0GB Normal
5 IMB-GenomicsLA3 6 48000.0GB 48000.0GB 8000.0GB Normal
6 MOL-EXTBMC 6 48000.0GB 48000.0GB 8000.0GB Normal
=====
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]#
```

## 7.4 Create volumegroups

```
[root@bmc-pcfs1 ~]# cli64 vsf create raid=1 level=6 fginit=N name=IMB-GenomicsCJR
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 vsf create raid=2 level=6 fginit=N name=FBV-MSImaging
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 vsf create raid=3 level=6 fginit=N name=IMB-GenomicsLA1
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 vsf create raid=4 level=6 fginit=N name=IMB-GenomicsLA2
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 vsf create raid=5 level=6 fginit=N name=IMB-GenomicsLA3
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 vsf create raid=6 level=6 fginit=N name=MOL-EXTBMC
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]# cli64 vsf info
# Name Raid Name Level Capacity Ch/Id/Lun State
=====
1 IMB-GenomicsCJR IMB-GenomicsCJR Raid6 16000.0GB 00/00/00 Initializing(0.0%)
2 FBV-MSImaging FBV-MSImaging Raid6 32000.0GB 00/00/01 Initializing(0.0%)
3 IMB-GenomicsLA1 IMB-GenomicsLA1 Raid6 32000.0GB 00/00/02 Initializing(0.0%)
4 IMB-GenomicsLA2 IMB-GenomicsLA2 Raid6 32000.0GB 00/00/03 Initializing(0.0%)
5 IMB-GenomicsLA3 IMB-GenomicsLA3 Raid6 32000.0GB 00/00/04 Need Init
6 MOL-EXTBMC MOL-EXTBMC Raid6 32000.0GB 00/00/05 Need Init
=====
GuiErrMsg<0x00>: Success.
[root@bmc-pcfs1 ~]#
```

## 7.5 Create partitions (for use with Btrfs)



```
[root@bmc-pcfs1 ~]# echo $(ls -l /dev/sd?)
/dev/sda /dev/sdb /dev/sdc /dev/sdd /dev/sde /dev/sdf /dev/sdg /dev/sdh
[root@bmc-pcfs1 ~]#

[root@bmc-pcfs1 ~]# for i in {c..h} ; do D=/dev/sd$i ; /sbin/parted -s $D mklabel gpt
; ( echo 'mkpart ext4 1 -1' ; echo 'quit' ) | /sbin/parted $D ; done
GNU Parted 3.1
Using /dev/sdc
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart ext4 1 -1
(parted) quit
Information: You may need to update /etc/fstab.

GNU Parted 3.1
Using /dev/sdd
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart ext4 1 -1
(parted) quit
Information: You may need to update /etc/fstab.

GNU Parted 3.1
Using /dev/sde
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart ext4 1 -1
(parted) quit
Information: You may need to update /etc/fstab.

GNU Parted 3.1
Using /dev/sdf
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart ext4 1 -1
(parted) quit
Information: You may need to update /etc/fstab.

GNU Parted 3.1
Using /dev/sdg
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart ext4 1 -1
(parted) quit
Information: You may need to update /etc/fstab.

GNU Parted 3.1
Using /dev/sdh
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mkpart ext4 1 -1
(parted) quit
Information: You may need to update /etc/fstab.
[root@bmc-pcfs1 ~]#
```

## 7.6 File system configuration for the HPC file server (XFS)

Create partitions on the hardware RAID volumes, and physical volumes and a volume group.

```
[root@fbv-one ~]# for i in {c..d} ; do D=/dev/sd$i ; /sbin/parted -s $D mklabel gpt ;
( echo 'mkpart lvm 1 -1' ; echo 'quit' ) | /sbin/parted $D ; pvcreate $D ; done
...
```



```
[root@bmc-one ~]# vgcreate vgpool /dev/sdc1 /dev/sdd1
```

Put the Logical Volumes on the separate hardware RAID volumes.

```
[root@fbv-one ~]# lvcreate -l 100%PVS -n CyanLV vgpool /dev/sdc1
Logical volume "CyanLV" created.
[root@fbv-one ~]# lvcreate -l 100%PVS -n MagentaLV vgpool /dev/sdd1
Logical volume "MagentaLV" created.
[root@fbv-one ~]#
```

Create XFS file systems on the logical volumes

```
[root@fbv-one ~]# mkfs.xfs /dev/mapper/vgpool-CyanLV -L Cyan
meta-data=/dev/mapper/vgpool-CyanLV isize=256    agcount=8, agsize=268435455 blks
         =                               sectsz=512   attr=2, projid32bit=1
         =                               crc=0      finobt=0
data      =                               bsize=4096  blocks=1953123328, imaxpct=5
         =                               sunit=0    swidth=0 blks
naming    =version 2                       bsize=4096  ascii-ci=0 ftype=0
log       =internal log                   bsize=4096  blocks=521728, version=2
         =                               sectsz=512  sunit=0 blks, lazy-count=1
realtime  =none                            extsz=4096  blocks=0, rtextents=0
[root@fbv-one ~]# mkfs.xfs /dev/mapper/vgpool-MagentaLV -L Magenta
meta-data=/dev/mapper/vgpool-MagentaLV isize=256  agcount=8, agsize=268435455 blks
         =                               sectsz=512  attr=2, projid32bit=1
         =                               crc=0      finobt=0
data      =                               bsize=4096  blocks=1953123328, imaxpct=5
         =                               sunit=0    swidth=0 blks
naming    =version 2                       bsize=4096  ascii-ci=0 ftype=0
log       =internal log                   bsize=4096  blocks=521728, version=2
         =                               sectsz=512  sunit=0 blks, lazy-count=1
realtime  =none                            extsz=4096  blocks=0, rtextents=0
[root@fbv-one ~]# mkdir /cyan /magenta
[root@fbv-one ~]# echo "LABEL=Cyan    /cyan    xfs    defaults,noatime    1 2"
>>/etc/fstab
[root@fbv-one ~]# echo "LABEL=Magenta /magenta xfs    defaults,noatime    1 2"
>>/etc/fstab
[root@fbv-one ~]# mount -a -t xfs
[root@fbv-one ~]# df -t xfs
Filesystem            1K-blocks    Used    Available    Use%    Mounted on
/dev/mapper/vgpool-CyanLV  7810406400  33056  7810373344    1%    /cyan
/dev/mapper/vgpool-MagentaLV 7810406400  33056  7810373344    1%    /magenta
[root@fbv-one ~]#
```

## 7.7 Create file systems (Btrfs)

```
[root@bmc-pcfs1 ~]# mkfs.btrfs -f -L IMG-GenomicsCJR /dev/sdc1
btrfs-progs v3.19.1
See http://btrfs.wiki.kernel.org for more information.
```

```
Turning ON incompat feature 'extref': increased hardlink limit per file to 65536
Turning ON incompat feature 'skinny-metadata': reduced-size metadata extent refs
fs created label IMG-GenomicsCJR on /dev/sdc1
    nodesize 16384 leafsize 16384 sectorsize 4096 size 14.55TiB
```

```
[root@bmc-pcfs1 ~]# mkfs.btrfs -L FBV-MSImaging /dev/sdd1
btrfs-progs v3.19.1
See http://btrfs.wiki.kernel.org for more information.
```

```
Turning ON incompat feature 'extref': increased hardlink limit per file to 65536
Turning ON incompat feature 'skinny-metadata': reduced-size metadata extent refs
fs created label FBV-MSImaging on /dev/sdd1
    nodesize 16384 leafsize 16384 sectorsize 4096 size 29.10TiB
```

```
[root@bmc-pcfs1 ~]# mkfs.btrfs -L IMB-GenomicsLA1 /dev/sde1
btrfs-progs v3.19.1
See http://btrfs.wiki.kernel.org for more information.
```

```
Turning ON incompat feature 'extref': increased hardlink limit per file to 65536
```





```
Turning ON incompat feature 'skinny-metadata': reduced-size metadata extent refs
fs created label IMB-GenomicsLA1 on /dev/sde1
   nodesize 16384 leafsize 16384 sectorsize 4096 size 29.10TiB
[root@bmc-pcfs1 ~]# mkfs.btrfs -L IMB-GenomicsLA2 /dev/sdf1
btrfs-progs v3.19.1
See http://btrfs.wiki.kernel.org for more information.
```

```
Turning ON incompat feature 'extref': increased hardlink limit per file to 65536
Turning ON incompat feature 'skinny-metadata': reduced-size metadata extent refs
fs created label IMB-GenomicsLA2 on /dev/sdf1
   nodesize 16384 leafsize 16384 sectorsize 4096 size 29.10TiB
[root@bmc-pcfs1 ~]# mkfs.btrfs -L IMB-GenomicsLA3 /dev/sdg1
btrfs-progs v3.19.1
See http://btrfs.wiki.kernel.org for more information.
```

```
Turning ON incompat feature 'extref': increased hardlink limit per file to 65536
Turning ON incompat feature 'skinny-metadata': reduced-size metadata extent refs
fs created label IMB-GenomicsLA3 on /dev/sdg1
   nodesize 16384 leafsize 16384 sectorsize 4096 size 29.10TiB
[root@bmc-pcfs1 ~]# mkfs.btrfs -L MOL-EXTBMC /dev/sdh1
btrfs-progs v3.19.1
See http://btrfs.wiki.kernel.org for more information.
```

```
Turning ON incompat feature 'extref': increased hardlink limit per file to 65536
Turning ON incompat feature 'skinny-metadata': reduced-size metadata extent refs
fs created label MOL-EXTBMC on /dev/sdh1
   nodesize 16384 leafsize 16384 sectorsize 4096 size 29.10TiB
```

## 7.8 Mount file systems (Btrfs)

With too many files and snapshots btrfs mount will time out unless changed.

```
[root@bmc-pcfs1 ~]# for i in IMB-GenomicsCJR FBV-MSImaging IMB-GenomicsLA{1,2,3} MOL-EXTBMC ; do
echo $i ; mkdir -p /data/$i ; echo LABEL=$i /data/$i btrfs compress,noatime,x-systemd.device-
timeout=0 1 2 >>/etc/fstab ; done
IMB-GenomicsCJR
FBV-MSImaging
IMB-GenomicsLA1
IMB-GenomicsLA2
IMB-GenomicsLA3
MOL-EXTBMC
[root@bmc-pcfs1 ~]#
[root@bmc-pcfs1 ~]# grep btrfs /etc/fstab
LABEL=IMB-GenomicsCJR /data/IMB-GenomicsCJR btrfs compress,noatime,x-systemd.device-timeout=0 1 2
LABEL=FBV-MSImaging /data/FBV-MSImaging btrfs compress,noatime,x-systemd.device-timeout=0 1 2
LABEL=IMB-GenomicsLA1 /data/IMB-GenomicsLA1 btrfs compress,noatime,x-systemd.device-timeout=0 1 2
LABEL=IMB-GenomicsLA2 /data/IMB-GenomicsLA2 btrfs compress,noatime,x-systemd.device-timeout=0 1 2
LABEL=IMB-GenomicsLA3 /data/IMB-GenomicsLA3 btrfs compress,noatime,x-systemd.device-timeout=0 1 2
LABEL=MOL-EXTBMC /data/MOL-EXTBMC btrfs compress,noatime,x-systemd.device-timeout=0 1 2
[root@bmc-pcfs1 ~]#
[root@bmc-pcfs1 ~]# mount -a -t btrfs
[root@bmc-pcfs1 ~]# df -h --si /data/*
Filesystem      Size  Used Avail Use% Mounted on
/dev/sdd1        32T   18M   32T   1% /data/FBV-MSImaging
/dev/sdc1        16T   18M   16T   1% /data/IMB-GenomicsCJR
/dev/sde1        32T   18M   32T   1% /data/IMB-GenomicsLA1
/dev/sdf1        32T   18M   32T   1% /data/IMB-GenomicsLA2
/dev/sdg1        32T   18M   32T   1% /data/IMB-GenomicsLA3
/dev/sdh1        32T   18M   32T   1% /data/MOL-EXTBMC
[root@bmc-pcfs1 ~]#
```



For a file server running Samba run this in /etc/rc.d/rc.local to this instead:

```
systemctl stop smb
for i in /data/* ; do mount $i & done
wait
systemctl start smb
```

## 7.9 Create CNAMEs to the shares in DNS for the common service PC file server

Send a mail to domainmaster@uu.se and request changes in DNS.

```
IMB-GenomicsCJR.files.uu.se. in CNAME bmc-pcfs1.bmc.uu.se.
FBV-MSImaging.files.uu.se. in CNAME bmc-pcfs1.bmc.uu.se.
IMB-GenomicsLA1.files.uu.se. in CNAME bmc-pcfs1.bmc.uu.se.
IMB-GenomicsLA2.files.uu.se. in CNAME bmc-pcfs1.bmc.uu.se.
IMB-GenomicsLA3.files.uu.se. in CNAME bmc-pcfs1.bmc.uu.se.
MOL-EXTBMC.files.uu.se. in CNAME bmc-pcfs1.bmc.uu.se.
```

## 7.10 Activating snapshots for Btrfs every hour (on primary file server)

```
wget http://www.update.uu.se/~jerker/extractor/btrfs.shadow.snapshots.txt -O
/usr/local/bin/btrfs.shadow.snapshots.sh
chmod +x /usr/local/bin/btrfs.shadow.snapshots.sh
echo "1 * * * * root /usr/local/bin/btrfs.shadow.snapshots.sh"
>/etc/cron.d/btrfs.shadow.snapshots
/usr/local/bin/btrfs.shadow.snapshots.sh
ls -la /data/MOL-EXTBMC/.snapshots/
...
drwxr-xr-x 1 root root 32 Sep 14 09:07 /data/MOL-EXTBMC/.snapshots/@GMT-2016.09.14-
07.22.05
```

## 7.11 Preparing for backups with the backup-script (primary file server) and expunge-script (for backup file server)

```
wget http://www.update.uu.se/~jerker/extractor/btrfs.backup.txt -O
/usr/local/bin/btrfs.backup.sh
chmod +x /usr/local/bin/btrfs.backup.sh

wget http://www.update.uu.se/~jerker/extractor/btrfs.expunge.snapshots.txt -O
/usr/local/bin/btrfs.expunge.snapshots.sh
chmod +x /usr/local/bin/btrfs.expunge.snapshots.sh

echo "1 1 * * * root /usr/local/bin/btrfs.expunge.snapshots.sh"
>/etc/cron.d/btrfs.expunge.snapshots
```

# 8 Join the Active Directory using Winbind

```
yum install -y oddjob-mkhomedir oddjob samba-winbind-clients samba-winbind samba-
common realmd samba samba-common samba-winbind-krb5-locator pam_krb5
```

```
ntpdate ntp.uu.se
systemctl enable ntpd
systemctl start ntpd
```

- This does not seem to work very well although it is recommended:

```
echo SECRET | realm join USER.UU.SE --automatic-id-mapping=no --
user=_jny25782@user.uu.se --computer-
ou=OU=Clients,OU=BMC1,OU=LocalIT,DC=user,DC=uu,DC=se --client-software=winbind --
membership-software=samba
```



- Old school authconfig works as expected:

```
authconfig --disablefingerprint '--winbindjoin=_jny25782%SECRET' --enablewinbind --
enablewinbindauth --smbsecurity=ads --smbrealm=USER.UU.SE --enablemkhomedir --
enablewinbindusedefaultdomain --smbworkgroup=USER --winbindtemplateshell=/bin/bash --
enablekrb5kdc dns --enablekrb5realmdns --enablekrb5 --krb5realm=USER.UU.SE --updateall
--enablewinbindoffline
```

- The computer account object shows up in **New Computers** in the USER-AD. Move it to the correct place.

## 9 Correcting the defaults in smb.conf to CentOS 6 default behavior

Uncomment the sections [homes] and [printers]  
Change this:

```
template homedir = /home/%D/%U
idmap config * : range = 1000-9999
```

Add these (emum might not be needed). The reason for using the RID backend instead of AD is that there is a delay between creating groups and users in USER-AD before they get their uidNumber and gidNumber. RID just works.

```
idmap config USER : backend = rid
idmap config USER : range = 10000000 - 1083741824
idmap config USER : base_rid = 0
winbind enum users = true
winbind enum groups = true
```

Create shares in smb.conf

```
for share in IMB-GenomicsLA1 IMB-GenomicsLA2 IMB-GenomicsLA3 IMB-
GenomicsCJR FBV-MSImaging ; do
(
echo "[$share]"
echo "vfs objects = shadow_copy2"
echo "shadow:sort = desc"
echo "comment = $share"
echo "path = /data/$share"
echo "public = yes"
echo "writable = yes"
echo "printable = no"
echo "browsable = no"
) >>/etc/samba/smb.conf
done
```

Start it up.

```
systemctl enable smb.service
systemctl start smb.service
```

Then reboot and check that things are working.



## I 0 Activating backup

Read the other SOP for how the backups work.

## I 1 Permissions

Add the groups ORG-Dir-Sharename-RW to USER-AD and include the correct Akka-groups.

Allow only that group access:

```
cd /data
for a in * ; do
  D=$(echo $a | sed 's/...-//')
  P=$(echo $a | sed 's/-.*//')
  GR=$P-Dir-$D-RW
  chgrp $GR $a/$D $a/.snapshots
  chmod 770 $a/$D $a/.snapshots
  chown root $a/$D $a/.snapshots
  ls -lad $a/$D $a/.snapshots
done
```

## I 2 Optional restricted rsync access

Activate *rsync*

```
cp -av /usr/share/doc/rsync-3.0.9/support/rrsync /usr/bin/rrsync
chmod +x /usr/bin/rrsync
```

Add user into */etc/ssh/sshd\_config* among *AllowUsers* and deny access using only password authentication.

```
AllowUsers root jny25782
Match User jny25782
  PasswordAuthentication no
```

Create *~/.ssh/authorized\_keys* for the user

```
su - jny25782
mkdir .ssh
emacs -nw .ssh/authorized_keys
```

Enter the key and set the command to use when using this key. This entry will restrict access to only the */data/* directory on the server when using the specified key.

```
command="/usr/bin/rrsync /data/",no-port-forwarding,no-x11-forwarding,no-agent-forwarding ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAQEAplq70LzWrtBY8zPMBvv36DnDorqwPrB4reLjsY2dMeHeXZGe
F1/A0sdFewpykGa8YE9bdbSsygCkQH3cH79X9uYG+K0ojTtwVEC/YBETMdqxaYu9ZKgKXYTVVda
6Dwr77jAbBNGF0WuA1qR/Aq1EHaJWAtT7h6aEgsY45VGmJWYEeFDqUZLZmBU/6WnkMjWoAplye3m
YAaXmnaiaFt1Xl/WZv18pi+umq4Vy1Dn4K+DN63bEmclGhrGaY0To6H2E1PFI1yxv1FylaAFwWg
3hFlvrkVpkEpYQ+heigMWNda5mr+YVWwERYFLI5MMWazlU50ViZii/hTcS45B5qnRw==
jerker@gforce.bmc.uu.se
```

## I 3 Optional Time Machine backup



This seems to work, but is maybe not really recommended. It is very hard to really sure that the backups are working as they should. It is better to keep the files on a file server than use backup for the clients. Our impression is that Time Machine is not really enterprise ready.

This setup however fixes one of the major drawbacks with time machine backup – after a snapshot has been taken, the client cannot change or delete the contents of the backup. This is very important for ransomware.

- Create a group in USER-AD for limiting which users have access to the AFP time machine shar. Example: **bmc-it-timemachine**
- Install dependencies

```
yum install -y rpm-build gcc make wget avahi-devel cracklib-devel dbus-devel dbus-glib-devel libacl-devel libattr-devel libdb-devel libevent-devel libgcrypt-devel bison docbook-style-xsl flex dconf perl-IO-Socket-INET6 perl-Socket
```

- Check Netatalk homepage for links to latest source package  
[http://netatalk.sourceforge.net/wiki/index.php/Netatalk\\_3.1.10\\_SRPMS\\_for\\_Fedora\\_and\\_CentOS](http://netatalk.sourceforge.net/wiki/index.php/Netatalk_3.1.10_SRPMS_for_Fedora_and_CentOS)
- Download and build and install

```
wget http://www003.upp.so-net.ne.jp/hat/files/netatalk-3.1.10-0.1.1.fc26.src.rpm  
rpm -ivh netatalk-3.1.10-0.1.1.fc26.src.rpm  
rpmbuild -bb rpmbuild/SPECS/netatalk.spec  
rpm -ivh ~/rpmbuild/RPMS/x86_64/netatalk-3.1.10-0.1.1.e17.centos.x86_64.rpm
```

- Prepare directory in the share for storing backups

```
mkdir /data/BMC-TimeMachine1/backups  
chgrp bmc-it-timemachine /data/BMC-TimeMachine1/backups  
chmod +t /data/BMC-TimeMachine1/backups/  
chmod 770 /data/BMC-TimeMachine1/backups/
```

- Create Netatalk configuration

```
cat >> /etc/avahi/services/afpd.service << EOF <?xml version="1.0" standalone='no'?>  
<!DOCTYPE service-group SYSTEM "avahi-service.dtd">  
<service-group>  
<name replace-wildcards="yes">%h</name>  
<service>  
<type>_afpovertcp._tcp</type>  
<port>548</port>  
</service>  
<service>  
<type>_device-info._tcp</type>  
<port>0</port>  
<txt-record>model=Xserve</txt-record>  
</service>  
</service-group>  
EOF
```



- Specify the directory and the valid group. Prefix group with @

```
cat >>/etc/netatalk/afp.conf <<EOF
[Time Machine]
path = /data/BMC-TimeMachine1/backups
valid users = @bmc-it-timemachine
time machine = yes
EOF
```

```
cat >> /etc/netatalk/afpd.conf << EOF
- -transall -uamlist uams_randnum.so,uams_dhx.so,uams_dhx2.so -nosavepassword -
advertise_ssh
EOF
```

- Specify the directory and group again

```
cat >> /etc/netatalk/AppleVolumes.default << EOF
/data/BMC-TimeMachine1/backups allow:@bmc-it-timemachine TimeMachine
options:usedots,upriv,tm dperm:0775 fperm:0660 cnidscheme:dbd volsizelimit:200000
EOF
```

- Enable and start servers.

```
systemctl enable avahi-daemon
systemctl enable netatalk
systemctl start avahi-daemon
systemctl start netatalk
```



- Remember to open port 548/tcp in router filters
- Activate Wide Area Bonjour in static DNS:

```
b._dns-sd._udp IN PTR @
lb._dns-sd._udp IN PTR @
db._dns-sd._udp IN PTR @
```

- Have fun with Bookmarks in Safari:

```
_http._tcp PTR BMC-IT._http._tcp
BMC-IT._http._tcp SRV 0 0 80 it.bmc.uu.se.
TXT "path="/" "url=http://it.bmc.uu.se/"
```

- Setup the shares:

```
_services._dns-sd._udp IN PTR _afpovertcp._tcp
_services._dns-sd._udp IN PTR _rfb._tcp
_services._dns-sd._udp IN PTR _device-info._tcp
_services._dns-sd._udp IN PTR _http._tcp

_rfb._tcp PTR bmc-pcfs1._rfb._tcp
bmc-pcfs1._rfb._tcp SRV 0 0 5901 bmc-pcfs1.bmc.uu.se.
bmc-pcfs1._rfb._tcp TXT ""
_afpovertcp._tcp PTR bmc-pcfs1._afpovertcp._tcp
bmc-pcfs1._afpovertcp._tcp SRV 0 0 548 bmc-pcfs1.bmc.uu.se.
bmc-pcfs1._afpovertcp._tcp TXT ""
_device-info._tcp PTR bmc-pcfs1._device-info._tcp
bmc-pcfs1._device-info._tcp SRV 0 0 0 bmc-pcfs1.bmc.uu.se.
bmc-pcfs1._device-info._tcp TXT "model=Xserve"
```

## 14 Optional OSC Inventory agent

Follow the instructions in the BMC-IT FAQ to install the OCS Inventory agent on the servers:

<http://it.bmc.uu.se/faq/?q=ocs.linux>