

Creating a Cray System Management Workstation (SMW) Bootable Backup Drive

This technical note provides the procedures to create a System Management Workstation (SMW) bootable backup drive. The purpose of this backup drive is to provide a bootable replacement drive to replace the primary drive in the event of the primary drive failure. These procedures support both the SMW hardware that is Restriction of Hazardous Substances (RoHS) compliant and the older SMW (nonRoHS-compliant) hardware. [Setting Up the Bootable Backup Drive as an Alternate Boot Device on page 8](#) of this document describes how to modify a bootable backup drive in order to boot from and run the SMW from the backup root partition.

Send e-mail to docs@cray.com with any comments that will help us to improve the accuracy and usability of this document. Be sure to include the title and number of the document with your comments. We value your comments and will respond to them promptly.

1.1 Creating an SMW Bootable Backup Drive

Complete the procedure in this section to create a bootable backup drive. The purpose of this backup drive is to provide a bootable replacement drive to replace the primary drive in the event of the primary drive failure.

When this procedure is completed, the backup drive on the SMW will be a bootable replacement for the primary drive when the backup drive is plugged in as or cabled as the primary drive.

Procedure 1. Creating an SMW bootable backup drive



Caution: The disk device names shown in this procedure are provided as examples only. Substitute the correct disk devices for your system. For example, On a RoHS-compliant SMW with three SMW disks, the boot disk is `/dev/sda` and the bootable backup disk is `/dev/sdc`; on a RoHS-compliant SMW with two SMW disks, the boot disk is `/dev/sda` and the bootable backup disk is `/dev/sdb`.

1. Log on to the SMW as `crayadm` and `su` to `root`.

```
crayadm@smw:~> su - root
smw:~ #
```

2. If the backup drive disk partition table already exists and the partition table on the backup drive matches the partition table that is on the primary boot drive, skip this step; otherwise, create the backup drive disk partition table.

Note: For optimal performance, the source and destination disks should be on different buses; drive slots 1 and 2 are on a different bus than drive slots 3 and 4.

In this example, the partition table consists of the following:

- Slice 1: 4 GB Linux swap partition
 - Slice 2: Balance of disk space used for the root file system
- a. Use the `fdisk` command to display the boot disk partition layout.

```
smw:~ # fdisk -lu /dev/sda
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes

   Device   Boot    Start        End    Blocks   Id  System
/dev/sda1             63      8401994    4200966   82  Linux swap / Solaris
/dev/sda2    *   8401995   625137344   308367675   83  Linux
```

- b. Use the `fdisk` command to configure the bootable backup disk partition layout. Set the bootable backup disk partition layout to match the boot disk partition layout. First, clear all of the old partitions using the `d` command within `fdisk`, then create a Linux swap and a Linux partition, and then write your changes to the disk. For help, type `m` within `fdisk` (see the following sample output).

```
smw:~ # fdisk -u /dev/sdc

The number of cylinders for this disk is set to 38913.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
 1) software that runs at boot time (e.g., old versions of LILO)
 2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK).

Command (m for help): p
Disk /dev/sdc: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes

   Device   Boot    Start        End    Blocks   Id  System
/dev/sdc1             63      8401994    4200966   82  Linux swap
/dev/sdc2             8401995   625105214   308351610   83  Linux
```

```
Command (m for help): d
Partition number (1-5): 2
Command (m for help): d
Selected partition 1
Command (m for help): n
Command action
```

```

e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First sector (63-625105215, default 63): (Press the Enter key)
Using default value 63
Last sector or +size or +sizeM or +sizeK (63-625105215, default 625105215): 8401994

Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 82
Changed system type of partition 1 to 82 (Linux swap / Solaris)

Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 2
First sector (8401995-625105215, default 8401995): (Press the Enter key)
Using default value 8401995
Last sector or +size or +sizeM or +sizeK (8401995-625105215, default 625105215): (Press the Enter key)
Using default value 625105215

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

```

c. Display the boot backup disk partition layout.

```

smw:~ # fdisk -lu /dev/sdc
Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders, total 625142448 sectors
Units = sectors of 1 * 512 = 512 bytes

```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdc1		63	8401994	4200966	82	Linux swap / Solaris
/dev/sdc2	*	8401995	625137344	308367675	83	Linux

3. Initialize the swap device.

```

smw:~ # mkswap /dev/sdc1

```

4. Standardize the /etc/fstab and grub disk device names.

The device names written into the /boot/grub/menu.lst file by the installation process are UDEV-based names (something like /dev/disk/by-id/scsi-SATA_ST3320620AS_922J3-part2 or /dev/disk/by-id/ata-ST3320620A_9QFA85PV-part2) rather than the more commonly used device names (something like /dev/sda2 or /dev/hda2). In the following procedures, edit the /boot/grub/menu.lst file to change ONLY the long UDEV-based name to the shorter, commonly used device name reflected in the output of the df command on your system.

If the device names have already been standardized, skip to [step 5](#).



Caution: Mistakes in the `/boot/grub/menu.lst` file will affect your ability to boot the SMW.

- a. SLES 10 SP1 sets up `/etc/fstab` and `/boot/grub/menu.lst` with UDEV-based names for the root device. For example:

```
smw:~ # head -2 /etc/fstab
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 /      ext3  acl,user_xattr  1 1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap  swap  defaults        0 0
```

```
smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title Cray SMW - SLES 10 SP1 base
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
    root=/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 vga=791 \
    resume=/dev/sdal splash=silent showopts
    initrd /boot/initrd-2.6.16.54-0.2.12-smp

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- Cray SMW SLES 10 SP1 base
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
    root=/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 vga=normal \
    showopts ide=nodma apm=off acpi=off noresume edd=off 3
    initrd /boot/initrd-2.6.16.54-0.2.12-smp
```

- b. Execute the `df` command to get the name of the device that you will use in the `/etc/fstab` and `/boot/grub/menu.lst` files to replace the long UDEV-based device name. Then edit your `/etc/fstab` and `/boot/grub/menu.lst` files appropriately.

- 1) Execute the `df` command to get the name of the device that you will use in the `/etc/fstab` and `/boot/grub/menu.lst` files to replace the long UDEV-based device name. For example:

```
smw:# df
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/sda2        303528624    40652904 247457340  15% /
udev             1030780         460    1030320   1% /dev
```

- 2) Save a copy of your `/etc/fstab` and `/boot/grub/menu.lst` files.

```
smw:# cp -p /etc/fstab /etc/fstab.save
smw:# cp -p /boot/grub/menu.lst /boot/grub/menu.lst.save
```

- 3) Edit your `/etc/fstab` file appropriately; you will use the device name (`dev`) you got from the `df` command output. In this example, the "1" and "2" refer to the partition names on the device. Change the following lines, which changes the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2` and changes `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1` to `sda1`. Ensure that your swap is on `sda1`:

```
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 /      ext3  acl,user_xattr  1 1
/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part1 swap  swap  defaults        0 0
```

to:

```
/dev/sda2 /      ext3  acl,user_xattr  1 1
/dev/sda1 swap  swap  defaults        0 0
```

- 4) Edit your `/boot/grub/menu.lst` file appropriately; you will use the device name (`dev`) you got from the `df` command output. Change the long name `disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2` to `sda2`. Change the following lines:

```
kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
root=/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 vga=791 \
resume=/dev/sda1 splash=silent showopts
```

to:

```
kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
root=/dev/sda2 vga=791 resume=/dev/sda1 splash=silent showopts
```

and change the following lines:

```
kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
root=/dev/disk/by-id/scsi-SATA_ST3320620AS_9QF922J3-part2 vga=normal \
showopts ide=nodma apm=off acpi=off noresume edd=off 3
```

to:

```
kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
root=/dev/sda2 vga=normal showopts ide=nodma apm=off acpi=off noresume edd=off 3
```

- 5) Verify that the edited files are correct and match the output of the `df` command.

```
smw:~ # head -2 /etc/fstab
/dev/sda2 /      ext3  acl,user_xattr  1 1
/dev/sda1 swap  swap  defaults        0 0
```

```
smw:~ # more /boot/grub/menu.lst
###Don't change this comment - YaST2 identifier: Original name: linux###
title Cray SMW - SLES 10 SP1 base
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
    root=/dev/sda2 vga=791 resume=/dev/sda1 splash=silent showopts
    initrd /boot/initrd-2.6.16.54-0.2.12-smp
```

```
###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- Cray SMW SLES 10 SP1 base
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.16.54-0.2.12-smp \
    root=/dev/sda2 vga=normal showopts ide=nodma apm=off acpi=off noresume edd=off 3
    initrd /boot/initrd-2.6.16.54-0.2.12-smp
```

- c. Update the grub device table to utilize the standardized drive names and recognize any new drives added since the initial operating system installation.

```
smw:~ # grub-install --recheck /dev/sda
```

The file `/boot/grub/device.map` will now be updated to reflect all drives, utilizing the standardized drive naming. This file can be viewed for verification; for example:

```
smw:~ # cat /boot/grub/device.map
(fd0)    /dev/fd0
(hd0)    /dev/sda
(hd1)    /dev/sdb
```

5. Create a new file system on the backup drive root partition by executing the `mkfs` command.

```
smw:~ # mkfs -t ext3 -I 128 /dev/sdX2
mke2fs 1.40.7.sun3 (28-Feb-2008)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
19292160 inodes, 77140113 blocks
3857005 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
2355 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

This filesystem will be automatically checked every 21 mounts or 180 days, whichever comes first. Use `tune2fs -c` or `-i` to override.

```
smw:~ #
```

6. Mount the new backup root file system on `/mnt`.

```
smw:~ # mount /dev/sdX2 /mnt
```

7. Confirm the running root file system device.

```
smw:~ # df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/sda2              303528624    6438700 281671544   3% /
udev                  1030332         116   1030216    1% /dev
/dev/sdX2              306128812    195568 290505224   1% /mnt
```

The running root file system device is the one mounted on /.

8. Dump the running root file system to the backup drive.

```
smw:~ # cd /mnt
smw:~ # dump 0f - /dev/sda2 | restore rf -
DUMP: WARNING: no file `/etc/dumpdates'
DUMP: Date of this level 0 dump: Fri Sep 11 08:21:16 2009
DUMP: Dumping /dev/sda2 (/) to standard output
DUMP: Label: none
DUMP: Writing 10 Kilobyte records
DUMP: mapping (Pass I) [regular files]
DUMP: mapping (Pass II) [directories]
DUMP: estimated 6432431 blocks.
DUMP: Volume 1 started with block 1 at: Fri Sep 11 08:22:55 2009
DUMP: dumping (Pass III) [directories]
DUMP: dumping (Pass IV) [regular files]
restore: ./lost+found: File exists
./tmp/rstmdir1252675276: (inode 26263561) not found on tape
./tmp/rstmode1252675276: (inode 26263568) not found on tape
DUMP: 39.48% done at 8464 kB/s, finished in 0:07
DUMP: 80.27% done at 8605 kB/s, finished in 0:02
DUMP: Volume 1 completed at: Fri Sep 11 08:35:22 2009
DUMP: Volume 1 6435690 blocks (6284.85MB)
DUMP: Volume 1 took 0:12:27
DUMP: Volume 1 transfer rate: 8615 kB/s
DUMP: 6435690 blocks (6284.85MB)
DUMP: finished in 747 seconds, throughput 8615 kBytes/sec
DUMP: Date of this level 0 dump: Fri Sep 11 08:21:16 2009
DUMP: Date this dump completed: Fri Sep 11 08:35:22 2009
DUMP: Average transfer rate: 8615 kB/s
DUMP: DUMP IS DONE
```

9. Install the GRUB boot loader.

To make the backup drive bootable, you must re-install the grub boot facility on that drive.

- a. Create a unique file on the backup drive to be used to identify that drive to grub boot facility.

```
smw:~ # touch /mnt/THIS_IS_SDx
```

- b. Invoke the grub boot utility. Within the grub boot utility:

- 1) Execute the find command to locate the drive designation that grub uses.
- 2) Select the drive to which the boot blocks will be installed with the root command.

- 3) Use the `setup` command to set up and install the `grub` boot blocks on that drive.

Note: The Linux `grub` utility and boot system ALWAYS refer to drives as `hd`, regardless of the actual type of drives.

For example:

```
smw:~ # grub
GNU GRUB version 0.97 (640K lower / 3072K upper memory)
[ Minimal BASH-like line editing is supported. For the first word, TAB^[]
lists possible command completions. Anywhere else TAB lists the possible
completions of a device/filename. ]
grub> find /THIS_IS_SDX
(hd1,1)
grub> root (hd1,1)
Filesystem type is ext2fs, partition type 0x83
grub> setup (hd1)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd1)"... 16 sectors are embedded.
succeeded
Running "install /boot/grub/stage1 (hd1) (hd1)1+16 p (hd1,1)/boot/grub/stage2
/boot/grub/menu.lst"... succeededDone.
grub> quit
```

10. Unmount the backup root partition.

```
smw:~ # umount /dev/sdx2
```

The drive is now bootable once plugged in or cabled as the primary drive.

1.2 Setting Up the Bootable Backup Drive as an Alternate Boot Device

The following procedure modifies a bootable backup drive, generated in [Procedure 1 on page 1](#), in order to boot from and run the SMW from the backup root partition.

Important: To boot from this backup drive, the primary boot drive must still be operable and able to boot the `grub` boot blocks installed.

Procedure 2. Setting up the bootable backup drive as an alternate boot device

Caution: The disk device names shown in this procedure are provided as examples only. Substitute the correct disk devices for your system. For example, On a RoHS-compliant SMW with three SMW disks, the boot disk is `/dev/sda` and the bootable backup disk is `/dev/sdc`; on a RoHS-compliant SMW with two SMW disks, the boot disk is `/dev/sda` and the bootable backup disk is `/dev/sdb`.

1. Mount the backup drive's root partition.

```
smw:~ # mount /dev/sdx2 /mnt
```

2. Create a new boot entry in the `/boot/grub/menu.lst` file. This entry should be a duplicate of the primary boot entry with the following changes:

- Modify the title to uniquely identify the backup boot entry.
- Modify the `root (hd0,1)` directive to reflect the Grub name of the backup drive. If you do not know the Grub name of the backup drive, it is provided in the `/boot/grub/device.map` file on the primary drive.
- Modify the `root=` and `resume=` specifications to reference the backup drive device.

An example `/boot/grub/menu.lst` file follows. Note the new entry at the end of the file. This example references `/dev/sda` as the primary drive and `/dev/sdc` as the backup drive.

```
smw:~ # cat /boot/grub/menu.lst
# Modified by YaST2. Last modification on Tue Aug 25 10:43:18 CDT 2009
default 0
timeout 8
gfxmenu (hd0,1)/boot/message
##YaST - activate

###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 10 SP1
  root (hd0,1)
  kernel /boot/vmlinuz-2.6.16.54-0.2.8-smp root=/dev/sda2 vga=0x317 \
resume=/dev/sda1 splash=silent showopts
  initrd /boot/initrd-2.6.16.54-0.2.8-smp

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 10 SP1
  root (hd0,1)
  kernel /boot/vmlinuz-2.6.16.54-0.2.8-smp root=/dev/sda2 vga=0x317\
showopts ide=nodma apm=off acpi=off noresume edd=off 3
  initrd /boot/initrd-2.6.16.54-0.2.8-smp

###Don't change this comment - YaST2 identifier: Original name: floppy###
title Floppy
  rootnoverify (hd0)
  chainloader (fd0)+1
### New entry allowing a boot of the back-up drive when the primary drive
### is still present.
title BACK-UP DRIVE - SUSE Linux Enterprise Server 10 SP1
```

```

root (hd1,1)
kernel /boot/vmlinuz-2.6.16.54-0.2.8-smp root=/dev/sdc2 vga=0x317 \
resume=/dev/sdc1 splash=silent showopts
initrd /boot/initrd-2.6.16.54-0.2.8-smp

```

3. Modify the backup drive's `/etc/fstab` file to reference the secondary drive slot rather than the first drive slot.

- a. Examine the backup drive's `fstab` file.

```

smw:~ # cat /mnt/etc/fstab
/dev/sda1 / ext3 acl,user_xattr 1 1
/dev/sda2 swap swap defaults 0 0
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
usbfs /proc/bus/usb usbfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0
/dev/fd0 /media/floppy auto noauto,user,sync 0 0

```

- b. Edit the `/mnt/etc/fstab` file, changing `/dev/sda1` and `/dev/sda2` to reference the backup drive. In the following example, the backup drive is `/dev/sdc`.

```

smw:~ # vi /mnt/etc/fstab
/dev/sdc1 / ext3 acl,user_xattr 1 1
/dev/sdc2 swap swap defaults 0 0
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
usbfs /proc/bus/usb usbfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0
/dev/fd0 /media/floppy auto noauto,user,sync 0 0

```

4. Unmount the backup drive.

```
smw:~ # umount /dev/sdX2
```

The SMW can now be shut down and rebooted. Upon display of the **Please select boot device** prompt, select the **BACK-UP DRIVE - SUSE Linux Enterprise Server 10 SP1** entry to boot the backup root partition.

1.3 SMW Recovery



Caution: Booting off the bootable backup disk is intended only for emergency use in the event of failure or loss of data on the primary disk.

Procedure 3. RoHS-compliant SMW recovery

To recover a RoHS-compliant SMW, you must reorder the drives at the front of the SMW. No BIOS or software configuration changes are required.

1. Shutdown the OS on the SMW, if possible.
2. Power the SMW off.

3. Unplug the power cord.
4. Open the disk drive access door, which is on the front of the SMW.
5. Remove the primary disk from its slot. The primary disk is located at the bottom of the column of disk drives at the front of the SMW.
6. Remove the bootable backup disk and place it in the primary disk slot.
7. Press the reset button, if required.
8. Boot the SMW.

Procedure 4. NonRoHS-compliant SMW recovery

To recover a nonRoHS-compliant SMW, you must re-cable the SMW drives. No BIOS or software configuration changes are required.

1. Perform a shutdown of the OS on the SMW, if possible.
2. Power the SMW off.
3. Unplug the power cord.
4. Remove the side panel (left panel as you face the front of the SMW) by removing the three panel screws.
5. Locate the two IDE disks inside the bottom, front of the SMW. The disk closest to the outside is the primary disk. The disk behind it is the bootable backup disk.
6. Disconnect the cable connection from the primary drive (primary master cable) and connect it to the backup drive. Depending upon its state, the primary drive can be left disconnected or it can be connected to the other cable (secondary master cable). These connections are delicate and care should be taken when handling them.
7. Place the cover back on the SMW, plug in the power cord, and turn on the power switch on the back panel.
8. Press the reset button, if required.
9. Boot the SMW.

© 2009, 2010 Cray Inc. All Rights Reserved. This document or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Inc.

GNU is a trademark of The Free Software Foundation. Linux is a trademark of Linus Torvalds. Solaris is a trademark of Sun Microsystems, Inc. in the United States and other countries. SUSE is a trademark of Novell, Inc. All other trademarks are the property of their respective owners.

Version 4.0 Published October 2009 Supports a System Management Workstation (SMW) running the SMW 4.0 or SMW 3.1 release software.

Version 4.01 Published April 2010 Contains an additional step to Procedure 1 to set up a swap area (mkswap) on the bootable backup drive. Supports a System Management Workstation (SMW) running the SMW 4.0 or SMW 3.1 release software.