



Windows® on SGI® Altix® UV Software
Install Guide

007-5762-002

COPYRIGHT

© 2011, SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is "commercial computer software" provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

SGI, Altix, and the SGI logo are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries.

Microsoft and Windows are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

New Features in This Manual

This rewrite of the *Windows on SGI Altix UV Software Install Guide* provides installation information for the Windows Server 2008 R2 Datacenter.

Major Documentation Changes

Performed the following:

- Updated Chapter 1, "Altix UV Software Installation" on page 1.
- Updated Chapter 2, "Disk Preparation for Windows" on page 13.
- Added Chapter 3, "Using the Drive Setup Utility" on page 17.

Record of Revision

Version	Description
001	March 2011 Original Printing.
002	April 2011 Updated.

Contents

About This Manual	xiii
Obtaining Publications	xiii
Related Publications and Other Sources	xiii
Conventions	xiv
Reader Comments	xiv
1. Altix UV Software Installation	1
Overview	1
Installing Windows	1
2. Disk Preparation for Windows	13
Disk Setup with Windows	13
Disk Setup with Linux	14
3. Using the Drive Setup Utility	17
Index	21

Figures

Figure 1-1	Altix UV 1000 Individual Rack Unit (IRU)	3
Figure 1-2	Altix UV 1000 BaseIO Riser Enabled Blade Front Panel Example	4
Figure 1-3	Altix UV 100 Individual Rack Unit (IRU)	4
Figure 1-4	Console Prompt for Boot Menu Screen	6
Figure 1-5	Device Manager Screen	7
Figure 1-6	Device Manager Advanced Option Screen	8
Figure 1-7	SGI Custom Configuration Option Example	8
Figure 1-8	Advanced Screen APIC Mode Option	9
Figure 1-9	APIC Mode Extended Example Screen	11

Procedures

Procedure 1-1	Installing Windows Server Software	1
Procedure 2-1	Disk Setup with Windows to Prepare Disk Drives	13
Procedure 2-2	Disk Setup with Linux to Prepare Disk Drives	14
Procedure 3-1	Using the MPT Setup Utility to Activate a RAID Volume	17

About This Manual

Your SGI Altix UV 100 or SGI Altix UV 1000 system comes with pre-installed software. This manual describes how to reinstall it if necessary.

Obtaining Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at: <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- You can also view man pages by typing `man title` on a command line.

Related Publications and Other Sources

This section describes documentation you may find useful, as follows:

- *SGI Altix UV System Management Node Administrator's Guide*

This guide describes the system management node (SMN) for SGI Altix UV 1000 and SGI Altix UV 100 series systems. It provides information on how to install, configure, and use software on the SMN to manage and monitor SGI Altix UV systems.

- *SGI Altix UV CMC Controller Software User's Guide*

Describes how to use the controller commands on your chassis manager controller (CMC) to monitor and manage SGI Altix UV 100 and SGI Altix UV 1000 systems.

- *SGI Altix UV 1000 System User's Guide*

This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix UV 1000 system. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.

- *SGI Altix UV 100 System User's Guide*

This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix UV 100 system. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:

techpubs@sgi.com

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

SGI
Technical Publications
46600 Landing Parkway
Fremont, CA 94538

SGI values your comments and will respond to them promptly.

Altix UV Software Installation

This chapter describes how to install Windows Server 2008 R2 Datacenter software on an SGI Altix UV 100 or SGI Altix UV 1000 system.

Note: This manual does not apply to SGI Altix UV 10 systems.

Overview

To install Windows Server 2008 R2 software on an SGI Altix UV 100 or SGI Altix UV 1000 system, you need to disable the `NL5_enable` setting (see Procedure 1-1 on page 1), ensure that the BaseIO BIOS version 1.45 (or later) is installed, and ensure that the system is either in `legacy` or `x2apic` mode. The install procedure assumes a clean disk. You need to completely wipe the partition tables, both MBR and GPT (see Chapter 2, "Disk Preparation for Windows" on page 13). Once the install is complete, you need to enable the `NL5_enable` setting and change the system to `extended` mode. If you do not completely remove the partition tables before doing the Windows server install procedure, the system may hang or fail on reboot.

Installing Windows

This section describes how to install Windows Server 2008 R2 Datacenter on an Altix UV 100 or Altix UV 1000 system.

Procedure 1-1 Installing Windows Server Software

To install Windows Server 2008 R2 software on an SGI Altix UV 100 or SGI Altix UV 1000 system, perform the following steps:

- 1.

Note: The Windows Server 2008 R2 SP1 release includes these hotfixes by default so you can skip this step if you are installing the SP1 release.

The hotfixes listed below are available from Microsoft, Inc. They are needed later in this procedure. The hotfixes have passwords which will be contained in emails

from Microsoft. The passwords are needed to extract the files from the zip archives. Put zipped software on a USB stick or burn them into a data CD. The hotfixes required are, as follows:

- <http://support.microsoft.com/kb/980598>
 - <http://support.microsoft.com/kb/2303458>
 - <http://support.microsoft.com/kb/2398906>
2. Make sure the system disk drives are setup for Windows. Windows 2008 R2 Datacenter, when booted from EFI, should allow you to manage the disks without any external help. Each disk must have a proper GPT label. If the disks have msdos labels, see Chapter 2, "Disk Preparation for Windows" on page 13.

If you run into problems, you may have to wipe out the master boot record (MBR), create a new msdos disk label, create a 100 MB system partition and a second boot/system partition (greater than or equal to 46 GB). Then format those targeted to Windows as new technology file system (NTFS). For more information, see "Disk Setup with Linux" on page 14.
 3. Attach a USB keyboard, USB mouse, and VGA monitor to the BaseIO blade (see Figure 1-1 on page 3 and Figure 1-2 on page 4).

Note: You can also use the display/keyboard and media redirection feature of the BaseIO BMC firmware (BMC firmware versions 1.2.2 or later). See "Enabling a Remote Console" in the *SGI Altix UV System Management Node Administrator's Guide* available at <http://docs.sgi.com>.




Figure 1-1 Altix UV 1000 Individual Rack Unit (IRU)




Figure 1-2 Altix UV 1000 BaseIO Riser Enabled Blade Front Panel Example




Figure 1-3 Altix UV 100 Individual Rack Unit (IRU)

4. Attach USB DVD-ROM drive if needed (see Figure 1-2 on page 4), and insert Windows Server 2008 R2 Datacenter DVD. For the location of the Altix UV 100 DVD drive, see Figure 1-3 on page 4.
5. Use the `ssh` command to connect to the Chassis Management Controller (CMC), similar to the following example:

Note: This assumes the system CMC SMN port is connected and an address is assigned. For more information, see Chapter 2 in the *SGI Altix UV CMC Controller Software User's Guide*.

```
% ssh root@hostname-cmc
```

```
SGI Chassis Manager Controller, Firmware Rev. 1.3.11
```

```
system-cmc CMC:r001i01c>
```

Typically, the default password set out of the factory is **root**. The CMC prompt appears. `r001i01c` refers to rack 1, IRU 1, CMC.

6. Use the `hwcfg` command on the console to disable NUMALink, as follows:

```
system-cmc CMC:r001i01c>hwcfg n15_enable=no
```

7. From the CMC, perform a power reset, as follows:

```
system-cmc CMC:r001i01c>power reset
```

8. When prompted on the console, hit the space bar for the **Boot Menu**, as shown in Figure 1-4 on page 6.




Figure 1-4 Console Prompt for Boot Menu Screen

From the boot manager screen, select **Device Manager**, as shown in Figure 1-5 on page 7.




Figure 1-5 Device Manager Screen

Navigate the **Device Manager**, as follows:

- F1=Scroll Help
 - Shift-6(^),v or arrow keys to Move Highlight
 - F9=Reset to Defaults
 - Enter=Select Entry
 - F10=Save
 - Esc= Exit without Save
 - Change setting values with plus (+) and minus (-) keys
 - Use CTRL-] q to exit the console.
9. From the **Device Manager** screen, as shown in Figure 1-6 on page 8, select **Advanced**.




Figure 1-6 Device Manager Advanced Option Screen

10. From the **Advanced** screen, select **SGI Custom Configuration**, as shown in Figure 1-7 on page 8.

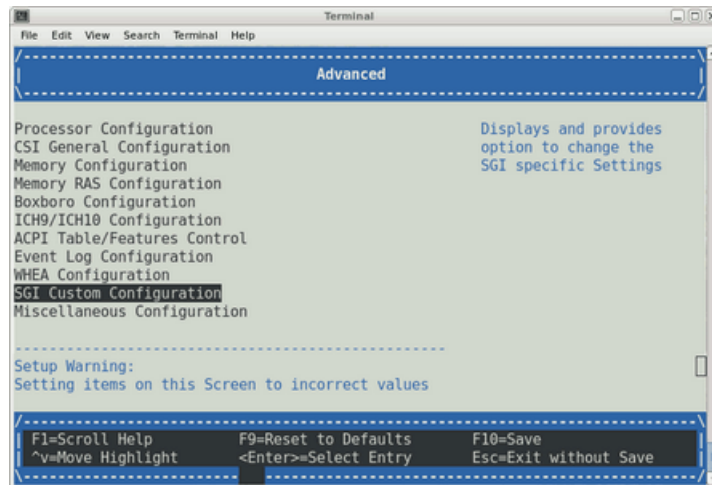


Figure 1-7 SGI Custom Configuration Option Example

From the Advanced screen, as shown in Figure 1-8 on page 9, perform the following:

- Change **Address Map Configuration Mvalue** from **37/7 (44bit)** to **36/4 (40bit)**. Use the plus and minus keys to change these values.
- Change **APIC Mode** from **Scalable** to **Legacy**.
- Use **F10** to save changes (**alt-0** in some terminals)

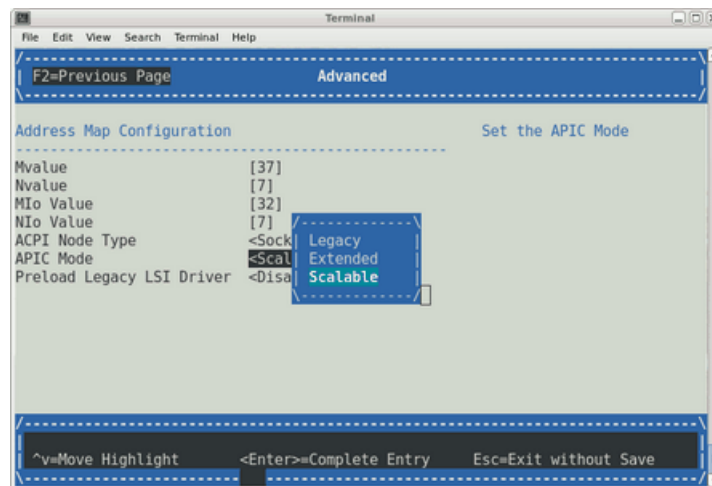


Figure 1-8 Advanced Screen APIC Mode Option

11. You can either boot to the EFI shell and run `\efi\boot\bootx64` or enter the boot menus and use the **Boot Manager** to select the boot loader. Be sure to navigate into the DVD's file hierarchy to find the Windows boot loader. Do not just boot from the DVD at the top level or you will get a legacy boot and Windows will not be able to use X2APIC mode or access more than eight Nehalem-EX sockets (or four Westmere-EX sockets.) Note that the DVD may not be at `fs0`, depending on what other bootable partitions are available on other drives.
12. Once the install completes and Windows has booted (using a VGA monitor and attached keyboard and mouse on the Altix UV system), copy the downloaded hotfixes onto the system. Unzip each one to the desktop.
- 13.

Note: The Windows Server 2008 R2 SP1 release includes the hotfixes by default so you can skip this step if you are install the SP1 release.

Install the first hotfix, 980598, by clicking on it and going through the dialog steps.

14. Reboot the system when prompted. The hotfix install completes after reboot. If the system appears to hang on reboot, after at least two minutes, reset the machine as follows:

```
system-cmc CMC:r001i01c>power reset
```

- 15.

Note: The Windows Server 2008 R2 SP1 release includes the hotfixes by default so you can skip this step if you are installing the SP1 release.

Install the first x2apic hotix - 2303458. Complete the dialogs and reboot when prompted. The hotfix install completes after reboot. Login as the administrator. Start a cmd shell and enter the following command:

```
bcdedit x2apicpolicy enable
```

16. Reboot the system. When prompted on the console, perform the following:
 - Hit the space bar for the **Boot Menu** (see Figure 1-4 on page 6).
 - Select **Device Manager** (see Figure 1-5 on page 7), **Advanced** (see Figure 1-6 on page 8), **SBI Custom Configuration** (see Figure 1-7 on page 8).
 - Change **Mvalue** back to 37 from 36. Change **Nvalue** back to 7 from 4. Change **APIC Mode** to **Extended** (see Figure 1-9 on page 11). Use F10 to save changes (alt-0 in some terminals). Use F2 to return to **Advanced** screen. Use F10 to save changes. Use ESC to return to **Device Manager** window. Select **Continue** and accept the reset request.




Figure 1-9 APIC Mode Extended Example Screen

17.

Note: The Windows Server 2008 R2 SP1 release includes the hotfixes by default so you can skip this step if you are installing the SP1 release.

Install the second x2apic hotfix - 2398906. Complete the dialog steps and reboot when prompted. The hotfix install completes after reboot. Shutdown windows to prepare to enable the rest of the system.

18. Use the `hwcfg` command on the console to enable NUMALink, as follows:

```
system-cmc CMC:r001i01c>hwcfg n15_enable=yes
```

19. From the CMC perform a power reset, as follows:

```
system-cmc CMC:r001i01c>power reset
```

20. Boot into Windows. This should take about 10 minutes on a 128 processor, one TB memory Altix UV 100 system. A 256 core, 2TB Altix UV 1000 system takes about 28 minutes to boot Windows. The disk activity LED on the Base I/O, just to the right of the SATA connector, should be flashing. This will be the only indication of activity for most of the wait time.

The `SAC>` command will show up on the serial console towards the end of the 10 minutes. It will take a few more minutes before the login screen appears on the VGA monitor. Windows will now take another 20 to 30 minutes, depending on

the speed of the root filesystem, to discover new hardware and write their configuration to disk. It will prompt for reboot when complete. Reboot once more to complete the install.

Disk Preparation for Windows

This chapter describes how to prepare disks before you install Windows Server 2008 R2 software on an SGI Altix UV 100 or SGI Altix UV 1000 system. Windows 2008 R2 Datacenter, when booted from EFI, will allow you to manage the disks with a proper GPT label without any external help. If you are having problems reinstalling Windows on a system originally installed with Windows, using the `dd` Linux command to wipe the master boot record (MBR) and GUID partition table (GPT) partition tables will allow you to be install Windows and bring the disks back online.

System disks are built into a RAID volume. You need to activate the volume (using the `lsiutil` configuration utility or the MPT Setup Utility, see Chapter 3, "Using the Drive Setup Utility" on page 17) before the system can see the system disks and boot the operating system. The `lsiutil` utility can be loaded from DVD or USB stick. The SGI Altix UV BaseIO BIOS firmware has an LSI driver embedded that allows for MPT Setup Utility access at EFI shell prompt using the `drvcfg -s` command (see Chapter 3, "Using the Drive Setup Utility" on page 17).

This section provides several alternative methods to setup disk drives to run Windows 2008 R2 Datacenter software.

Disk Setup with Windows

Procedure 2-1 Disk Setup with Windows to Prepare Disk Drives

1. If there is a a problem where the disk is not seen by the Windows Installation disk, then boot a minimal live Linux distribution that supports NTFS. See http://www.sysresccd.org/Main_Page or <http://mirrors.kernel.org/gentoo/releases/amd64/current-iso/> and use the `dd` command to clear the disk, as described in the next step.
2. The first 51 to 64 blocks on a disk need to be cleared to get rid of all the data that Windows install program sees. Use the `dd` command to clear the first several dozen sectors of the drive, as follows:

```
% dd if=/dev/zero of=/dev/sdX bs=512 count=32
```

3. Start the Windows install program and use the installer to online, partition, and format the disk. It is easier to wipe the drive and use Windows software to

partition the disk, making extra partitions for non-Windows operating systems you may want to install. Linux will install within a pre-made partition.

Disk Setup with Linux

In summary, you need to wipe out the master boot record (MBR), create a new `msdos` label, create 100MB system partition and a second boot/system partition greater than or equal to 46GB. Then format the partitions targeted to Windows as NTFS.

Procedure 2-2 Disk Setup with Linux to Prepare Disk Drives

If you have been running Linux on your Altix UV 100 or Altix UV 1000 system and want to prepare disk drives for Windows, perform the following steps:

1. Boot a minimal live Linux distribution that supports NTFS. See http://www.sysresccd.org/Main_Page or <http://mirrors.kernel.org/gentoo/releases/amd64/current-iso/>.
- 2.



Caution: Before you run the following commands, make sure that nothing on the drive is mounted or the running, Linux may get confused and crash. And make extra sure that you are wiping the correct drive. Physically pull any drives you do not want to overwrite, if you are not sure.

Assuming the drive is Serial Advanced Technology Attachment (SATA) or Serial attached SCSI (SAS), use the `dd` command, as follows:

```
% dd if=/dev/zero of=/dev/sdX bs=512 count=1
```

3. Use the `parted` disk partitioning and partition resizing command, as follows:

```
% parted /dev/sdX
```

4. Use the `mklabel` command to make an MS-DOS label, as follows:

```
% mklabel msdos
```

5. Use the `mkpart` command to create a partition, as follows:

```
% mkpart p fat32 0 100MB
```

6. Create another partition (whatever size you need > 46GB, as follows:

```
% mkpart p fat32 100MB 100GB
```

7. Run `mkpart` for any additional partitions.

8. The boot partition should be selected with the "boot" flag. Only one boot partition can be selected, as follows:

```
(parted) set 2 boot on
```

9. Use the `mkfs.ntfs` command to create an NTFS file system for each partition that Windows will see, as follows:

```
% mkfs.ntfs /dev/sdX
```

Replace the "X" with the letter of the drive on your system, for example,
`/dev/sda`

Using the Drive Setup Utility

The SGI Altix UV BaseIO BIOS firmware has an LSI driver embedded that allows for MPT Setup Utility access at EFI shell prompt using the `drvcfg -s` command. This section describes how to use it to activate a RAID Volume

Procedure 3-1 Using the MPT Setup Utility to Activate a RAID Volume

System disks are built into a RAID volume. You need to activate the volume using the before the system can see the system disks and boot the operating system, as follows:

1. Use the `drvcfg -s` command from EFI shell to activate the volume, as follows:

```
Shell> drvcfg -s

/-----\tions.
| LSI Logic MPT Setup Utility          3.04.00.00          |
| Adapter List                               |
| Adapter      PCI  PCI  PCI  FW Revision      Status      |
|              Bus  Dev  Fnc                    |
| SAS1064E     04   00   00   1.29.00.00-IR     Enabled      ^|
|                                                    *|
|                                                    v|
| Esc = Exit Menu          F1/Shift+1 = Help          |
| Enter = Adapter Properties                          |
\-----/
```

2. Press **Enter** and the **Adapter Properties** screen appears, as follows:

```
/-----\tions.  
| LSI Logic MPT Setup Utility          3.04.00.00  
| Adapter Properties -- SAS1064E  
|  
| Adapter                               SAS1064E  
| PCI Address(Bus/Dev/Func)            04:00:00  
| MPT Firmware Revision                1.29.00.00-IR  
| SAS Address                          50800691:5C624000  
| Status                               Enabled  
| Boot Support                         [Enabled BIOS & OS]  
|  
| RAID Properties  
|  
| SAS Topology  
|  
| Advanced Adapter Properties  
|  
| Esc = Exit Menu          F1/Shift+1 = Help  
| Enter = Select Item     -/+ = Change Item  
|-----\
```

- Use the down arrow key to highlight **RAID Properties** and press the Enter key.
- From the **Select New Array Type** screen, with **View Existing Array** selected, press Enter to view the existing array.

```

/-----\tions.
| LSI Logic MPT Setup Utility      3.04.00.00
| Select New Array Type -- SAS1064E
|
| View Existing Array View the existing configuration.
|
| Create IM Volume   Create Integrated Mirror Array of 2
|                   disks plus an optional hot spare. Data
|                   on the primary disk may be migrated.
|
| Create IME Volume  Create Integrated Mirrored Enhanced
|                   Array of 3 to 8 disks including an
|                   optional hot spare.
|                   ALL DATA on array disks will be DELETED!
|
| Create IS Volume   Create Integrated Striping array of
|                   2 to 8 disks.
|                   ALL DATA on array disks will be DELETED!
|
| Esc = Exit Menu      F1/Shift+1 = Help
| Enter = Choose array type to create
|                               Esc = Return to Adapter Properties
\-----/
```

5. Use the down arrow key to highlight **Activate Array** and press Enter. Note that the field is greyed out if the array is currently active.

```
/-----\tions.  
| LSI Logic MPT Setup Utility          3.04.00.00 |  
| Manage Array -- SAS1064E            |  
|                                     |  
| Identifier          LSILOGICLogical Volume 3000 |  
| Type                IM                |  
| Scan Order          0                  |  
| Size(MB)            139898            |  
| Status              Optimal            |  
|                                     |  
| Manage Hot Spare   |  
| Synchronize Array |  
| Activate Array     |  
| Delete Array       |  
|                                     |  
| Esc = Exit Menu    F1/Shift+1 = Help |  
| Enter = Select Item |  
\-----/
```

6. After activating the array it will be necessary to reset the system, as follows:

```
Shell> power reset
```

7. You can speed up the reset process to get to EFI shell by just booting the first blade from the CMC, as follows:

```
cmc> hwcfg NL5_ENABLE=no
```

After activating the volume, you can revert to the original configuration, as follows:

```
cmc> hwcfg NL5_ENABLE=yes
```

And then reset the system.

Index

A

- activating a RAID volume, 13, 17
- APIC Mode
 - Legacy, 9
 - Scalable, 9

B

- BIOS setup utility, 5
- boot menu, 5

C

- command
 - hwcfg, 5
 - power reset, 5
- commands
 - dd, 13
 - mkfs.ntfs, 15
 - mklablel, 14
 - mkpart, 15
 - parted, 14
- console
 - exiting, 7

D

- dd command, 13
- Device Manager
 - Advanced option, 7
 - See also "BIOS setup utility", 5
 - SGI Custom Configuration option, 8
- disk drive

- SAS, 14

- SATA, 14

- disk partition requirements, 14
- disk preparation, 13
- disk setup to prepare disk drives for Windows Server software, 14
- disk setup with Linux, 14
- disk setup with Windows, 13
- drvcfg -s, 17

E

- exit console, 7

H

- hwcfg command, 5

I

- installing Windows, 1
- introduction, 1

L

- lsiutil configuration utility, 13

M

- making system disks visible to install CD, 13, 17
- master boot record (MBR), 14
- mkfs.ntfs command, 15

mklable command, 14
mkpart command, 15
MPT Setup Utility, 13

N

navigate the device manager, 7

P

parted command, 14

power reset command, 5

U

using the drive setup utility, 17