

HP IO Accelerator Version 3.2.3 Windows User Guide

Abstract

This document describes software requirements for all relevant HP IO Accelerators using Microsoft Windows operating systems. This document is intended for system administrators who plan to install and use HP IO Accelerators with a Microsoft Windows operating system. It is helpful to have previous experience with HP IO Accelerators and a Microsoft Windows operating system. This user guide is intended for IO Accelerator software release 3.2.3 or later.



Part Number: 647099-003
February 2013
Edition: 3

© Copyright 2011, 2013 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Microsoft®, Windows®, Windows® XP, Windows Vista®, and Windows Server® are U.S. registered trademarks of Microsoft Corporation.

Contents

About this guide	6
Contents summary	6
Introduction	7
Overview	7
Product naming	7
Performance attributes	8
Required operating environment	9
Supported firmware revisions	9
Supported hardware	9
Windows Server environments	12
Installation overview	12
Uninstalling a previously-installed driver	12
Installing software on a Windows operating system	13
Using the Setup Wizard	14
Upgrading device firmware from VSL 1.x.x or 2.x.x to 3.x.x	19
Upgrading procedure	20
Manual installation on Windows Server	21
Manually installing on Windows Server 2003	22
Manually installing on Windows Server 2008 and 2012	23
Upgrading the device firmware using Windows	23
Viewing the firmware version	23
Performing the upgrade	23
Upgrading driver software using Windows	24
Upgrading driver software using Windows in a non-RAID configuration	24
Upgrading driver software using Windows in a RAID configuration	25
Defragmentation	26
Silent install option	26
Outdated firmware check	26
Enabling PCIe power	27
IO Accelerator naming	27
Adding a file system, formatting, and performing multi-disk configuration	28
Creating a RAID configuration	29
Using the IO Accelerator as swap with Windows operating system	29
Setting up SNMP for Windows operating systems	29
SNMP details for Windows operating systems	29
Using test mode registry values	30
SNMP MIB fields supporting Windows Server	31
Using installation logs	33
Windows Installer logging options	33
Creating an installation log	34
Creating an uninstall log	34
Creating a patch install log	34
Automated logging with the Windows Installer Logging Policy	35
Troubleshooting event log messages	35

Maintenance	37
Maintenance tools	37
Common maintenance tasks	37
Uninstalling the IO Accelerator driver package.....	37
Disabling AutoAttach	37
Enabling Autoattach	38
Unmanaged shutdown issues	38
Enabling PCIe power override	38
Enabling the override parameter	39
Utilities	41
Utilities reference.....	41
fio-attach.....	41
fio-beacon.....	42
fio-bugreport.....	42
fio-config.....	43
fio-config options	44
fio-detach	45
fio-format	46
fio-pci-check	47
fio-status.....	47
fio-sure-erase	50
fio-trim-config.....	51
fio-update-iodrive	52
Monitoring IO Accelerator health	54
NAND flash and component failure.....	54
Health metrics.....	54
Health monitoring techniques.....	54
Flashback substitution events.....	55
Software RAID and health monitoring	55
Trim support.....	57
Introduction to Trim.....	57
Trim platforms.....	57
Using Trim.....	57
Starting and stopping Trim.....	58
Enabling Trim.....	58
Controlling Trim aggressiveness.....	58
Trim configurations	59
Using Windows page files with the IO Accelerator.....	60
Introduction to Windows page files	60
Configuring IO Accelerator paging support.....	60
RAM consumption	60
Non-paged memory pool	61
Enabling and disabling paging support.....	61
Windows page file management.....	62
Setting up page files	62
System drive paging file configuration.....	63
Guaranteeing minimum committable memory	63
Verifying page file operation	64
Virtual Memory performance	64
Performance and tuning.....	65

Introduction to performance and tuning	65
Disabling DVFS	65
Limiting APCI C-states	65
Setting NUMA affinity	66
Setting the interrupt handler affinity	66
NUMA configuration.....	67
Introduction to NUMA architecture	67
FIO_AFFINITY parameter	67
Advanced configuration example	68
Checking the log for errors	69
Resources	70
Subscription service	70
For more information	70
Support and other resources	71
Before you contact HP	71
HP contact information	71
Customer Self Repair	71
Regulatory information	79
Safety and regulatory compliance	79
Turkey RoHS material content declaration	79
Ukraine RoHS material content declaration	79
Warranty information	79
Acronyms and abbreviations.....	80
Documentation feedback	82
Index.....	83

About this guide

Contents summary

- Instructions on downloading and installing the approved driver and utilities
- Instructions on maintaining the IO Accelerator
- Description of the following IO Accelerator models:
 - HP IO Accelerator for BladeSystem c-Class
 - HP PCIe IO Accelerator
 - HP PCIe IO Accelerator Duo



CAUTION: Before upgrading to 3.x.x software and firmware, back up all data on the IO Accelerator. The 3.2.3 software and firmware reformat the drive, which causes data to be lost if not backed up. The 3.2.3 software is not backward compatible with 1.2.x or 2.x software.

Introduction

Overview

Designed around ioMemory, a revolutionary storage architecture, HP IO Accelerator is an advanced NAND flash storage device. With performance comparable to DRAM and storage capacity on par with hard disks, the IO Accelerator increases performance so that every server can contain internal storage that exceeds the I/O performance of an enterprise SAN.

HP IO Accelerator is the first data accelerator designed specifically to improve the bandwidth for I/O-bound applications.

In addition to the hardware driver, the IO Accelerator also includes a VSL. This hybrid of the RAM virtualization subsystem and the disk I/O subsystem combines the best features of both systems. VSL functions as a disk to interface well with block-based applications and software, while also running like RAM underneath to maximize performance. This feature produces the following benefits:

- **Performance:** The VSL offers direct and parallel access to multiple CPU cores, enabling near linear performance scaling, consistent performance across different read/write workloads, and low latency with minimal interruptions and context switching.
- **Extensibility:** The VSL enables flash-optimized software development, making each IO Accelerator module a flexible building block for creating a flash-optimized data center.

Product naming

HP IO Accelerator Generation 1 devices include:

- AJ876A: HP 80GB IO Accelerator for BladeSystem c-Class
- AJ877A: HP 160GB IO Accelerator for BladeSystem c-Class
- AJ878A: HP 320GB IO Accelerator for BladeSystem c-Class
- AJ878B: HP 320 GB IO MLC Accelerator for BladeSystem c-Class
- BK836A: HP 640GB IO MLC Accelerator for BladeSystem c-Class



IMPORTANT: Generation 1 IO accelerators for BladeSystem c-Class are only compatible with G7 and earlier server blades.

- 600278-B21: HP 160GB Single Level Cell PCIe ioDrive for ProLiant Servers
- 600279-B21: HP 320GB Multi Level Cell PCIe ioDrive for ProLiant Servers
- 600281-B21: HP 320GB Single Level Cell PCIe ioDrive Duo for ProLiant Servers
- 600282-B21: HP 640GB Multi Level Cell PCIe ioDrive Duo for ProLiant Servers
- 641027-B21: HP 1.28TB Multi Level Cell PCIe ioDrive Duo for ProLiant Servers

HP IO Accelerator Generation 2 devices include:

- QK761A: HP 365GB IO MLC Accelerator for BladeSystem c-Class

- QK762A: HP 785GB IO MLC Accelerator for BladeSystem c-Class
- QK763A: HP 1.2 TB IO MLC Accelerator for BladeSystem c-Class



IMPORTANT: Generation 2 IO accelerators for BladeSystem c-Class are only compatible with Gen8 and later server blades.

- 673642-B21: HP 365GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673644-B21: HP 785GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673646-B21: HP 1205GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673648-B21: HP 2410GB Multi Level Cell G2 PCIe ioDrive2 Duo for ProLiant Servers
- 721458-B21: HP 3.0TB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers

Performance attributes

IO Accelerator capacity Models AJ878B and BK836A	320GB	640GB
NAND type	MLC (Multi Level Cell)	MLC (Multi Level Cell)
Read Bandwidth (64kB)	735 MB/s	750 MB/s
Write Bandwidth (64kB)	510 MB/s	550 MB/s
Read IOPS (512 Byte)	100,000	93,000
Write IOPS (512 Byte)	141,000	145,000
Mixed IOPS* (75/25 r/w)	67,000	74,000
Access Latency (512 Byte)	30 μ s	30 μ s
Bus Interface	PCI-Express x4	PCI-Express Gen1 x4

IO Accelerator capacity Models QK762A and QK763A	785GB	1.2TB
NAND type	MLC (Multi Level Cell)	MLC (Multi Level Cell)
Read Bandwidth (1MB)	1.5 GB/s	1.5 GB/s
Write Bandwidth (1MB)	1.1 GB/s	1.3 GB/s
Read IOPS (Seq. 512 Byte)	443,000	443,000
Write IOPS (Seq. 512 Byte)	530,000	530,000
Read IOPS (Rand. 512 Byte)	141,000	143,000
Write IOPS (Rand. 512 Byte)	475,000	475,000
Read Access Latency	68 μ s	68 μ s
Write Access Latency	15 μ s	15 μ s
Bus Interface	PCI-Express Gen2 x4	PCI-Express Gen2 x4

*Performance achieved using multiprocessor enterprise server

- Enterprise data integrity
- Field upgradeability
- Green footprint, 7.5W nominal per device

NOTE: MSI was disabled to obtain these statistics.

Required operating environment

The IO Accelerator is supported for use in the following operating environments:

- Windows Server 2003 (x86_64-bit only) SP2 (not supported on Gen8 servers)
- Windows Server 2008 (x86_64-bit only) R1 with SP2 or higher
- Windows Server 2008 (x86_64-bit only) R2
- Windows Server 2012 (x86_64-bit only)



IMPORTANT: All operating systems must be 64-bit architecture.

NOTE: IO Accelerators cannot be used as hibernation devices.

Supported firmware revisions

After February 19, 2013, all IO Accelerators ship with firmware version 7.1.13.109322 or higher. This firmware version only works with VSL 3.2.2 or higher. If you are installing a recently purchased or a replacement IO Accelerator into a system that already has IO Accelerators installed, then you must upgrade the firmware on the previously installed devices to 7.1.13.109322 or higher. You must also upgrade the VSL to 3.2.2 or higher. Upgrading the firmware and VSL on cards that were running firmware versions 6.x.x or higher and VSL 3.x.x or higher is NOT data destructive. However, HP recommends that you back up any data on the device before performing the upgrade. The latest supported version of the firmware and VSL can be found on the HP website (<http://www.hp.com>).

Release	Firmware revision
1.2.4	17350
1.2.7	36867 or 42014
1.2.8.4	43246
2.2.x	43674
2.2.3	101583
2.3.1	101971_4 or 101971_6
3.1.1	107004 or greater
3.2.3	109322

Supported hardware

HP IO Accelerator for BladeSystem c-Class

BladeSystem c-Class IO Accelerators have two distinct designs for the respective server product lines. The G1 through G7 IO Accelerator adapter is provided in a c-Class Type 1 Mezzanine card form factor. It can be installed in both Type 1 and Type 2 mezzanine slots within the c-Class blade G1 through G7 servers,

enabling a total of two cards in a half-height server blade, and three cards in a full-height server blade and up to 6 in a double-high, double-wide server (BL680c).

The Gen8 adapter is provided in a c-Class Type B Mezzanine card form factor. It can only be installed in Type B mezzanine slots within the Gen 8 or later servers, enabling one IO Accelerator in a half-height Gen8 server.

The Type I mezz card and the Type B mezz card are distinguished by the mezzanine connector. The type B card is slightly larger than a Type I mezz card.

The amount of free RAM required by the driver depends on the size of the blocks used when writing to the drive. The smaller the blocks, the more RAM is required. The table below lists the guidelines for each 80GB of storage. For the latest information, see the *QuickSpecs* sheet to the *HP IO Accelerator for HP BladeSystem c-Class* at HP Customer Support (<http://www.hp.com/support>).

The Remote Power Cut Module for the c-Class blade mezzanine card provides a higher level of protection in the event of a catastrophic power loss (for example, a user accidentally pulls the wrong server blade out of the slot). The Remote Power Cut Module ensures in-flight writes are completed to NAND flash in these catastrophic scenarios. Without the Remote Power Cut Module, write performance is slower. Writes are not acknowledged until the data is written to the NAND module, thereby slowing performance. When the Remote Power Cut Module is installed, writes are acknowledged by the IO Accelerator controller to the driver. The IO Accelerator controller then completes the write to the NAND module.

The IO Accelerators (QK761A, QK762A, and QK763A) for Gen 8 BladeSystem c-Class have the power cut functionality embedded on the card. They offer the protection without requiring the remote power cut module.

NOTE: The Remote Power Cut Module is used only in the AJ878B and BK836A models. Without the Remote Power Cut Module, write performance is slower.

HP PCIe IO Accelerator minimum requirements

- An open PCI-Express slot—The accelerator requires a minimum of one half-length, half-height slot with a x4 physical connector. All four lanes must be connected for full performance. HP PCIe IO Accelerator Duo requires a minimum of a full-height, half-length slot with an x8 physical connection. If your system is using PCI 1.1, all x8 signaling lanes must be connected for full performance. If your system is using PCI 2.0, for full performance you only have to connect x4 signaling lanes.

NOTE: For PCIe IO Accelerators, using PCIe slots greater than x4 does not improve performance.

NOTE: The power cut feature is built into PCIe IO Accelerators; therefore, no Remote Power Cut Module is necessary.

- 300 LFM of airflow at no greater than 50°C. To protect against thermal damage, the IO Accelerator also monitors the junction temperature of its controller. The temperature represents the internal temperature of the controller, and it is reported in `fio-status` report. The IO Accelerator begins throttling write performance when the junction temperature reaches 78°C. If the junction temperature continues to rise, the IO Accelerator shuts down when the temperature reaches 85°C.

NOTE: If you experience write performance throttling due to high temperatures, see your computer documentation for details on increasing airflow, including fan speeds.

- Sufficient RAM to operate—The amount of RAM that the driver requires to manage the NAND flash varies according to the block size you select when formatting the device (filesystem format, not low-level format). For a virtual machine using an IO Accelerator directly (using PCI pass-through), consult the user guide for the installed operating system. The following table lists the amount of RAM required per

100GB of storage space, using various block sizes. The amount of RAM used in driver version 3.0 is significantly less than the amount used in version 1.2.x.

Average block size (bytes)	RAM usage for each 80 GB IO Accelerator (Megabytes)	RAM usage for each 100 GB IO Accelerator (Megabytes)	Minimum system RAM requirement for 320 GB MezzIO Accelerator*
8,192	250	280	1 GB
4,096	400	530	1.6 GB
2,048	750	1,030	3 GB
1,024	1,450	2,000	5.8 GB
512	2,850	3,970	11.4 GB

Average block size (bytes)	Minimum system RAM requirement for 640 GB MezzIO Accelerator*	Minimum system RAM requirement for 785 GB MezzIO Accelerator*	Minimum system RAM requirement for 1.2 TB MezzIO Accelerator*
8,192	2 GB	2.2 GB	3.4 GB
4,096	3.2 GB	4.2 GB	6.4 GB
2,048	6 GB	8.1 GB	12.4 GB
1,024	11.6 GB	15.7 GB	24 GB
512	22.8 GB	31.2 GB	47.6 GB

* For IO Accelerator use only. Additional RAM is needed for system operating system and applications.

HP PCIe IO Accelerator Duo requirements

In addition to the IO Accelerator cooling and RAM requirements listed in the previous table, the IO Accelerator Duo requires at least:

- A PCIe Gen1 x8 slot or a PCIe Gen2 x4 slot
- A minimum of a full-height, half-length slot with a x8 physical connection. For systems with PCI 1.1, all eight signaling lanes must be active for full IO Accelerator Duo performance. For systems with PCIe 2.0, only four signaling lanes must be active for full performance.

NOTE: With driver version 3.1.0 and later, the driver detects in the BIOS if the PCIe slot supports a 75W power draw. If the slot supports up to 75W, the IO Accelerator device draws up to that amount of power. However, if an external power cable is used, power is only supplied by that cable.

To verify whether a slot is supplying 75W, view the system logs or use the `fio-pci-check` utility.

Windows Server environments

Installation overview

-
- △ **CAUTION:** The firmware for all IO Accelerator devices in a system running HP IO Accelerator VSL 3.1.x or later must be upgraded to the latest version of the firmware.
-
- △ **CAUTION:** If you have IO Accelerator devices configured for HP IO Accelerator VSL 2.x or earlier, you must upgrade the firmware before installing new devices in the system. For upgrade procedures, see "Upgrading device firmware from VSL 1.x.x or 2.x.x to 3.x.x (on page 19)."
1. If you are installing this version of HP IO Accelerator VSL on a system with IO Accelerator devices configured for VSL 2.x or earlier, follow the instructions in "Upgrading device firmware from VSL 1.x.x or 2.x.x to 3.x.x (on page 19)."
 2. Install ("Installing software on a Windows operating system" on page 13) the latest version of the IO Accelerator driver.
 3. If necessary, upgrade the firmware ("Upgrading the device firmware using Windows" on page 23) to the latest version.
 4. Configure the devices by adding a file system ("Adding a file system, formatting, and performing multi-disk configuration" on page 28), creating a RAID configuration (on page 29), and so on.

Uninstalling a previously-installed driver

-
- △ **CAUTION:** Upgrading an IO Accelerator from Version 1.2.x to Version 2.x is a one-way process. New firmware is required that is not compatible with 1.2.x systems. Be sure all data is backed up before continuing with the upgrade.
-

Before installing a new driver, uninstall the previous version.

1. Go to the **Start** menu.
2. Open the **Control Panel**.
3. Select **Add or Remove Programs**.
4. Select the package or program.

NOTE: Uninstall IO Manager and HP StorageWorks IO Accelerator software.

5. Select **Remove**.
6. Reboot the server, and then log in as a user with Administrative privileges.

If previous (downrev) drivers (such as driver files `fiodrive.sys`) exist on the system (typically located in `C:\Windows\System32\drivers`) after you have uninstalled the software, then upon reboot, Windows

recognizes those drivers for the IO Accelerator. This situation causes issues with the installation of Version 3.x.

After the software is uninstalled and the system is rebooted, Windows operating systems display a `New Hardware Found` message and normally do not find a driver for the new device.

If a driver is found, go to **Control Panel > Device Manager** and search for **HP IO Accelerator**. View the driver details and check the revision of the driver. If the HP IO Accelerator appears in the Device Manager and the driver revision is not 3.x or later, manually uninstall the driver for each IO Accelerator in the system, and then reboot the system.

If no driver is available, the HP IO Accelerator appears in the Device Manager as a mass storage controller. The system must be in this state before you install the 3.x software.

Installing software on a Windows operating system

1. Complete the hardware installation.
2. Log in to your computer as Administrator or have Administrator rights.
3. Download the Windows® IO Accelerator installation program from the HP website (<http://www.hp.com/support>) to your desktop or a convenient directory.
4. Run the IO Accelerator installation program (Setup Wizard) as explained in "Using the Setup Wizard."

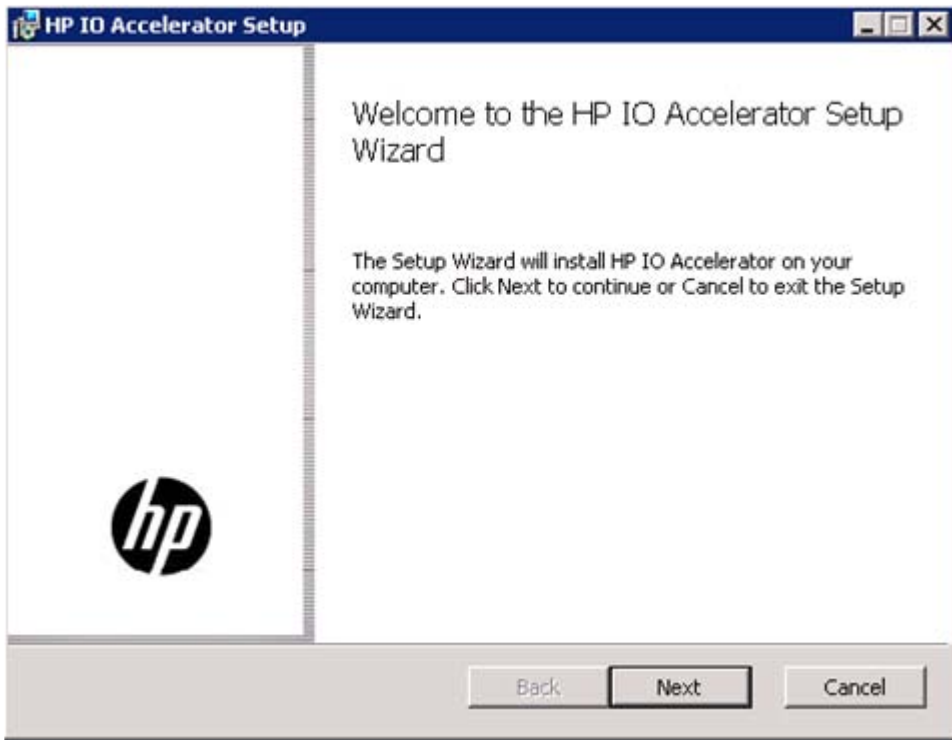


CAUTION: The installation program does not remove previous versions of the software. If a previous version is removed after the newest version is installed, the driver does not load after a restart. If this occurs, run the **Repair** option in the installation program from **Programs and Features** (or **Add or Remove Programs**) in the **Control Panel**, and then restart the computer.


NOTE: Screens might appear slightly different (for example, the word "StorageWorks" might not appear in the name or pathname with version 2.x).

Using the Setup Wizard


1. Download the IO Accelerator installation program. The **Welcome** screen appears.



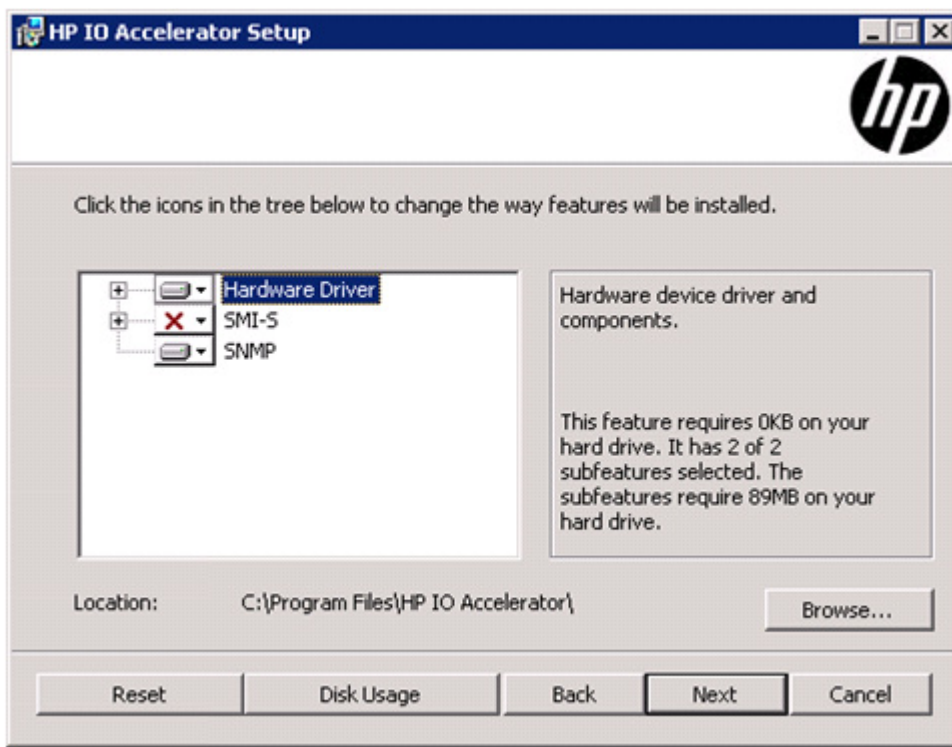
2. Click **Next**. The **Installation Recommendations** screen appears.



3. If you have a previous version of the IO Accelerator software installed, follow the directions on the screen to uninstall it. Otherwise, click **Next**. The **Select Installation Folder** screen appears.



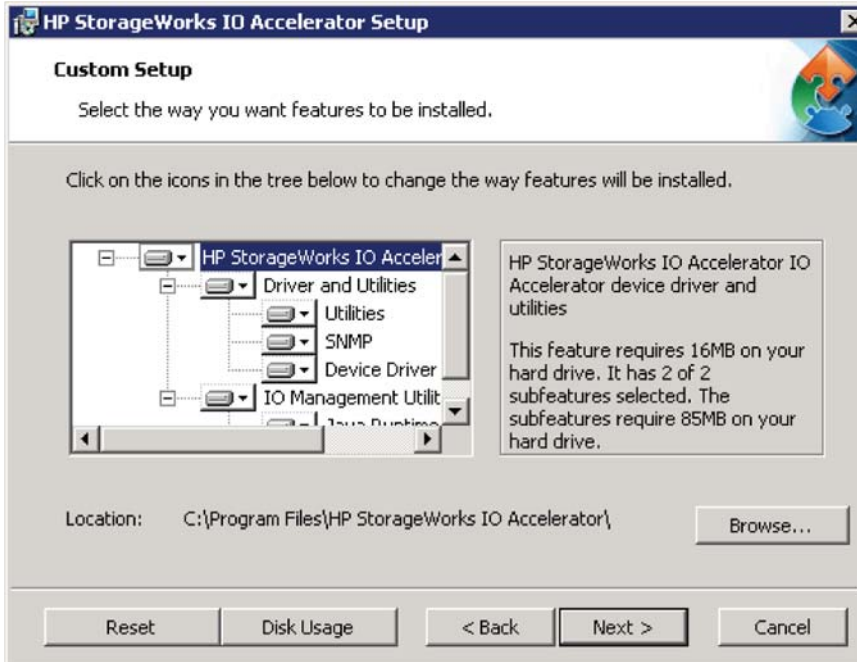
4. If you want to use the default installation location (C:\Program Files\HP IO Accelerator), click **Next**. If you want to change the default location, click **Browse**, navigate to the folder you want, and then click **Next**. The **Custom Setup** screen appears.



5. If you want to use the default program features for the installation, click **Next**. If you want to customize the features, click the plus sign next to **Driver and Utilities** or **IO Management Utilities**, and then select the customizations you want from the list.



CAUTION: Changing the default features might affect the performance or functionality of the IO Accelerator software.




The following features are available:

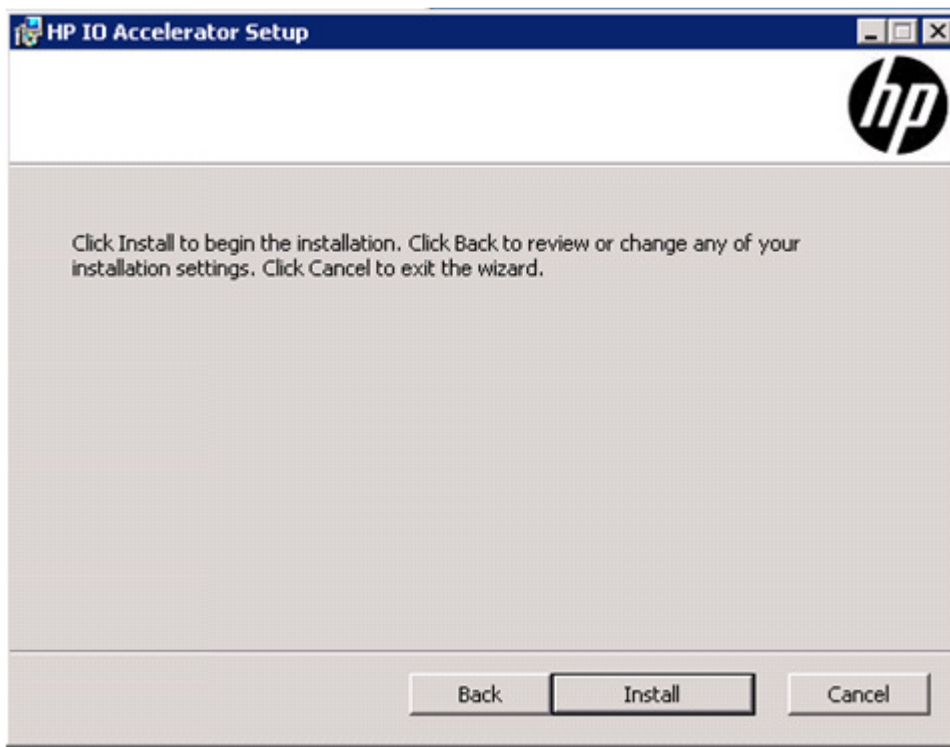
- Device Driver (required for the drive to be recognized in Windows® operating system)
- HP IO Accelerator Management Tool (Java® Runtime Software)
- SNMP (requires Microsoft® SNMP to be installed)
- Utilities (Command Line Utilities)

HP recommends that you select **will be installed on local hard drive** for each feature. If you change your mind later, you can use the **Repair** option in **Programs and Features**, or **Add or Remove Programs** in the Control Panel.

The **Configure Shortcuts** screen appears.




6. If you want to accept the default program shortcuts, click **Next**. Otherwise, clear any options you do not want, and then click **Next**. The **Ready to Install** screen appears.




7. Click **Install** to begin installation. An installation status screen appears. When installation is complete, the **Completed** screen appears.

NOTE: You might be prompted to stop installation and start the Microsoft® SNMP service during the software installation.



8. Click **Finish** to exit the Wizard and launch the IO Accelerator software.



The installation program performs the following:

- Creates a folder for the IO Accelerator software components (the default path is C:\Program Files\HP IO Accelerator).

- Installs and loads the IO Accelerator Windows® driver. (When complete, you will need to reboot the system.)
- Installs support for SNMP (if installed, and if Microsoft® SNMP is installed and the SNMP service is running).

When the installation program creates the IO Accelerator folder on the drive, it also creates these sub-folders:

- **<driver version>\Driver**—Installs manually using Device Manager.
- **Firmware**—Contains the latest IO Accelerator firmware.
- **SMI-S**—Currently not used; reserved for future SMI-S use.
- **SNMP**—Contains the SNMP components.
- **Utils**—Contains the IO Accelerator command line utilities.

To continue installation, go to "Upgrading the device firmware using Windows (on page 23)."

Upgrading device firmware from VSL 1.x.x or 2.x.x to 3.x.x

CAUTION: You cannot downgrade an HP IO Accelerator device firmware to an earlier version after you have upgraded the device.

CAUTION: Upgrading IO Accelerator devices that were previously configured for VSL 1.x.x or 2.x.x to work with VSL 3.x.x requires a low-level media format of the device. No user data is maintained during the media format process. Be sure to backup all data on your IO Accelerator device as instructed before upgrading the firmware.

Version 3.1.x of the HP IO Accelerator VSL supports new features, including the latest generation of IO Accelerator architecture and improved Flashback protection. These features require the latest version of the firmware. Every IO Accelerator device in a system running 3.1.x or later must be upgraded to the latest version of the firmware.

For example, if you have a system running 2.3.1 HP IO Accelerator VSL with IO Accelerator devices previously installed, and you want to install new IO Accelerator Gen2 devices (that require the latest version of the firmware), then you will need to upgrade all of the existing devices to the latest firmware version.

Upgrade path

Depending on the current version of your HP IO Accelerator device, to preserve the internal structure of the device, you might have to perform multiple upgrades. The following path is the minimum upgrade path that you must follow. Upgrade the HP IO Accelerator VSL software on the system, and upgrade the firmware to the compatible version in the following order:

1.2.4 > 1.2.7 > 2.1.0 > 2.2.3 > 3.1.x

For VSL upgrade information for the HP IO Accelerator, see the *HP IO Accelerator Release Notes* on the HP website (<http://www8.hp.com/us/en/support-drivers.html>). General upgrade instructions, including the firmware update instructions, are available in the *HP IO Accelerator User Guide* for each operating system.

Overformatting not supported

The `-o overformat` option is not supported in the 3.x.x VSL software. All upgraded HP IO Accelerator devices are formatted to the maximum advertised capacity, regardless of whether the device was overformatted prior to the upgrade.

Upgrading procedure

Be sure to follow the upgrade path and make sure that all previously installed IO Accelerator devices are updated with the appropriate 2.3.1-compatible firmware.

1. Prepare each existing IO Accelerator device for upgrade.
 - a. Backup user data on each IO Accelerator device.

CAUTION: Upgrading IO Accelerator devices that were previously configured for VSL 1.x.x or 2.x.x to work with VSL 3.x.x requires a low-level media format of the device. No user data is maintained during the media format process. Be sure to backup all data on your IO Accelerator device as instructed before upgrading the firmware.

Do not back up the data onto another IO Accelerator device on the same system. The backup must be to a local disk or to an externally attached volume.

- b. Run the `fio-bugreport` command-line utility and save the output. The output captures the device information for each device in the system. This device information will be useful in troubleshooting any upgrade issues. For example:

```
fio-bugreport
```
 - c. Detach the ioDrive devices. For example:


```
fio-detach /dev/fct*
```


For more information, see `fio-detach` (on page 45).
2. Unload the current IO Accelerator driver. For example:

```
$ modprobe -r iomemory-vsl
```
 3. Uninstall the 2.x HP IO Accelerator VSL software.
 - a. Go to **Start > Control Panel**.
 - b. Click **Administrative Tools**.
 - c. Click **Computer Management**.
 - d. Click **Device Manager** in the console tree at the left.
 - e. Expand the IO Accelerator Devices item.
 - f. Right-click the desired IO Accelerator device.
 - g. Click **Uninstall** or **Delete the driver software for this device** (depending on your version of Windows®).
 - h. Restart the computer.
 4. Install the new VSL.
 - a. Download the IO Accelerator installation program for Microsoft Windows from the HP website (<http://www.hp.com/support>) to your desktop or a convenient directory.
 - b. Run the IO Accelerator installation program. The installation program presents a custom setup tree-view with options for installation.

Mouse over a component in the tree view to see the description.
 - c. Click **Next**.

- d. To select a different folder for the installation, browse to the folder and click **OK**. The default folder is C:\Program Files\HP IO Accelerator.
The uninstaller file is placed in the root of the HP IO Accelerator (default install folder).
 - e. Follow the onscreen prompts to complete the install.
 - f. Choose **Reboot Now** on the finish screen of the installer.
For full installation instructions, see "Installing software on a Windows system ("[Installing software on a Windows operating system](#)" on page 13)."
5. Update the firmware on each device to the latest version using `fio-update-iodrive`.

 **CAUTION:** Do not turn off the power during a firmware upgrade, because this might cause device failure. If a UPS is not in place, consider adding one to the system before performing a firmware upgrade.

For all warnings, alerts, and options pertaining to this command utility, see "`fio-update-iodrive`."


Sample syntax:

```
fio-update-iodrive <iodrive_version.fff>
```

Where `<iodrive_version.fff>` is the path to the firmware archive. This command updates all of the devices to the selected firmware. To update specific devices, consult the utility reference for more options.

6. Reboot the system
If you run the `fio-status` command, a warning appears, indicating that the upgraded devices are missing a `lebmap`. This is a customary warning, and the issue will be corrected in the next step.
7. Format each device using `fio-format`. For example:

```
fio-format <device>
```

 **CAUTION:** Use this utility with care since it deletes all user information on the IO Accelerator.

You are prompted to confirm that you want to erase all data on the device. The format might take an extended period of time, depending on the wear on the device.

8. Attach all IO Accelerator devices. For example:

```
fio-attach /dev/fct*
```
9. Verify the status of all devices using `fio-status`. For example:

```
fio-status -a
```

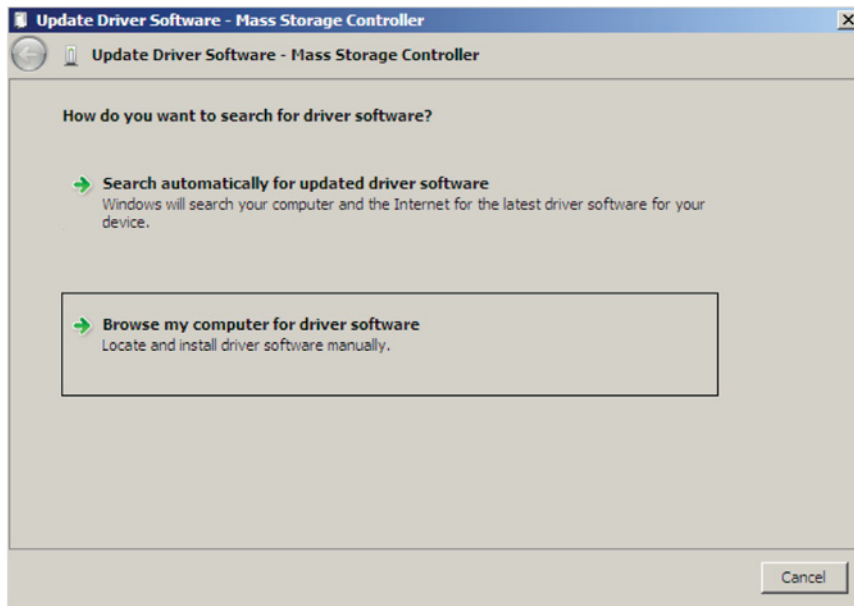
Your IO Accelerator devices are now upgraded for this version of the HP IO Accelerator. You can now install any IO Accelerator Gen2 devices.

Manual installation on Windows Server

The Windows® Setup program installs your driver and software. However, if the driver does not install, you can perform a manual installation (or upgrade) using the following steps.

Manually installing on Windows Server 2003


The Windows® Driver Wizard automatically detects the new IO Accelerator and begins to locate its driver after you restart the system. First, Windows® prompts you to locate the software driver.



1. If you have not done so, download and run the IO Accelerator Windows® Setup program from the HP website (<http://www.hp.com/support>).
2. Return to the **Update Driver** dialog box.
3. Select **Install from a list or specific location**. Click **Next**.
4. Select **Don't Search. I will choose the driver to install**. Click **Next**.
5. To open a browsing dialog, click **Have Disk**.
6. Browse to the folder that contains the IO Accelerator driver. The default is C:\Program Files\HP IO Accelerator\Driver.
7. Click **OK**. The correct driver installs the device software. When the driver installation completes, the following message appears:
The wizard has finished installing the software for: ioMemory VSL Device
8. To continue installation, see "Upgrading the device firmware using Windows (on page 23)."

Manually installing on Windows Server 2008 and 2012

The Windows Driver Wizard automatically detects the new IO Accelerator and begins to locate its driver after you restart the system. First, the OS prompts you to locate the software driver.



1. Click **Browse** next to the **path** field. A file dialog box appears.
2. Select the folder that contains the IO Accelerator driver. The default is C:\Program Files\HP StorageWorks IO Accelerator\Drivers.
3. Click **OK**.
4. Click **Next**. The correct driver installs the device software.
5. When the driver installation completes, restart the computer.
6. To continue installation, see "Upgrading the device firmware using Windows (on page 23)."

Upgrading the device firmware using Windows

CAUTION: Upgrade the device firmware only if the System Event Log reports out-of-date firmware, or if you are instructed to do so by HP Customer Support or the Release Notes document.

Viewing the firmware version

The firmware version can be found in the Windows® System Event Log. The version number is reported by the IO Accelerator Management Tool and the `fio-status` command line utility. For more information, see the *HP IO Accelerator Management Tool User Guide* or "`fio-status` (on page 47)."

Performing the upgrade



CAUTION: Do not turn off the power during a firmware upgrade, because this might cause device failure. If a UPS is not in place, consider adding one to the system before performing a firmware upgrade.

NOTE: You might be required to upgrade the IO Accelerator driver if you upgrade the device firmware. For more information, see the Release Notes.

Upgrading the firmware might take several minutes. The update command displays a progress bar to indicate the pace of the upgrade. This progress bar might appear to stall at points, but let the firmware update command complete.

Typically, you use the HP IO Accelerator Management Tool console to perform firmware upgrades. You can also perform the upgrade using the `fio-update-iodrive` command line utility.

To upgrade the firmware on the IO Accelerator using `fio-update-iodrive`:

1. Back up the data on the IO Accelerator prior to any upgrade as a precaution.
2. Open a command prompt and navigate to the `C:\Program Files\HP IO Accelerator\Utils` directory.
3. Run the `Detach` command for your device at the command prompt.

For example:

```
fio-detach /dev/fct0
```

detaches the IO Accelerator labeled `fct0`.

a. Specify the name of the drive you want to detach. Each IO Accelerator is labeled `/dev/fct#` where `#` is its bus number.

b. Use the `fio-status` command at a command prompt to view this number.

The `fio-detach` command shows a progress bar and percentage as it completes the operation.

4. Type the following command at the command prompt:

```
fio-update-iodrive <path> ioaccelerator_<version>.fff -d /dev/fct#
```

where `<path>` is the location of the IO Accelerator firmware (the default is `C:\Program Files\HP IO Accelerator\Firmware`) and `<version>` is the numerical version of the driver (higher is newer.) The `-d /dev/fct#` option identifies the IO Accelerator to update, where `#` is its bus number. Use the `fio-status` command at a command prompt to view this number.

All three external LED indicators light up during the update process.


5. Restart the computer to finish the firmware upgrade.

Upgrading driver software using Windows

Upgrading the driver software in Windows® operating systems depends on whether you are using a non-RAID or a RAID configuration.

Upgrading driver software using Windows in a non-RAID configuration


1. See the Release Notes for your version of the driver for details on any additional steps to perform the upgrade.

 **CAUTION:** Be sure to read the Release Notes document that comes with each new release as well as these installation instructions to ensure no loss of data when performing upgrades.

2. Follow the steps in "Uninstalling the IO Accelerator driver" to uninstall the existing driver.
3. Download the latest driver from the HP website (<http://www.hp.com/support>).
4. Either unzip or run the Windows® package to copy the files to a convenient directory.
5. Go to **Start > Control Panel**.
6. Click **Administrative Tools**.
7. Click **Computer Management**.
8. In the console tree at the left, click **Device Manager**.
9. Expand **HP IO Accelerators**.
10. Select **System Devices with pre-2.1 drivers**.
11. Right-click an IO Accelerator.
12. Click **Update Driver Software**. If necessary, see "Manual installation on Windows Server (on page 21)" for details on the remaining steps to install the updated driver. The Windows® operating system now detects your IO Accelerator.

Upgrading driver software using Windows in a RAID configuration

1. See the Release Notes for your version of the driver for details on any additional steps to perform the upgrade.

 **CAUTION:** Be sure to read the Release Notes document that comes with each new release as well as these installation instructions to ensure no loss of data when performing upgrades.

2. Shut down any applications that are accessing the IO Accelerator devices.
3. Open the IO Accelerator Utilities folder. (The default location for this release is C:\Program Files\HP IO Accelerator\Utils.)
4. Double-click the **AutoAttachDisable.reg** file to add a key to the Windows® registry. This key prevents the IO Accelerator device from automatically attaching the next time you restart the computer.
5. Uninstall the IO Accelerator software in Windows® **Add/Remove Programs**.
6. Restart the computer.
7. Download the latest driver from the HP website (<http://www.hp.com/support>).
8. Either unzip or run the Windows® package to copy the files to a convenient directory.
9. Open the IO Accelerator Utilities folder. (The default location is C:\Program Files\HP IO Accelerator\Utils.)
10. To reset the key in the Windows® registry, double-click the **AutoAttachEnable.reg** file. Your drive will automatically attach the next time you restart the computer.
11. Update the firmware of the devices. Follow the steps in "Upgrading the device firmware using Windows (on page 23)."

Restart the computer after the firmware upgrade is complete. The IO Accelerator Check Utility will run at next boot.

Your Windows® operating system now detects the devices in the RAID configuration with the upgraded software.

Defragmentation

The IO Accelerator does not need to be defragmented. However, some Windows® versions automatically run defragmentation as a scheduled task. If necessary, turn off automatic defragmentation.

Silent install option

Uninstalling the previous version

If you have a version of the HP IO Accelerator previously installed, you must uninstall it first. You must manually reboot the computer after installing the new version with the silent install option. This step must be performed prior to using any HP IO Accelerator utilities or functionality.

If you are installing remotely or with scripts, you can use the silent install option (`/quiet`) when you run the installation program in the command-line interface.

In the command-line interface, navigate to the folder that contains the `.exe` installer file, and then run the following command:

```
<installname>.exe /quiet
```

Where `<installname>.exe` is the name of the installer file.

This option installs the HP IO Accelerator using the default settings and eliminates the need to click Next or select settings during installation.

Installing components

The default command-line installation installs the HP IO Accelerator (including the command-line utilities) and the firmware file. You can add additional components using the following command:

```
<installname>.exe /quiet ADDLOCAL=SNMP,SDK,SMIS
```

Remove any component (and leading or trailing comma) that you do not want to install.

Performing a silent uninstall

You can silently uninstall the HP IO Accelerator with this command:

```
<installname>.exe /uninstall /quiet
```

Outdated firmware check

Checking by using the command-line interface

1. Run `fio-status` and examine the output:
 - o If any device is in minimal mode, then the firmware is outdated.
 - o If the firmware listed for any device is a lower number than the latest firmware version as noted in the release notes, then the firmware is old, but not outdated.
2. If the firmware is old or outdated, update it using the `fio-update-iodrive` utility.

Checking by using the optional GUI interface

You can use the (optional) HP IO Accelerator Management Tool GUI program to check for outdated firmware.

To check for outdated or old firmware:

1. Launch the IO Accelerator Management Tool application and look for any devices that have a warning symbol.
2. To ensure that the alert is from outdated firmware, click any devices with a warning symbol.
3. Select all devices requiring firmware update and use the IO Accelerator Management Tool to update the firmware. For details, see the *HP IO Accelerator Management Tool User Guide*.

To check for old but not outdated firmware:

1. Find the name of the latest firmware version as noted in the release notes.
2. Use the IO Accelerator Management Tool application to check the firmware version of each IO Accelerator device against the latest firmware version.

For instructions on how to update the firmware, see the *HP IO Accelerator Management Tool User Guide*.

Enabling PCIe power

For PCIe IO Accelerators, if you have installed any dual IO Accelerator devices, such as the HP ioDrive2 Duo, then the device might require additional power than the minimum 25 W provided by PCIe Gen2 slots to properly function.

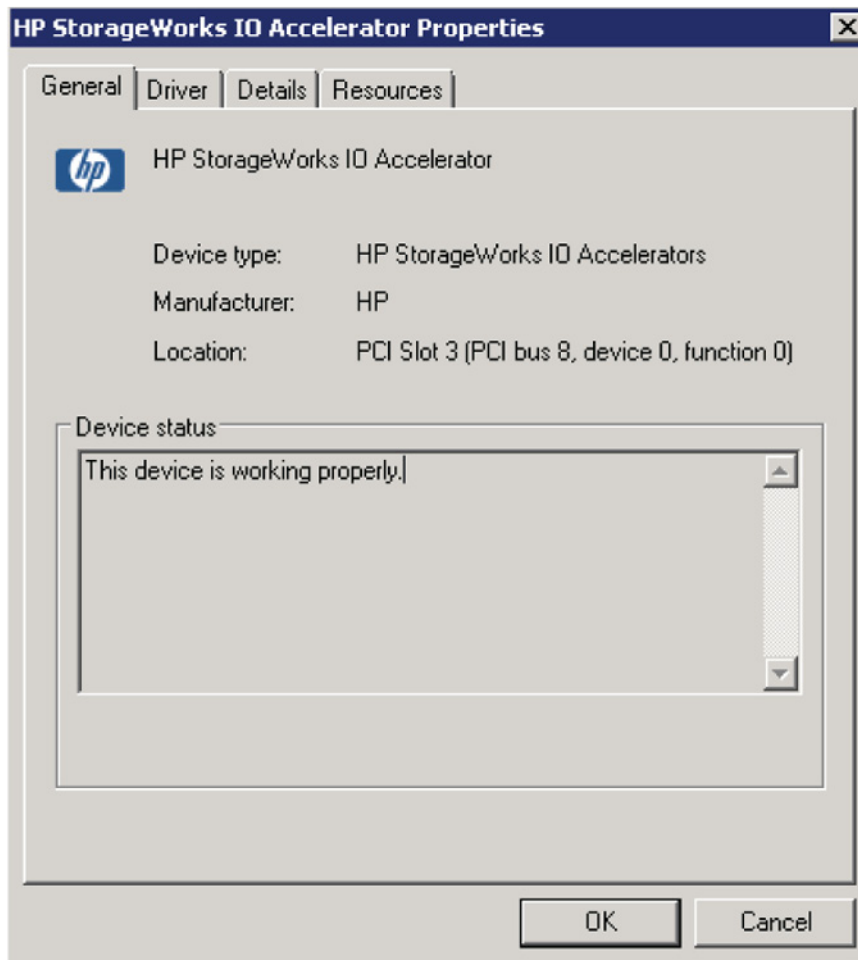
For instructions on enabling the device to draw additional power from the PCIe slots, see "Enabling PCIe power override (on page 38)."

IO Accelerator naming

The IO Accelerator receives an identification name and number as part of the installation process. The syntax is `fctx`, where `x` is the number of the bus where you installed the IO Accelerator. Use the IO Accelerator Management Tool to view this bus number, or follow these steps:

1. Select **Start > Control Panel > System > Hardware > Device Manager**.

2. Expand **HP StorageWorks IO Accelerators**, and then double-click **HP StorageWorks IO Accelerator**. The **Properties** dialog box appears. The **Location** field shows the bus number for your device (fct8 in the following example).



NOTE: HP assigns bus numbers, which begin at 0. These numbers might or might not reflect the physical location of the bus. For example, the second slot from the edge of the motherboard may be Bus 2, but it might also be Bus 16, or another arbitrary number. Checking Device Manager is one way to confirm the specific bus number for your installation. You can also use IO Accelerator Management Tool to view this number.

Adding a file system, formatting, and performing multi-disk configuration

If you want to use a standard file system such as NTFS, you can create and format partitions using standard disk management utilities.

For example, to initialize the IO Accelerator:

1. From the **Start** Menu, select **Control Panel > Administrative Tools > Computer Management**, and then from the **Storage** branch of the tree, select **Disk Management**.
2. Find the IO Accelerator in the list of storage devices, right-click it, and then select **Initialize Disk**.

3. After the device is initialized, you can add partitions or volumes. For more information, see the Disk Management documentation.

You can create dynamic volumes, which can be used in multi-disk RAID configurations (spanned, striped, or mirrored). For more information, see the Disk Management documentation.

Creating a RAID configuration

You can use your IO Accelerator as part of a RAID configuration. To do so, you must format the IO Accelerator as a dynamic volume. You can then use this dynamic volume to create multi-disk RAID configurations (spanned, striped, or mirrored).

For specific steps to perform a RAID configuration, see the Windows® Disk Management Utility documentation.

NOTE: If you are using RAID1 mirroring and one device fails, enter the `fio-format` command on the replacement device (not the existing good device) before rebuilding the RAID. HP recommends that you do not use a RAID5 configuration.

Using the IO Accelerator as swap with Windows operating system


To safely use the IO Accelerator as swap space, you must use the `fio-config` utility to pass a special pre-allocation parameter.

For example:

```
fio-config -p FIO_PREALLOCATE_MEMORY 1072,4997,6710,10345
```

where `1072,4997,6710,10345` are serial numbers obtained from `fio-status`. Be sure to use serial numbers for the IO Accelerator modules and not the adapter.

A 4K sector size format is required for swap. This reduces the driver memory footprint.

 **CAUTION:** You must have 400MB of free RAM per 80GB of IO Accelerator capacity (formatted to 4KB block size) to enable the IO Accelerator with pre-allocation enabled for swap. Attaching an IO Accelerator with pre-allocation enabled and insufficient RAM might result in the loss of user processes and system instability.

Setting up SNMP for Windows operating systems

Use the following procedures to set up SNMP for Windows® operating systems.

SNMP details for Windows operating systems

You can set up a test mode with your IO Accelerator Windows® SNMP agent. Doing so enables you to set test values in a Windows® registry and force SNMP traps without having to create the actual conditions on the device.

For example, you can use the SNMP test mode to change the non-writeable indicator and generate a trap, or simulate a change to the physical or logical size of the device, and so forth.

NOTE: To use SNMP test mode, you must have installed the SNMP option with your PCIe IO Accelerator.

Using test mode registry values

The following table shows the registry entries included for SNMP test values. Entries marked by an asterisk (*) generate SNMP traps when set to the indicated values, and the `cpqDrvMIBCondition` and `cpqDrvInfoStatus` MIB variables might be affected because of the changes.

NOTE: All entries, except those marked by **, reflect your registry changes immediately. Entries marked by ** require a restart of the Windows® SNMP agent for the changes to take effect.

SNMP test registry entry	Description
<code>cpqioDrvExtnAvailLogCapacityL</code>	Lower word of the available logical capacity in bytes
<code>cpqioDrvExtnAvailLogCapacityU</code>	Upper word of the available logical capacity in bytes
<code>cpqioDrvExtnBytesReadL</code>	Lower word of the total number of bytes read since the device was formatted
<code>cpqioDrvExtnBytesReadU</code>	Upper word of the total number of bytes read since the device was formatted
<code>cpqioDrvExtnLogBytesWrittenL</code>	Lower word of the number of user data bytes written
<code>cpqioDrvExtnLogBytesWrittenU</code>	Upper word of the number of user data bytes written
<code>cpqioDrvExtnPhysBytesWrittenL</code>	Lower word of the total physical bytes written
<code>cpqioDrvExtnPhysBytesWrittenU</code>	Upper word of the total physical bytes written
<code>cpqioDrvExtnTotalLogCapacityL</code>	Lower word of the total logical capacity in bytes as formatted
<code>cpqioDrvExtnTotalLogCapacityU</code>	Upper word of the total logical capacity in bytes as formatted
<code>cpqioDrvExtnTotalPhysCapacityL</code>	Lower word of the total logical capacity in bytes as formatted
<code>cpqioDrvExtnTotalPhysCapacityU</code>	Upper word of the total logical capacity in bytes as formatted
<code>cpqioDrvExtnUsablePhysCapacityL</code>	Lower word of the useable physical capacity in bytes. This space is holding valid data, or is erased and ready for writing, or is waiting to be reclaimed via garbage collection.
<code>cpqioDrvExtnUsablePhysCapacityU</code>	Upper word of the useable physical capacity in bytes. This space is holding valid data, or is erased and ready for writing, or is waiting to be reclaimed via garbage collection.
<code>cpqioDrvInfoInternalTemp*</code>	Current temperature of the device in degrees Celsius. If this value is set above 78°C for ioDimm cards, a trap is generated. If set above 90°C for HP Mezzanine cards, a trap is generated.
<code>cpqioDrvInfoFlashbackIndicator*</code>	1 = flashback redundancy is degraded; 2 = false
<code>cpqioDrvInfoNonWritableIndicator*</code>	1 = device is no longer writable because it has surpassed the read-only threshold; 2 = false
<code>cpqioDrvInfoPercentLifeRemaining</code>	Upper word of the total logical capacity in bytes as formatted
<code>cpqioDrvInfoState*</code> (trap generated if state = 4)	Current state of the attached client device: <ul style="list-style-type: none"> • unknown(0) • detached(1) • attached(2) • minimal(3) • error(4) • detaching(5) • attaching(6) • scanning(7)

SNMP test registry entry	Description
	<ul style="list-style-type: none"> • formatting(8) • updating(9) • attach(10) • detach(11) • format(12) • update(13)
cpqioDrvInfoWearoutIndicator*	Boolean: True = device has surpassed the wearout threshold.
cpqioDrvTestMode	Set test mode on or off
cpqioDrvTestModeIndex	Number indicating the selected IO Accelerator
mib	Name of the MIB in use
pathname	Path to the driver, set at installation
traps**	Set trap generation on or off
updateDelay**	Number of milliseconds to wait until getting the next value from the driver to generate a trap

SNMP MIB fields supporting Windows Server

- cpqloDimmMibRevMajor
- cpqloDimmMibRevMinor
- cpqloDimmMibCondition
- cpqloDimmInfoIndex
- cpqloDimmInfoStatus
- cpqloDimmInfoName
- cpqloDimmInfoSerialNumber
- cpqloDimmInfoPartNumber
- cpqloDimmInfoSubvendorPartNumber
- cpqloDimmInfoSparePartNumber
- cpqloDimmInfoAssemblyNumber
- cpqloDimmInfoFirmwareVersion
- cpqloDimmInfoDriverVersion
- cpqloDimmInfoUID
- cpqloDimmInfoState
- cpqloDimmInfoClientDeviceName
- cpqloDimmInfoBeacon
- cpqloDimmInfoPCIAddress
- cpqloDimmInfoPCIDeviceID
- cpqloDimmInfoPCISubdeviceID
- cpqloDimmInfoPCIVendorID
- cpqloDimmInfoPCISubvendorID

- cpqloDimmInfoPCISlot
- cpqloDimmInfoWearoutIndicator
- cpqloDimmInfoFlashbackIndicator
- cpqloDimmInfoWritableIndicator
- cpqloDimmInfoInternalTemp
- cpqloDimmInfoHealthPercentage
- cpqloDimmInfoMinimalModeReason
- cpqloDimmInfoReducedWriteReason
- cpqloDimmInfoMilliVolts
- cpqloDimmInfoMilliVoltsPeak
- cpqloDimmInfoMilliVoltsMin
- cpqloDimmInfoMilliWatts
- cpqloDimmInfoMilliWattsPeak
- cpqloDimmInfoMilliAmps
- cpqloDimmInfoMilliAmpsPeak
- cpqloDimmInfoAdapterType
- cpqloDimmInfoAdapterPort
- cpqloDimmInfoAdapterSerialNumber
- cpqloDimmInfoAdapterExtPowerPresent
- cpqloDimmInfoPowerlossProtectDisabled
- cpqloDimmInfoInternalTempHigh
- cpqloDimmInfoAmbientTemp
- cpqloDimmInfoPCIBandwidthCompatibility
- cpqloDimmInfoPCIPowerCompatibility
- cpqloDimmInfoActualGoverningLevel
- cpqloDimmInfoLifespanGoverningLevel
- cpqloDimmInfoPowerGoverningLevel
- cpqloDimmInfoThermalGoverningLevel
- cpqloDimmInfoLifespanGoverningEnabled
- cpqloDimmInfoLifespanGoverningTgtDate
- cpqloDimmExtnIndex
- cpqloDimmExtnTotalPhyslCapacityU
- cpqloDimmExtnTotalPhysCapacityL
- cpqloDimmExtnTotalLogCapacityU
- cpqloDimmExtnTotalLogCapacityL
- cpqloDimmExtnBytesReadU

- cpqloDimmExtnBytesReadL
- cpqloDimmExtnBytesWrittenU
- cpqloDimmExtnBytesWrittenL
- cpqloDimmExtnFormattedBlockSize
- cpqloDimmExtnCurrentRAMUsageU
- cpqloDimmExtnCurrentRAMUsageL
- cpqloDimmExtnPeakRAMUsageU
- cpqloDimmExtnPeakRAMUsageL
- cpqloDimmWearoutTrap
- cpqloDimmNonWritableTrap
- cpqloDimmFlashbackTrap
- cpqloDimmTempHighTrap
- cpqloDimmTempOkTrap
- cpqloDimmErrorTrap
- cpqloDimmPowerlossProtectTrap

Using installation logs

You can create a Windows® installation log to track various types of messages displayed when installing the driver.

Windows Installer logging options

The Windows® Installer handles its installations through **Msiexec.exe**. The logging command is issued with the `/L` parameter. Options are indicated with an asterisk. For example, the following command turns on logging with verbose output (`*v`):

```
msiexec /i "C:\MyPackage\Example.msi" /L*v "C:\log\example.log"
```

The following logging options enable you to create different types of logs, depending on the information you need about the installation.

Option	Description
*a	Startup of actions
*c	Initial UI parameters
*e	All error messages
*i	Status messages
*m	Out of memory or fatal exit information
*o	Out-of-disk-space messages
*p	Terminal properties
*r	Action-specific records
*u	User requests
*v	Verbose output
*w	Nonfatal warnings

Option	Description
*x	Extra debugging information
*+	Append to existing log file
*!	Flush each line to the log
**	Log all information except for v and x options

Creating an installation log

1. Determine the path of the MSI file (for example, `C:\MyPackage\Example.exe`).
2. Decide the path of the log (for example, `C:\log\example.log`).
3. Open `cmd.exe`. You can use any command shell.
4. Use the `msiexec` command line to launch the MSI with logging parameters.

Any installation logging command line must have this form:

```
msiexec /i <path_to_msi> /L*v <path_to_log>
```

The following is a sample command for creating the installation log. The `/i` parameter launches the MSI package. Be sure to specify the actual paths to your MSI file and log file.

```
msiexec /i "C:\MyPackage\Example.msi" /L*v "C:\log\example.log"
```

If you want the log to be created next to the MSI, you can specify only the name of the log file:

```
msiexec /i "C:\MyPackage\Example.msi" /L*v "example.log"
```

After the installation is finished, the log is complete.

Using an .exe bootstrapper

When the installation package is included in an .exe bootstrapper, do not use `msiexec` in the command line. For example:

```
"C:\MyPackage\Setup.msi" /L*v "example.log"
```

Creating an uninstall log

Launch the installer with the logging command, for example:

```
"C:\MyPackage\Setup.exe" /L*v "example.log"
```

This command puts the package into maintenance mode, and you can choose to uninstall it. Because the package was launched with logging, an uninstall log is generated.

Creating a patch install log

To create a log for a patch installation, follow the procedure in "Creating an installation log (on page 34)," but replace the `/i` parameter with `/p`, for example:

```
msiexec /p "C:\MyPackage\Patch.msp" /L*v "C:\log\patch.log"
```

Using the .exe bootstrapper

Another approach is to create a log file by using the `/L*v` parameters in the command line of the Installer Bootstrapper. Also, these parameters can be always passed to the MSI when the package is launched through the .exe bootstrapper.

If you want your installation package to create a log, follow these steps:

1. Open the Installer project of your installation package.
2. Go to the **Media** page and select the **Bootstrapper** tab.
3. Select the **Create EXE setup file** option.
4. Set the MSI Command Line field to `/L*v "C:\package.log"`.

When the user launches the installation through the bootstrapper, a log of the installation ("package.log") will be created automatically in the "C:\\" drive.

The command line received by the bootstrapper overrides the command line in the **MSI Command Line** field. Therefore, if you launch an EXE package with logging parameters, these parameters will be used for creating the log.

Automated logging with the Windows Installer Logging Policy

To enable automated logging for debugging purposes, create a registry value.

1. From a command line, run `regedit.exe`.
2. Locate the following registry key:
`HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\Windows\Installer`
3. Create a new STRING Entry named **Logging**.
4. Change its value using the logging options discussed previously. Note that you cannot use the wildcard (*) for this approach.

The .LOG file will be created in the Temp folder of the current user account with a name of this form:
MSI.LOG.

NOTE: Do not leave this option active because every install or uninstall operation on an .msi file will create a new log. This action results in unnecessarily occupied disk space and makes your system crowded. Only use this option for debugging purposes.

The logging policy is particularly useful for troubleshooting Active Directory and Group Policy deployments. In these cases, the installation is carried out without a user interface, and you cannot specify a command line for the .msi.

Troubleshooting event log messages

The Windows® System Event Log displays `fiordrive` messages concerning the IO Accelerator. These messages can be either informational and warnings.

NOTE: Each IO Accelerator is numbered from 0 upwards. These numbers reflect the bus number where you installed the device. Use the IO Accelerator Management Tool to view this number for your device.

To open the Windows Event Viewer:

1. Click **Start**.
2. Right-click **Computer**, and select **Manage** from the list.
3. Expand **System Tools**.
4. Expand **Event Viewer**.
5. Expand **Windows Logs**.
6. Select **System**.

The following table lists common Event Log error messages, along with suggested solutions.

Message	Suggested solution
<p>Error: ioDrive(x) firmware is too old. The firmware must be updated.</p>	<p>To update the firmware, see "Upgrading the device firmware using Windows (on page 23)."</p>
<p>Error: The ioDrive(x) initialization failed with error code 0xerrorcode (where errorcode is a number which might vary)</p>	<ol style="list-style-type: none"> 1 Reinstall the Windows driver. 2 Remove and reseat the IO Accelerator device. 3 Remove the IO Accelerator and insert it in a different slot.
<p>Error: The ioDrive(x) was not attached. Use the fio-attach utility to rebuild the drive.</p>	<p>This error might appear after an unmanaged shutdown. You can use either the fio-attach command line utility or HP IO Manager to re-attach the IO Accelerator. This attach process might take up to ten minutes as the utility performs a consistency check on the drives.</p>
<p>Warning: The ioDrive(x) was not attached because auto-attach is disabled.</p>	<p>The IO Accelerator must attach to the Windows operating system to be available to users and applications. This attach normally occurs at boot time. As part of this attach process, the IO Accelerator driver verifies the presence of an AutoAttach parameter in the Windows® registry. If you create this Registry parameter to disable auto-attach, the attach operation does not complete.</p> <p>To attach an unattached IO Accelerator:</p> <ol style="list-style-type: none"> 1 Run the HP IO Accelerator Management Tool. 2 Select your unattached IO Accelerator from the Device Tree. 3 Click Attach. 4 Confirm the Attach operation. <p>Your IO Accelerator now attaches to the Windows® operating system. To re-enable auto attach at boot time, see "Enabling auto attach ("Enabling Autoattach" on page 38)."</p>

Maintenance

Maintenance tools

The IO Accelerator includes both software utilities to maintain the device and external LED indicators to display the status. You can also install SNMP as a monitoring option.

The following are the most common tasks for maintaining your IO Accelerator. You can also use the IO Accelerator Management Tool application to perform firmware upgrades. For more information, see the *HP IO Accelerator Management Tool User Guide*.

Common maintenance tasks

The following tasks are the most common tasks for maintaining your IO Accelerator.

NOTE: All commands require administrator privileges. Log in as **root** or use `sudo` to run the commands.

Uninstalling the IO Accelerator driver package

1. Go to **Start > Control Panel**.
2. Click **Administrative Tools**.
3. Click **Computer Management**.
4. In the console tree at the left, click **Device Manager**.
5. Expand **HP IO Accelerators**.
6. Right-click an IO Accelerator.
7. Click **Uninstall**. The Windows® operating system uninstalls the driver.

Disabling AutoAttach

When the IO Accelerator driver is installed, it is configured to automatically attach any devices when the driver is loaded. When necessary, disable the AutoAttach feature.

To disable auto attach:

1. Open the **IO Accelerator Utilities** folder. The default location is `C:\Program Files\HP IO Accelerator\Utils\`.
2. Double-click the `autoattachdisable.reg` file.
3. If necessary, confirm that you want to modify the registry. This creates a new `DWORD` parameter registry key called `AutoAttach` in `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\fiodrive\Parameters`.
Your IO Accelerator does not automatically attach the next time you reboot the computer.

4. When you finish troubleshooting the driver issue, use the IO Accelerator Management Tool to attach the IO Accelerators and make them available to Windows® operating system.

Enabling Autoattach

1. Open the **IO Accelerator Utilities** folder. The default location is C:\Program Files\HP IO Accelerator\Utils\.
2. Double-click the `autoattachenable.reg` file.
3. If necessary, confirm that you want to modify the registry. This action resets the `AutoAttach` parameter in the Registry. The next time that you reboot your Windows® operating system, your IO Accelerator automatically attaches.

Unmanaged shutdown issues

Unmanaged shutdowns, due to power loss or other circumstances, force the IO Accelerator to perform a consistency check during the reboot. This reboot might take several minutes or more to complete and is indicated by a progress percentage during a startup.

You can cancel this consistency check by pressing **Esc** during the first 15 seconds after the **Fusion-io Consistency Check** message appears at the prompt. However, if you choose to cancel the check, the IO Accelerators remains unavailable to users until the check is completed. You can perform this check later using the IO Accelerator Management Tool **Attach** function.

Although data written to the IO Accelerator is not lost due to unmanaged shutdowns, important data structures might not have been properly committed to the drive. This consistency check repairs these data structures.

Enabling PCIe power override

For PCIe IO Accelerators, if you have installed any dual IO Accelerator devices, such as the HP ioDrive2 Duo, then the device might require additional power than the minimum 25 W provided by PCIe Gen2 slots to properly function. Even if additional power is not required for your device, all dual IO Accelerator devices that receive additional power might benefit with improved performance.

HP ioDrive2 Duo devices must have additional power to properly function. For more information on which devices require additional power, see the *HP PCIe IO Accelerator for ProLiant Servers Installation Guide*.

Additional power can be provided in two ways:

- **External power cable**—This cable ships with all dual ioMemory devices. For information on installing this cable, see the *HP PCIe IO Accelerator for ProLiant Servers Installation Guide*.

NOTE: When a power cable is used, all of the power is drawn from the cable and no power is drawn from the PCIe slot.

- **Enabling full slot power draw**—Some PCIe slots provide additional power (often up to 75 W of power). If your PCIe slot is rated to provide at least 55 W, then you can allow the device to draw full power from the PCIe slot by setting a VSL module parameter. For more information on enabling this override parameter, see "Enabling the override parameter (on page 39)."



CAUTION: If the PCIe slot is not capable of providing the needed amount of power, then enabling full power draw from the PCIe slot might result in malfunction or even damage to server hardware. The user is responsible for any damage to equipment due to improper use of the override parameter. HP expressly disclaims any liability for damage arising from improper use of the override parameter. To confirm the power limits and capabilities of each slot, as well as the entire system, contact the server manufacturer. For information about using the override parameter, contact HP Customer Support.

NOTE: The override parameter overrides the setting that prevents devices from drawing more than 25 W from the PCIe slot. The parameter is enabled by device using the device serial numbers. Once the setting is overridden, each device might draw up to the full 55 W needed for peak performance.

Before you enable the override parameter, ensure that each PCIe slot is rated to provide enough power for all slots, devices, and server accessories. To determine the power slot limits, consult the server documentation, BIOS interface, setup utility, or use the `fio-pci-check` command.

Important considerations

- If you are installing more than one dual IO Accelerator device and enabling the override parameter for each device, be sure the motherboard is rated to provide 55W power to each slot used. For example, some motherboards safely provide up to 75W to any one slot, but run into power constraints when multiple slots are used to provide that much power. Installing multiple devices in this situation might result in server hardware damage. Consult with the manufacturer to determine the total PCIe slot power available.
- The override parameter persists in the system and enables full power draw on an enabled device even if the device is removed and then placed in a different slot within the same system. If the device is placed in a slot that is not rated to provide 55W of power, your server hardware could experience a power drag.
- The override parameter is a setting for the IO Accelerator VSL software by server and is not stored in the device. When moved to a new server, the device defaults to the 25 W power limit until an external power cable is added or the override parameter is enabled for that device in the new server. To determine the total PCIe slot power available for the new server, consult the manufacturer.

Enabling the override parameter

1. Use one of the following methods to determine the serial number of each device to be installed in a compatible slot.

- Enter the `fio-status` command:

Sample output:

```
fio-status
```

```
...
```

```
Adapter: Dual Controller Adapter
```

```
Fusion-io ioDrive2 DUO 2.41TB, Product Number:F01-001-2T41-CS-0001, FIO
```

```
SN:1149D0969
```

```
External Power: NOT connected
```

```
PCIe Power limit threshold: 24.75W
```

```
Connected ioMemory modules:
```

fct2: SN:1149D0969-1121

fct3: SN:1149D0969-1111

In this example, 1149D0969 is the adapter serial number.

If you have multiple IO Accelerator devices installed on your system, use the `fio-beacon` command to verify where each device is physically located.

- Inspect the adapter serial number labels on the IO Accelerator devices to determine the serial numbers. However, HP recommends confirming that each serial number is an adapter serial number by running the `fio-status` command. The adapter serial number label resides on the back of all HP ioDrive Duo devices and HP ioDrive2 Duo devices. On ioDrive Duos, the serial number is located on the PCB component that is attached to the PCIe connector.
2. Set the module parameter for the `fio-config` command-line utility. For example:

```
fio-config -p FIO_EXTERNAL_POWER_OVERRIDE <value>
```

Where the `<value>` for this parameter is a comma-separated list of adapter serial numbers. For example:

```
1149D0969,1159E0972,24589
```

The `-p` option makes the parameter persistent after reboot.

```
fio-config /dev/fct<x> -p PCIe_GLOBAL_SLOT_POWER_LIMIT 74750
```

Where `<x>` is the device number for the device relative to the `<value>` specified adapter serial number. For example:

```
fio-config /dev/fct2 -p PCIe_GLOBAL_SLOT_POWER_LIMIT 74760
```

3. To enforce any parameter changes, you must reboot or unload and then load the drivers.

Utilities

Utilities reference

The IO Accelerator installation packages include various command line utilities, installed by default to the C:\Program Files\HP IO Accelerator\Utils folder. These utilities provide a number of useful manners to access, test, and manipulate your device.

Utility	Purpose
fio-attach	Makes an IO Accelerator available to the OS
fio-beacon	Lights the IO Accelerator external LEDs
fio-bugreport	Prepares a detailed report for use in troubleshooting issues
fio-config	Enables configuration parameters for device operation
fio-detach	Temporarily removes an IO Accelerator from OS access
fio-format	Performs a low-level format of an IO Accelerator
fio-pci-check	Searches for errors on the PCIe bus tree, specifically for PCIe IO Accelerators
fio-status	Displays information about the device
fio-sure-erase	Securely removes information from an IO Accelerator device
fio-trim-config	Enables or disables the Trim feature
fio-update-iodrive	Updates the IO Accelerator firmware

NOTE: All utilities have `-h` (Help) and `-v` (Version) options.

NOTE: To run these utilities from a command line using Windows Server®, you must either change directories to the directory that contains them (by default, C:\Program Files\HP IO Accelerator\Utils) or add that directory to your system path. See the documentation for your Windows® version for information about adding a directory to the system path.

NOTE: The command-line utilities require administrator rights to run under Windows® operating systems. Right-click the **Command Prompt** menu item and select **Run as administrator**.

fio-attach

Description

Attaches the IO Accelerator device and makes it available to Windows® operating systems. You can then partition the IO Accelerator device, or set it up as part of a RAID array, using the Windows® Disk Management utility. This command displays a progress bar and percentage as it completes the attach process.

In most cases, the driver automatically attaches the device when loading and does a scan. You only need to run `fio-attach` if you ran `fio-detach` or if you set the `auto_attach` parameter to 0.

If the IO Accelerator device is in minimal mode, then `auto-attach` is disabled until the cause of the device being in minimal mode is fixed.

Syntax

```
fiio-attach <device> [options]
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the PCIe bus number where you installed the IO Accelerator. For example, the name `/dev/fct4` refers to the IO Accelerator device installed in PCIe Bus 4 in your Windows® operating system. Use the IO Accelerator Management Tool or `fiio-status` to view this bus number.

You can specify multiple IO Accelerator devices. For example, `/dev/fct /dev/fct2` indicates the IO Accelerator devices installed in PCIe Buses 1 and 2 in your Windows® operating system.

Options

Option	Description
<code>-c</code>	Attach only if clean.
<code>-q</code>	Quiet: disables the display of the progress bar and percentage.

fiio-beacon

Description

The `fiio-beacon` utility turns on all three LEDs of the specified IO Accelerator for identification. You must detach the IO Accelerator device and then run `fiio-beacon`.

This utility is not applicable for the mezzanine format of the IO Accelerator, which has no externally visible LEDs.

Syntax

```
fiio-beacon <device> [options]
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the device number.

Options

Option	Description
<code>-0</code>	Off: Turns off the three LEDs
<code>-1</code>	On (default): Lights the three LEDs
<code>-p</code>	Prints the PCI bus ID of the device at <code><device></code> to standard output. Usage and error information might be written to standard output rather than to standard error.

fiio-bugreport

Description

Prepares a detailed report of the device for use in troubleshooting issues.

This utility captures the current state of the device. When a performance or stability problem occurs with the device, run the `fio-bugreport` utility and contact HP support (<http://www.hp.com/support>) for assistance in troubleshooting.

`fio-bugreport` runs several information-gathering utilities and combines the resulting data into a text file. The results are saved in the `utils` directory (default installation path is `C:\Program Files\HP IO Accelerator\Utils\`) in a `.cab` file that indicates the date and time the utility was run.

NOTE: If the utility recommends that you contact Fusion-io support, disregard that message and contact HP support (<http://www.hp.com/support>) instead.

Syntax

```
fio-bugreport
```

Sample output

```
C:\Users\>"\Program Files\HP IO Accelerator\Utils\fio-bugreport.exe"  
Generating bug report. Please wait, this may take a while...
```

```
-----
```

```
Gathering all Windows Event Logs...DONE  
Gathering Fusion-io Windows Event Logs...DONE  
Gathering System Information...DONE  
Running fio utilities...DONE  
Compressing to CAB file...DONE  
Bug report has successfully been created:  
fio-bugreport-20100222_192621.cab.  
Please e-mail this file to HP support
```

For example, a bug report file named `fiobugreport-20090921.192621.cab` indicates the following:

- Date (20090921)
- Time (192621, or 19:26:21)

fio-config

Description

Sets and gets configuration parameters for device operation. For a list of parameters, see "fio-config options (on page 44)."

Syntax

```
fio-config <device> [options]
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the PCIe bus number where you installed the IO Accelerator. For example, the name `/dev/fct4` refers to the IO Accelerator device installed in PCIe Bus 4 in your Windows® operating system. Use the IO Accelerator Management Tool or `fio-status` to view this bus number.

Options

Option	Description
-e	Enumerates configuration parameter names and values.
-g <name>	Gets the configuration parameter.
-p <name>	Set and persist the configuration parameter.
-s <name>	Set the configuration parameter in memory only.
-V	Print verbose information.
-v	Print version information.

fio-config options

The following table describes the driver parameters you can set with the `fio-config` utility.



IMPORTANT: `fio-config` options must be entered in uppercase to function properly.



IMPORTANT: MSI (Message Signaled Interrupts) is enabled by default and it cannot be disabled using `fio-config`.

Option	Default (min/max)	Description
AUTO_ATTACH	1 (0,1)	Always attach the device on driver load (1).
IODRIVE_TINTR_HW_WAIT	0 (0, 255)	Interval (microseconds) to wait between hardware interrupts
FIO_EXTERNAL_POWER_OVERRIDE	No devices selected	Allows selected devices to draw full power from the PCIe slot. Where the <value> for this parameter is a comma-separated list of adapter serial numbers.
FORCE_MINIMAL_MODE	0 (0,1)	Force minimal mode on the device (1), this parameter is set to false (0) by default.
PARALLEL_ATTACH	0 (0,1)	Enable parallel attach of multiple devices (1), this parameter is set to false (0) by default.
FIO_PREALLOCATE_MEMORY	0	For the selected device, preallocate all memory necessary to have the drive usable as swap space. For example: <pre>fio-config /dev/fct0 -p FIO_PREALLOCATE_MEMORY = "1234,54321"</pre> where "1234" and "54321" are serial numbers obtained from <code>fio-status</code> .
WIN_DISABLE_ALL_AFFINITY	0 (affinity is enabled)	When <code>WIN_DISABLE_ALL_AFFINITY</code> is set to 0, the driver will enable interrupt and worker thread affinity in the driver. When <code>WIN_DISABLE_ALL_AFFINITY</code> is set to 1, the driver will disable all affinity settings. This is an override of any other affinity settings. The driver must be reloaded for this parameter to take effect.
WIN_DISABLE_DEFAULT_NUMA_AFFINITY	0 (default; NUMA affinity is enabled)	When <code>WIN_DISABLE_DEFAULT_NUMA_AFFINITY</code> is set to 0, during initialization, the driver will query the Windows® operating system for the affinity settings assigned to the adapter by the OS. This setting is known as the "default NUMA affinity".

Option	Default (min/max)	Description
		After the affinity is queried correctly, the driver sets the affinity of the adapter interrupt and associated worker threads to the default operating system setting. This action has the effect of setting the affinity of the interrupt and worker threads to all processors on a single NUMA node in the system. When WIN_DISABLE_DEFAULT_NUMA_AFFINITY is set to 1, the driver will ignore the affinity settings assigned to the adapter by the OS. The driver must be reloaded for this parameter to take effect.
FIO_AFFINITY	N/A	FIO_AFFINITY is a list of <affinity specification> triplets to specify the affinity settings of all adapters in the system. Each item in the triplet is separated by a comma, and each triplet set is separated by a semicolon. For more information on using this parameter, see "NUMA configuration (on page 67)."

fio-detach

Description

Detaches and removes the corresponding `/dev/fctx` IO Accelerator block device. The `fio-detach` command waits until the device completes all read/write activity before executing the detach process. The command displays a progress bar and percentage as it completes the process.

NOTE: Before using this utility, be sure that the device you want to detach is not currently mounted and in use.

Syntax

```
fio-detach <device> [options]
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the PCIe bus number where you installed the IO Accelerator. For example, the name `/dev/fct4` refers to the IO Accelerator device installed in PCIe Bus 4 in your Windows® operating system. To view this bus number, use the IO Accelerator Management Tool or `fio-status`.

You can specify multiple IO Accelerator devices. For example, `/dev/fct1 /dev/fct2` indicates the devices installed in PCIe Buses 1 and 2 in your Windows® operating system.

Options

Option	Description
-f	Force: Causes an immediate detach (does not save metadata).
-i	Immediate: Causes a forced immediate detach (does not save metadata). This will fail if the device is in use by the OS.
-q	Quiet: Disables the display of the progress bar and percentage.



IMPORTANT: Although the `-f` (force) option causes the IO Accelerator device to detach, even in a RAID setup, it is strongly recommended to take the drives/volume offline using the Windows® Disk Management plug-in and then perform the detach. Forcing the detach may result in loss of data.

Notes

Attempting to detach an IO Accelerator device might fail with an error indicating that the device is busy. This might occur if the device is part of a software RAID (0, 1, 5) volume. The Windows® operating system refuses the request to detach the drive associated with the IO Accelerator device because it is part of a RAID volume and might cause the volume to fail. This does not occur with simple volumes (such as a single IO Accelerator device). To detach a simple volume, take the volume offline using the Disk Management MMC plug-in application.

fiio-format

Description

The `fiio-format` command performs a low-level format of the IO Accelerator. This format is different from a format performed by the operating system. The utility displays a progress bar and percentage as it completes the format.

The IO Accelerator ships pre-formatted, so you only have to run the `fiio-format` command if you must change the logical size of the device.



CAUTION: Use this utility with care since it deletes all user information on the IO Accelerator.

Syntax

```
fiio-format <device> [options] <device>
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the PCIe bus number where you installed the IO Accelerator. For example, the name `/dev/fct4` refers to the IO Accelerator device installed in PCIe Bus 4 in your Windows® operating system. To view this bus number, use the IO Accelerator Management Tool or `fiio-status`.

Options

Option	Description
<code>-b <size B K></code>	Set the block (sector) size, in bytes or KiBytes (base 2). The default is 512 bytes. For example: <code>-b 512B</code> or <code>-b 4K</code> (B in 512B is optional).
<code>-f</code>	Force the format size, bypassing normal checks and warnings. This option might be needed when <code>fiio-format</code> does not proceed properly. The "Are you sure?" prompt still appears unless you use the <code>-y</code> option. This option can only be used with the <code>-o</code> option.
<code>-q</code>	Quiet mode: Disable the display of the progress-percentage indicator.
<code>-s <size M G T %></code>	Set the device capacity as a specific size (in TB, GB, or MB) or as a percentage of the advertised capacity: <ul style="list-style-type: none"> • T Number of terabytes (TB) to format • G Number of gigabytes (GB) to format • M Number of megabytes (MB) to format • % Percentage, such as 70% (the percent sign must be included)
<code>-R</code>	Disable fast rescan on unclean shutdowns at the cost of some capacity.

Option	Description
-y	Auto-answer "yes" to all queries from the application (bypass prompts).

fio-pci-check

Description

Checks for errors on the PCI bus tree, specifically for IO Accelerators. This utility displays the current status of each IO Accelerator. It also prints the standard PCI Express error information and resets the state.

It is normal to see a few errors (perhaps as many as five) when `fio-pci-check` is initially run. Subsequent runs should reveal only one or two errors during several hours of operation.

The IO Accelerator driver must be loaded to run this utility. Some PCI errors cannot be reset in your Windows® operating system.

Syntax

```
fio-pci-check [options]
```

Options

Option	Description
-d <value>	1 = Disable the link; 0 = bring the link up
-f	Scan every device in the system.
-r	Force the link to retrain.
-v	Verbose: Print extra data about the hardware.
-y	"Yes" is forced when the user is asked to continue.

fio-status

Description

The `fio-status` command provides detailed information about the specified or all installed IO Accelerators, either `fctx` or `fiox` devices. The `fio-status` command must run with root privileges and requires that the driver be loaded. If it finds no device, it displays an index number instead.

`fio-status` provides alerts for certain error modes, such as a minimal-mode, read-only mode, and write-reduced mode, describing what is causing the condition.

Syntax

```
fio-status [<device>] [options]
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the PCIe bus number where you installed the IO Accelerator. For example, the name `/dev/fct4` refers to the IO Accelerator device installed in PCIe Bus 4 in your Windows® operating system. To view this bus number, use the IO Accelerator Management Tool or `fio-status`.

If you do not specify a name, `fio-status` displays information on all installed IO Accelerators. If the IO Accelerator is not loaded, this parameter is ignored.

Options

Option	Description
-a	Report all available information for each device.
-e	Show all errors and warnings for each device. This option is for diagnosing issues, and it hides other information such as format sizes.
-c	Count: Report only the number of ioMemory devices installed.
-d	Show basic information set plus the total amount of data read and written (lifetime data volumes). This option is not necessary when the -a option is used.
-fj	Format JSON: creates the output in JSON format.
-fx	Format XML: creates the output in XML format.
-u	Show unavailable fields. Only valid with -fj or -fx.
-U	Show unavailable fields and details why. Only valid with -fj or -fx.
-F<field>	Print the value for a single field (see the next option for field names). Requires that a device be specified. Multiple -F options may be specified.
-l	List the fields that can be individually accessed with -F.

Starting with version 3.0.0 and later, the standard formatting of the `fiio-status` output changed. This change affects any custom management tools that used the output of this utility.

Basic Information: If no options are used, `fiio-status` reports the following basic information:

- Number and type of cards installed in the system
- Software version

Adapter information:

- Adapter type
- Product number
- External Power Supply
- PCI power limit threshold (if available)
- Connected IO Accelerator devices

Block device information:

- Attach status
- Product name
- Product number
- Serial number
- PCIe slot number
- Firmware version
- Size of the device, out of total capacity
- Internal temperature (average and maximum, since IO Accelerator load) in degrees Centigrade
- Health status: healthy, nearing wearout, write-reduced or read-only
- Reserve capacity (percentage)
- Warning capacity threshold (percentage)

Data Volume Information: If the `-d` option is used, then the following data volume information is reported in addition to the basic information:

- Physical bytes written
- Physical bytes read

All Information: If the `-a` option is used, then all information is printed, which includes the following information in addition to basic and data volume information:

Adapter information:

- Manufacturer number
- Date of manufacture
- Power loss protection status
- PCIe bus voltage (avg, min, max)
- PCIe bus current (avg, max)
- PCIe power limit threshold (watts)
- PCIe slot available power (watts)
- PCIe negotiated link information (lanes and throughput)

Block device information:

- Part number
- Manufacturer's code
- Manufacturing date
- Vendor and sub-vendor information
- Size of the device, out of total capacity
- Format status and sector information (if device is attached)
- FPGA ID and Low-level format GUID
- PCIe slot available power
- PCIe negotiated link information
- Card temperature, in degrees Centigrade
- Internal voltage: avg. and max.
- Auxiliary voltage: avg. and max.
- Percentage of good blocks, data and metadata
- Lifetime data volume statistics

Error Mode Information: If the IO Accelerator is in minimal mode, read-only mode, or write-reduced mode when `fio-status` is run, the following differences occur in the output:

- `Attach status` is `Status unknown: Driver is in MINIMAL MODE:`
- The reason for the minimal mode state is displayed (such as `Firmware is out of date. Update firmware.`)
- `Geometry and capacity information not available.` appears.
- No media health information appears.

fio-sure-erase



CAUTION: Before using this utility, back up any important data.



IMPORTANT: As a best practice, do not use this utility if you have any IO Accelerator devices installed in the system that you do not want to clear or purge. First remove any devices that you do not want to accidentally erase. After the data is removed with this utility it is not recoverable.



IMPORTANT: If the device is in Read-only mode, perform a format using `fio-format` before running `fio-sure-erase`. If the device is in Minimal mode, then `fio-sure-erase` cannot erase the device. Updating the firmware may take the device out of Minimal Mode. If the device remains in Minimal mode, contact HP Customer Support for further assistance.



IMPORTANT: After using `fio-sure-erase`, format the device using `fio-format` before using the device again.

To run the `fio-sure-erase` utility, the block device must be detached. See `fio-detach` for more information.

Description

The `fio-sure-erase` is a command-line utility that securely removes data from IO Accelerator devices. It complies with the Clear and Purge level of destruction from the following standards:

- DOD 5220.22-M—Comply with instructions for Flash EPROM
- NIST SP800-88—Comply with instructions for Flash EPROM

Registry requirement

To create a registry key to configure the driver for ECC-bypass mode:

1. Locate the following key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\fiodrives\Parameters
```

2. Create a DWORD key underneath it called `BypassECC` and set the value to 1.
3. Restart the computer before running the utility.

Syntax

```
fio-sure-erase [options] <device>
```

where `<device>` is the name given to the IO Accelerator during the installation process. The syntax is `/dev/fctx`, where `x` indicates the PCIe bus number where you installed the IO Accelerator. For example, the name `/dev/fct4` refers to the IO Accelerator device installed in PCIe Bus 4 in your Windows® operating system. To view this bus number, use the IO Accelerator Management Tool or `fio-status`.

The `fio-sure-erase` syntax works on individual IO Accelerator devices. For example, if you are planning to purge an IO Accelerator Duo device, you must perform this operation on each of the IO Accelerator devices.

Options

If `fio-sure-erase` is run with no options, a Clear operation is performed.

Option	Description
-p	Purge instead of Clear: Performs a write followed by an erase. Purging the device might take several hours to complete, depending on the size of the device.
-y	No confirmation: Does not require a yes/no response to execute the utility.
-q	Quiet: Do not display the status bar.

Clear support

A Clear is the default state of running `fio-sure-erase` (with no options), and is the act of performing full low-level erase (every cell pushed to 1) of the entire NAND media, including retired erase blocks.

Metadata that is required for operation will not be destroyed (media event log, erase counts, physical bytes read/written, performance and thermal history), but any user-specific metadata will be destroyed.

A Clear operation performs the following:

1. Creates a unity map of every addressable block (this allows `fio-sure-erase` to address every block, including previously unmapped bad blocks).
2. Performs an erase cycle for each block (every cell is pushed to 1).
3. Restores the bad block map.
4. Formats the device to make the device usable again. The utility erases all of the headers during the clear.

Purge support

A Purge operation is implemented by using the `-p` option with `fio-sure-erase`. Purge is the act of first overwriting the entire NAND media (including retired erase blocks) with a single character (every cell written to logical 0), and then performing a full chip erase (every cell pushed to 1) across all media (including retired erase blocks).

Metadata that is required for operation will not be destroyed (media event log, erase counts, physical bytes read/written, performance and thermal history), but any user-specific metadata will be destroyed.

A Purge operation performs the following:

1. Creates a unity map of every addressable block (this allows `fio-sure-erase` to address every block, including previously unmapped bad blocks).
2. Performs a write cycle for each block (every cell written to 0).
3. Performs an erase cycle for each block (every cell pushed to 1).
4. Restores the bad block map.
5. Formats the drive to make the drive usable again. The utility erases all of the headers during the clear.

fio-trim-config

Description

Enables or disables the Trim feature, which reclaims available space from the file system. It is generally recommended that Trim be enabled for best performance. TRIM is enabled by default. Running this utility with no options displays the current Trim status. For more information, see "Trim support (on page 57)."

NOTE: This utility affects all IO Accelerators in the system. You cannot enable or disable Trim for only selected devices.

Syntax

```
fio-trim-config [options]
```

Options

Option	Description
-d	Disable Trim on this computer.
-e	(Default) Enable Trim on this computer.

fio-update-iodrive

Description

The `fio-update-iodrive` utility updates the HP IO Accelerator firmware. This utility scans the PCI-Express bus for all HP IO Accelerators and updates them all sequentially.

You must detach all IO Accelerator devices before updating the firmware. Detach the device from the system using the `fio-detach` utility. See "fio-detach (on page 45)" for more information.

To update one or more specific drives (rather than all drives):

- Make sure the IO Accelerator driver is loaded.
- Use the `-d` or `-s` option.

CAUTION: Do not turn off the power during a firmware upgrade, because this might cause device failure. If a UPS is not in place, consider adding one to the system before performing a firmware upgrade.

CAUTION: Note that when running multiple upgrades in sequence (such as going from driver version 1.2.7 to 2.1.0 to 2.3.1), it is important to load the IO Accelerator driver after each firmware upgrade step. Otherwise the on-drive format will not be changed, and there will be data loss.

CAUTION: Do not use this utility to downgrade the IO Accelerator to an earlier version of the firmware. Doing so might result in data loss. If you need to downgrade your firmware, contact HP customer support. For more information, see "Technical support."

NOTE: If you receive an error message when updating the firmware instructing you to update the midprom information, contact Customer Support. For more information, see "Technical support."

The default action (without using the `-d` or `-s` option) is to upgrade all IO Accelerators with the firmware contained in the `<iodrive_version.fff>` file. Confirm that all devices need the upgrade prior to running the update. If in doubt, use the `-p` (Pretend) option to view the possible results of the update.

Syntax

```
fio-update-iodrive <iodrive_version.fff> [options]
```

where `<iodrive_version.fff>` is the path and firmware archive file provided by HP. This is typically `C:\Program Files\HP IO Accelerator\Firmware`. This parameter is required.

Option	Description
-d	Updates the specified devices (by <code>fctx</code> , where <code>x</code> is the number of the device shown in <code>fio-status</code>). If this option is not specified, all devices are updated.

Option	Description
-f	Force upgrade. Used primarily to downgrade to an earlier firmware version. If the IO Accelerator driver is not loaded, this option also requires the -d option.
-l	Lists firmware available in the archive.
-c	Clears locks placed on a device.
-p	Pretend. Shows what updates are done. The firmware is not modified.
-q	Runs the update process without displaying the progress bar or percentage.
-y	Confirm all warning messages.
-s	Updates devices in selected slots, using '*' as a wildcard for boards. The slots are in the following format: [[[<domain>]: [<func>]]<="" [<slot>]="" [.=""]=""]<bus>:="" td=""> </domain>]:>

All three external LED indicators light up during the update process.



CAUTION: Use the -f option with care because it might damage your board.



CAUTION: Use the -d and -s options with care because updating a device in the wrong slot might damage it.

Monitoring IO Accelerator health

NAND flash and component failure

The IO Accelerator is a highly fault-tolerant storage subsystem that provides many levels of protection against component failure and the loss nature of solid state storage. However, as in all storage subsystems, component failures might occur.

When a large enough number of data blocks is retired due to error, the NAND flash media is considered worn out. By properly selecting NAND flash media for the hosted application and proactively monitoring device age and health, you can assure reliable performance over the intended product life.

Health metrics

The IO Accelerator driver manages block retirement using pre-determined retirement thresholds. The IO Accelerator Management Tool and the `fio-status` utility show a health indicator that starts at 100 and counts down to 0. As certain thresholds are crossed, various actions are taken.

At the 10% healthy threshold, a one-time warning is issued. For methods for capturing this alarm event, see "Health monitoring techniques (on page 54)."

At 0%, the device is considered unhealthy. It enters write-reduced mode, which somewhat prolongs its lifespan so data can be safely migrated. In this state, the IO Accelerator behaves normally except for the reduced write performance.

At some point after the 0% threshold, the device enters read-only mode. Any attempt to write to the IO Accelerator causes an error. Some file systems might require special mount options to mount a read-only block device, beyond specifying that the mount should be read-only. For example, under Linux, `ext3` requires that `-o ro,noLoad` is used. The `noLoad` option tells the file system not to try to replay the journal.

Read-only mode should be considered a final opportunity to migrate data off the device since device failure is more likely with continued use.

The IO Accelerator might enter failure mode. In this case, the device is offline and inaccessible. This can be caused by an internal catastrophic failure, improper firmware upgrade procedures, or device wears out.

Health monitoring techniques

fio-status

Output from the `fio-status` utility shows the health percentage and drive state. These items are in bold in the following sample output.

```
Found 1 ioDrive in this system
Fusion-io driver version: 2.1.0 build 19032
Adapter: ioDrive
HP 160GB SLC PCIe ioDrive for ProLiant Servers, Product Number:600278-B21
```

...

Media status: Healthy; Reserves: 100.00%, warn at 10.00%; Data: 99.12%

Lifetime data volumes:

...Physical bytes written: 6,423,563,326,064

...Physical bytes read: 5,509,006,756,312

The following Health Status messages are produced by the `fio-status` utility:

- Healthy
- Low metadata
- Read-only
- Reduced-write
- Unknown

HP IO Accelerator Management Tool

In the **Device Report** tab, look for the **Reserve Space** percentage in the right column. The higher the percentage, the healthier the drive is likely to be.

SNMP

For details on how to configure SNMP health indicators, see "Setting up SNMP for Windows operating systems (on page 29)."

Flashback substitution events

The IO Accelerator is equipped with Flashback substitution capability. This capability allows for predictive NAND flash failures to occur and keep the IO Accelerator functioning by substituting out bad locations or even a whole NAND chip. The output of the `fio-status -a` command will show if any flashback substitution events have occurred.

The following is an example of the output after entering the `fio-status -a` command. The numbers represent the bank (0-3) and chip number that was substituted.

First flashback event -> No action to take as infant mortality condition on NAND chip was caught per design.

Second flashback event -> Run `fio-bugreport`. Return the card (if under warranty) when you see two or more flashback substitution events.

- Flashback active on (bank a/chip m) and (bank b/chip n) (multi-bank failure)
- Flashback active on (bank a/chip m) and (bank b/chip m) (chip failure)

On the second flashback event, include the output from running the `fio-bugreport` command when you return the card.

Software RAID and health monitoring

Software RAID stacks are typically designed to detect and mitigate the failure modes of traditional storage media. The IO Accelerator attempts to fail as gracefully as possible, and its new failure mechanisms are compatible with existing software RAID stacks. A drive in write-reduced mode participating in a write-heavy workload is evicted from a RAID group for failure to receive data at a sufficient rate. A drive in read-only

mode is evicted when write I/Os are returned from the device as failed. Catastrophic failures are detected and handled just as though they were on traditional storage devices.

Trim support

Introduction to Trim

Trim is used to address a unique property of solid state drives. The issue occurs when a user deletes a file, and then the drive that contains the file does not recognize that it can reclaim that space. Instead, the drive reads the data as valid until the system informs the drive it can overwrite that data. With a typical hard drive you can continually write to the same sector of the drive without significant degradation. However, a solid state drive writes to different flash memory areas when the same logical sector is written to.

Because of the way flash memory handles writes, having a full SSD can cause significant reduction in write throughput. Currently available storage stacks and file systems were not designed with this caveat in mind. The SSD software must assume that all sectors contain valid data, and therefore treat the drive as full. In reality, a file system does not normally contain a full drive of data.

Trim is an enhancement to existing file systems that informs the SSD software of logical sectors that do not contain valid user data. Trim retrieves this information from a file system when a file has been deleted, and it informs the drive that it can reclaim the space held by that file. This feature enables the wear-leveling software to reclaim that space as reserve to handle future write operations.

For Trim to perform this task, it continually runs in the background and monitors the file systems being used on IO Accelerator devices. The process is meant to run in the background so as not to noticeably impact the performance of the computer.

Trim platforms

The Trim feature is available on IO Accelerators for Windows® operating systems starting with Windows® XP and Windows Server® 2003.

NOTE: Starting with Windows® 7 and Windows Server® 2008 R2, Trim capabilities are built into the operating system. IO Accelerators support Windows® 7 Trim, which means that IO Accelerator Trim is unnecessary under Windows® 7 and Windows Server® 2008 R2, and it will not run when either of these is present. No system changes are needed. The IO Accelerator Trim service automatically shuts off if it is attempting to run on a Windows® 7 or Windows Server® 2008 R2 or later operating system.

Using Trim

Trim runs quietly in the background as a Windows® Service a few minutes after the computer is started, so as not to slow down the boot process. However, you can also fine-tune Trim by using the configurable settings.

Starting and stopping Trim

To start or stop the Trim service through the Windows® Services Manager:

1. Search for **IO Accelerator Trim Service**.
2. Click **Start service** or **Stop service**.

Stopping the service only stops it for that session. Trim begins running again when the computer is restarted.

Enabling Trim

1. Open a registry editor.
2. Navigate to
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\fiodrive\Parameters`
3. Add the `TrimEnabled` value as a DWORD (see **TrimEnabled** below). Stop and restart the Trim service, or restart the computer.

TrimEnabled

This registry value controls whether Trim runs, regardless of the user starting and stopping Trim. If this value is present and set to 1, then Trim runs normally (unless you have stopped the service). If this value is set to 0, then Trim does not run, regardless of whether you have started the service. If this value is not present, Trim runs unless you stop it. The `TrimEnabled` value persists across sessions, so if you restart the computer and set this value to 0, Trim does not run when the computer is restarted.

Controlling Trim aggressiveness

You can control how aggressive the Trim service is. The default settings for Trim minimize impact on computer performance while enabling the IO Accelerators to perform optimally. However, if a drive is being used with a write-intensive application, then you can set Trim to be more aggressive. To do this, you must add registry entries.

1. Open a registry editor.
2. Navigate to
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\fiodrive\Parameters`
3. Add the `TrimTimeDelayInSeconds` value as a DWORD (see the **TrimTimeDelayInSeconds** description below).
4. Add the `BitmapRetrievalDelayInMilliseconds` value as a DWORD (see the **BitmapRetrievalDelayInMilliseconds** description below).
5. Stop and restart the Trim service, or restart the computer.

TrimTimeDelayInSeconds

This registry value controls how quickly Trim starts running again after it has finished processing all the volumes on a system. For example, if this value is set to 600 (10 minutes in seconds), then Trim will process all the volumes on a system, sleep for 10 minutes, and then process the volumes again and continue this cycle. This value can be any number between 0 and 4,294,967,295. If this value is not set or present in the registry, then Trim assumes a default of 900 seconds (or 15 minutes).

BitmapRetrievalDelayInMilliseconds

This value affects how much impact Trim has on the CPU. The default is 25 milliseconds but can be any number between 0 (most aggressive) and 4,294,967,295 (not aggressive at all).

Trim configurations

Trim can be used with the following configurations and features:

- RAID volumes (mirrored, spanned, or striped).

NOTE: RAID 5 (striped with parity) is not currently supported.

- Simple volumes (no RAID).
- Any combination of the above RAID levels across multiple devices, so long as there is at least one IO Accelerator in the RAID set.
- Multiple partitions on the same IO Accelerator device.
- NTFS and FAT32 file systems.
- Volumes with mount points (no `<drive letter>` is defined).
- Compressed volumes.
- Different cluster sizes, packet sizes, and sector sizes.
- Extended and shrunk volumes.

Using Windows page files with the IO Accelerator

Introduction to Windows page files

This section describes how to effectively use paging (swap) files on IO Accelerator devices with Windows® operating systems.

Using a page file with a traditional disk drive places practical limits on the usable size of the page file and virtual memory due to the poor performance of disk drives in relation to RAM. Placing the operating system paging file on one or more IO Accelerators enables much larger page files and usable virtual memory. This is due to the much faster response times and bandwidth on IO Accelerators versus hard disks. IO Accelerator software prior to version 2.0.2 does not support paging files.

Configuring IO Accelerator paging support

The IO Accelerator 2.0.2 driver and later versions can be configured to support paging files on one or more IO Accelerators. Configuring paging support requires that each IO Accelerator used with a paging file preallocates the worst-case amount of memory it might need in any possible I/O scenario. Memory is pre-allocated on a per-adapter (ioDIMM) instance.

Because of the extra host RAM memory use, paging must be enabled only on IO Accelerators that actually hold a paging file. You can place a single paging file on more than one IO Accelerator. Windows® operating systems stripe paging I/O across all available paging files, providing additional performance to the VM subsystem.

RAM consumption

The amount of RAM preallocated per IO Accelerator depends on the total size of the device and the sector (block) size selected when formatting the drive (with `fiio-format`). The following table lists the estimated memory consumption for the various device and sector sizes.

IO Accelerator size (GB)	512 byte sector	4096 byte sector
80	2,278	284
160	4,556	569
320	9,132	1,141
640	18,264	2,282

Host RAM Consumption in Mebibytes (MiB); $1024 \times 1024 = 1 \text{ MiB}$

Using a larger sector size significantly reduces the amount of host memory consumption needed for paging support. HP recommends using a 4K sector size because:

- that is generally the natural size of a host memory page, and
- it minimizes overall host memory consumption.

On Windows® operating systems, NTFS will generally use a cluster size of 4K, so formatting to 512 is not useful except for applications that are compatible only with 512-byte sector sizes (such as Windows® XP and Windows Server® 2003).

The indicated amount is needed per IO Accelerator that supports paging. You must carefully plan which IO Accelerators are used to hold a paging file.

Non-paged memory pool

Pre-allocated memory for the IO Accelerator comes from the Windows kernel non-paged memory pool. This pool dynamically grows as system components consume additional kernel memory. The maximum size of this pool is restricted as follows:

- **Windows Server 2003, 2008 R1/R2, 2012**—75% of RAM up to a maximum of 128GB.
- **Windows Vista and Windows 7**—40% of RAM up to a maximum of 128GB.

The amount of in-use non-paged pool memory should be noted when planning page file usage. This is because the IO Accelerator preallocates RAM and that reduces the available physical non-paged memory. The driver fails to load if the total preallocated memory plus the in-use non-paged memory exceeds the maximum non-paged memory pool.

To determine the total non-paged memory pool use for two IO Accelerators, use the following example:

- One IO Accelerator is 80GB and the other is 160GB
- Both are formatted with a 4K sector size
- Both will support paging files

The current allocated non-paged pool is obtained from Task Manager and, in this example, has a value of 576 MiB. (Values shown in Task Manager are in MiB (1024x1024 = 1 MiB)). The total RAM on the system is 8,196 MiB and the operating is Windows Server 2008 R2.

To calculate the total available non-paged pool, use the following formula:

```
(8196 MiB x 0.75) - 576 - 284 (80GB IO Accelerator) - 569 (160GB IO Accelerator)
```

This still leaves 4,718 MiB available for non-paged pool.

Enabling and disabling paging support

Memory preallocation occurs during IO Accelerator driver initialization. To enable paging support, you must enable the `FIO_PREALLOCATE_MEMORY` configuration item. This can be done using the `fio-config` command-line utility. This parameter is assigned a string with a list of decimal serial numbers of the IO Accelerators that will support a paging file. The driver performs memory preallocation for those instances.

The following example illustrates using the `fio-config` utility, which enables paging and preallocation on two IO Accelerators with serial numbers 1234 and 17834. Serial number information can be obtained using the `fio-status` utility.

```
fio-config -p FIO_PREALLOCATE_MEMORY "1234,17834"
```

To disable paging support on all drives, use a value of 0 for `FIO_PREALLOCATE_MEMORY`:

```
fio-config -p FIO_PREALLOCATE_MEMORY "0"
```

To query the current value, run this command:

```
fio-config -g FIO_PREALLOCATE_MEMORY
```

An alternate method to manage (enable or disable) paging support is to use the IO Accelerator Management Tool.

NOTE: You must reload the IO Accelerator driver for the new preallocation setting to take effect. Typically this can be done by restarting the machine or using **disable/enable** within **Device Manager** for each IO Accelerator instance.

Using the **Windows System Properties** to change paging file configuration requires a system restart before the properties are applied. Therefore, you can change both FIO_PREALLOCATE_MEMORY and the system page file configuration and then apply both with a single restart.

Windows page file management

By default, the IO Accelerator driver disables support for page files. The previous section described how to enable support for page files on one or more IO Accelerators. The following section describes how to work with the built-in Windows® control panels to configure and set up paging files on IO Accelerators.

Setting up page files

To set up page files in Windows® operating systems:

1. In the **Control Panel**, double-click **System**.
2. On the **Task** pane, click **Advanced system settings**.
3. On the **Advanced** tab, click **Settings**. The **Performance Options** dialog box opens.
4. On the **Advanced** tab, click **Change**. The **Virtual Memory** dialog box opens.
5. Using this dialog box, configure a page file for each available drive in the system. Selecting the **Automatically manage paging file size for all drives** checkbox causes Windows® operating systems to create a single page file on the system drive, which is the drive the operating system starts from. This checkbox should be cleared when using an IO Accelerator with a paging file.

Windows® operating systems support up to 16 distinct paging files. To enable a page file on an IO Accelerator:

1. Choose the IO Accelerator from the device list.
2. Select the **Custom size** button.
3. Provide values in the **Initial size** and **Maximum size** fields.
4. Click **Set** to save the setting. Do not omit this step, or your changes will be lost.
5. Click **OK**.
6. When prompted to restart, click **Yes**. This is necessary for the new page file settings to take effect.

To remove a paging file on the drive, follow the steps above but select **No paging file**. For performance reasons, typically you will remove all paging files on any system hard disk.

NOTE: The Virtual Memory dialog box enables pages files to be configured on available IO Accelerators, even if the IO Accelerator has not been configured to support a page file. Even though the dialog box enables the page file, following the required restart, no page file is created on the device. Follow the directions in "Enabling and disabling paging support (on page 61)" to properly enable page file support on one or more IO Accelerators.

System drive paging file configuration

By default, Windows® operating systems create and manage a page file on the system boot drive (typically a hard disk), which is where the operating system is installed. Maintaining a regular page file on the system hard disk is typically not optimal, because the hard disk I/O performance is many orders of magnitude slower than an IO Accelerator. To resolve this issue, you can eliminate or minimize the size of the system boot drive page file, as explained later. Enabling page files on IO Accelerators but not the system drive improves VM subsystem performance since the VM manager stripes I/O across all available page files. Additionally, the IO Accelerators act as a large memory store, which can greatly improve memory usage for large applications.

The Windows® kernel uses the system disk page file to store crash dumps. Crash dumps might be small (mini-dumps) or large (full-kernel memory dumps). Typically, running without dump file support or with a small dump file is adequate. There are several possible system drive page file configurations:

- Eliminate all page files on any hard disks, including the system boot drive. Although this maximizes paging I/O on IO Accelerators, no post-mortem crash dump file will be available if a system crash occurs. However, you might be able to re-enable a page file on the system drive and then reproduce the crash scenario.
- Create a minimal-size page file on the system boot drive. The recommended minimum size is 16MB, although the operating system might warn that a minimum 400MB page file is needed.
- Create a page file large enough for a full-kernel memory dump. This typically requires a page file at least the size of installed RAM, with some recommending the size equal to RAM x 1.5.

To view or change the crash dump configuration:

1. Go to the **System Properties** dialog box.
2. Click the **Advanced** tab.
3. In the **Startup and Recovery** section, click **Settings**. The **Startup and Recovery** dialog box opens.
4. In the **System Failure** section, you can change settings to handle the system log, restart, and debugging information.

Guaranteeing minimum committable memory

If you enable **System managed size** or set a **Custom size** in the **Virtual Memory** dialog box, do so with care. If the initial size is less than the desired amount of committable virtual memory, this size might cause an application to have memory allocation failures if the amount of committed memory exceeds the currently allocated page file size or the initial size value. When committed memory exceeds the current page file size, a request to allocate additional memory will fail. The Windows Virtual Memory manager will slowly increase the size of the paging file up to the available size of its drive or to the **Maximum size** custom setting, whichever is smaller.

If you want to use a large amount of committed virtual memory (more than 1.5 times the amount of RAM) and avoid application memory allocation errors, the initial and maximum committed memory must be explicitly set for the expected application committed memory usage. These values should generally be the same.

The following articles explain in detail how to size the page file appropriately.

- Main Article Link: Pushing the Limits of Windows (<http://blogs.technet.com/b/markrussinovich/archive/2008/07/21/3092070.aspx>)
- Specific section that documents virtual memory: Pushing the Limits of Windows: Virtual Memory (<http://blogs.technet.com/b/markrussinovich/archive/2008/11/17/3155406.aspx>)

Verifying page file operation

To verify that a page file is actively placed on an IO Accelerator, you can browse for hidden files at the drive root. For example, run the following command at a prompt:

```
dir c: /ah
```

In the output listing there should be a file called `pagefile.sys`. If no page file is present, recheck the page file configuration in the **Virtual Memory** dialog box to verify that page file support has been enabled on the queried IO Accelerator.

Virtual Memory performance

Using the IO Accelerator as the paging store can improve overall Virtual Memory system performance. Actual benefits will vary widely with the application virtual memory usage and with hardware platform and performance.

Performance and tuning

Introduction to performance and tuning

HP IO Accelerator devices provide high bandwidth and high IOPS and are specifically designed to achieve low latency.

As IO Accelerator devices improve in IOPS and low latency, the device performance may be limited by operating system settings and BIOS configuration. To take advantage of the revolutionary performance of IO Accelerator devices, you might have to tune these settings.

While IO Accelerator devices generally perform well out of the box, this section describes some of the common areas where tuning may help achieve optimal performance.

Disabling DVFS

DVFS is a power management technique that adjusts the CPU voltage and frequency to reduce power consumption by the CPU. These techniques help conserve power and reduce the heat generated by the CPU, but they adversely affect performance while the CPU transitions between low-power and high-performance states.

These power-savings techniques are known to have a negative impact on I/O latency and maximum IOPS. When tuning for maximum performance, you might benefit from reducing or disabling DVFS completely, even though this might increase power consumption.

DVFS, if available, should be configurable as part of your operating systems power management features as well as within your system BIOS interface. Within the operating system and BIOS, DVFS features are often found under the ACPI sections. Consult your computer documentation for details.

Limiting ACPI C-states

Newer processors have the ability to go into lower power modes when they are not fully utilized. These idle states are known as ACPI C-states. The C0 state is the normal, full power, operating state. Higher C-states (C1, C2, C3, and so on) are lower power states.

While ACPI C-states save on power, they are known to have a negative impact on I/O latency and maximum IOPS. With each higher C-state, typically more processor functions are limited to save power, and it takes time to restore the processor to the C0 state.

These power savings techniques are known to have a negative impact on I/O latency and maximum IOPS. When tuning for maximum performance you might benefit from limiting the C-states or turning them off completely, even though this might increase power consumption.

If your processor has ACPI C-states available, you can typically limit or disable them in the BIOS interface (sometimes referred to as a Setup Utility). ACPI C-states might be part of the ACPI menu. For details, see your computer documentation.

Setting NUMA affinity

Servers with a NUMA (Non-Uniform Memory Access) architecture require special installation instructions in order to maximize ioMemory device performance. These servers include the HP ProLiant DL580 and HP DL980 Servers.

On servers with NUMA architecture, during system boot, the BIOS on some systems will not distribute PCIe slots evenly among the NUMA nodes. Each NUMA node contains multiple CPUs. This imbalanced distribution means that, during high workloads, half or more of the CPUs might remain idle while the rest are 100% utilized. To prevent this imbalance, you must manually assign IO Accelerator devices equally among the available NUMA nodes.

For information on setting NUMA affinity, see "NUMA configuration (on page 67)."

Setting the interrupt handler affinity

Device latency can be affected by placement of interrupts on NUMA systems. HP recommends placing interrupts for a given device on the same NUMA socket that the application is issuing I/O from. If the CPUs on this socket are overwhelmed with user application tasks, in some cases it might benefit performance to move the interrupts to a remote socket to help load balance the system.

Many operating systems will attempt to dynamically place interrupts for you and generally make good decisions. Hand tuning interrupt placement is an advanced option that requires profiling of application performance on any given hardware. For information on how to pin interrupts for a given device to specific CPUs, see your operating system documentation.

NUMA configuration

Introduction to NUMA architecture

Servers with NUMA (Non-Uniform Memory Access) architecture require special installation instructions in order to maximize IO Accelerator device performance. These servers include the HP DL580 and the HP DL980 server.

On servers with NUMA architecture, during system boot, the BIOS on some systems will not distribute PCIe slots evenly among the NUMA nodes. Each NUMA node contains multiple CPUs. This imbalanced distribution means that, during high workloads, half or more of the CPUs will remain idle while the rest are 100% utilized. To prevent this imbalance, you must manually assign IO Accelerator devices equally among the available NUMA nodes.

Configuring the IO Accelerator devices for servers with NUMA architecture requires the use of the `FIO_AFFINITY` parameter with the `fio-config` utility.

FIO_AFFINITY parameter

The `FIO_AFFINITY` parameter is a list of `<affinity specification>` triplets that specify the affinity settings of all adapters in the system. Each item in the triplet is separated by a comma, and each triplet set is separated by a semicolon.

Syntax

```
fio-config -p FIO_AFFINITY <affinity specification>[;<affinity specification>...]
```

Where each `<affinity specification>` has the following syntax:

```
<fct number>,[g|n]<group or node number>[,<hex mask>]
```

If there is no `g` or `n` character before the group or node number, then the number is assumed to be a group number.

The hex mask is optional. If it is not present, the mask is assumed to be `0xffffffffffffffff`. Also, the `0x` prefix is optional.

If the hex mask is a node mask, then the mask is relative to the node, not the group to which the node belongs.

Simple example

```
fio-config -p FIO_AFFINITY 4,n1,0xf;5,n0;7,g1;9,g2,0xff0
```

This command creates the following:

Device	Node/Group	Processor Affinity
fct4	node 1	Processors 0 to 3 in the node (mask 0xf)
fct5	node 0	All processors in the node (no hex mask)
fct7	group 1	All processors in the group (no hex mask)
fct9	group 2	Processors 4 to 11 in the group (mask 0xff0)

Advanced configuration example

This example server has four NUMA nodes with eight hyper-threaded cores per node (16 logical processors per node, a total of 64 logical processors in the system). This system also uses the expansion configuration and has 11 PCIe expansion slots. During system boot, the system BIOS assigns PCIe slots 1-6 to NUMA node 2 and PCIe slots 7-11 to NUMA node 0. NUMA nodes 1 and 3 have no assigned PCIe slots. This configuration creates a load balancing problem in the system when IO Accelerator devices are under heavy traffic. During these periods of high use, half of the processors in the system sit idle while the other half of the processors are 100% utilized, thus limiting the throughput of the IO Accelerator devices.

To avoid this situation, you must manually configure the affinity of the IO Accelerator devices using the `FIO_AFFINITY` configuration parameter to distribute the work load across all NUMA nodes. This parameter overrides the default behavior of the IO Accelerator driver. For more information about the `FIO_AFFINITY` configuration parameter, refer to the syntax explanation below.

Syntax:

The following is an example of how to configure 10 HP IO Accelerator ioDrive Duo devices (each with two IO Accelerator devices) in a HP DL580 G7 system manually as described in the preceding paragraphs. Slot 1 is a Generation 1 PCIe slot, so it is not compatible with an ioDrive Duo device. Therefore you can fill slots 2-11 with ioDrive Duo devices. Because each ioDrive Duo device has two IO Accelerator devices, each ioDrive Duo devices has two device numbers (one for each IO Accelerator device). Each slot has two device numbers.

The following tables list the default BIOS NUMA node assignments.

BIOS assigned NUMA node	PCIe slots	FCT device numbers	Processor Affinity
0	7-11	8,9,13,14,18,19,23,24,28,29	All processors in the node
1	None	None	None
2	2-6	135,136,140,141,145,146,150,151,155,156	All processors in the node
3	None	None	None

Assigned NUMA node	PCIe slots	FCT device numbers	Processor Affinity
0	7-9	8,9,13,14,18,19	All processors in the node (no hex mask)
1	10-11	23,24,28,29	All processors in the node (no hex mask)
2	2-3	135,136,140,141	All processors in the node (no hex mask)
3	4-6	145,146,150,151,155,156	All processors in the node (no hex mask)

In this example, the BIOS creates a load imbalance by assigning the cards to only two NUMA nodes in the system. To balance the work load, enter the following settings:

To configure the IO Accelerator driver with these override settings, run `fio-config` with the following string:

```
fio-config -p FIO_AFFINITY
8,n0;9,n0;13,n0;14,n0;18,n0;19,n0;23,n1;24,n1;28,n1;29,n1;135,n2;136,n2;140,
n2;141,n2;145,n3;146,n3;150,n3;151,n3;155,n3;156,n3
```

No `<hex mask>` was stipulated for any of the devices in this example (making each `<affinity specification>` a couplet rather than a triplet). Therefore, each device is shared among all of the

processors in the specified nodes. To fine tune the configuration, you can stipulate the hex mask in each `<affinity specification>` and specify the processors for each device.

Checking the log for errors

If you enter a configuration that is not valid, the settings will be disabled and an error will be display in the system logs.

Example:

```
fio-config -p FIO_AFFINITY 5,g0,0xf;6,0xf
```

In this example, the affinity for device `fc6` is set incorrectly, because there is no group or node number before the mask. The following errors appear in the system log:

```
2011-09-09T12:22:15.176086800Z - ERROR - FusionEventDriver - FIO_AFFINITY:
Invalid group or node number
2011-09-09T12:22:15.176086800Z - ERROR - FusionEventDriver - Invalid
FIO_AFFINITY parameter syntax at character 13: "5,g0,0xf;6,0x". Manual
affinity settings are disabled!
```

Resources

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website (<http://www.hp.com/support>).

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

For more information

For additional information, see the following HP websites:

- HP BladeSystem technical resources (<http://www.hp.com/go/bladepystem/documentation>) (white papers and support documents)
- HP BladeSystem components (<http://h18004.www1.hp.com/products/ blades/ components/ c-class-compmatrix.html>)
- HP support (<http://www.hp.com/support>)

Support and other resources

Before you contact HP

Be sure to have the following information available before you call HP:

- Active Health System log (HP ProLiant Gen8 or later products)
Download and have available an Active Health System log for 3 days before the failure was detected. For more information, see the *HP iLO 4 User Guide* or *HP Intelligent Provisioning User Guide* on the HP website (<http://www.hp.com/go/ilo/docs>).
- Onboard Administrator SHOW ALL report (for HP BladeSystem products only)
For more information on obtaining the Onboard Administrator SHOW ALL report, see the HP website (<http://h20000.www2.hp.com/bizsupport/TechSupport/Document.jsp?lang=en&cc=us&objectID=c02843807>).
- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Product identification number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

HP contact information

For United States and worldwide contact information, see the Contact HP website (<http://www.hp.com/go/assistance>).

In the United States:

- To contact HP by phone, call 1-800-334-5144. For continuous quality improvement, calls may be recorded or monitored.
- If you have purchased a Care Pack (service upgrade), see the Support & Drivers website (<http://www8.hp.com/us/en/support-drivers.html>). If the problem cannot be resolved at the website, call 1-800-633-3600. For more information about Care Packs, see the HP website (<http://pro-aq-sama.houston.hp.com/services/cache/10950-0-0-225-121.html>).

Customer Self Repair

HP products are designed with many Customer Self Repair (CSR) parts to minimize repair time and allow for greater flexibility in performing defective parts replacement. If during the diagnosis period HP (or HP service

providers or service partners) identifies that the repair can be accomplished by the use of a CSR part, HP will ship that part directly to you for replacement. There are two categories of CSR parts:

- **Mandatory**—Parts for which customer self repair is mandatory. If you request HP to replace these parts, you will be charged for the travel and labor costs of this service.
- **Optional**—Parts for which customer self repair is optional. These parts are also designed for customer self repair. If, however, you require that HP replace them for you, there may or may not be additional charges, depending on the type of warranty service designated for your product.

NOTE: Some HP parts are not designed for customer self repair. In order to satisfy the customer warranty, HP requires that an authorized service provider replace the part. These parts are identified as "No" in the Illustrated Parts Catalog.

Based on availability and where geography permits, CSR parts will be shipped for next business day delivery. Same day or four-hour delivery may be offered at an additional charge where geography permits. If assistance is required, you can call the HP Technical Support Center and a technician will help you over the telephone. HP specifies in the materials shipped with a replacement CSR part whether a defective part must be returned to HP. In cases where it is required to return the defective part to HP, you must ship the defective part back to HP within a defined period of time, normally five (5) business days. The defective part must be returned with the associated documentation in the provided shipping material. Failure to return the defective part may result in HP billing you for the replacement. With a customer self repair, HP will pay all shipping and part return costs and determine the courier/carrier to be used.

For more information about HP's Customer Self Repair program, contact your local service provider. For the North American program, refer to the HP website (<http://www.hp.com/go/selfrepair>).

Réparation par le client (CSR)

Les produits HP comportent de nombreuses pièces CSR (Customer Self Repair = réparation par le client) afin de minimiser les délais de réparation et faciliter le remplacement des pièces défectueuses. Si pendant la période de diagnostic, HP (ou ses partenaires ou mainteneurs agréés) détermine que la réparation peut être effectuée à l'aide d'une pièce CSR, HP vous l'envoie directement. Il existe deux catégories de pièces CSR:

Obligatoire - Pièces pour lesquelles la réparation par le client est obligatoire. Si vous demandez à HP de remplacer ces pièces, les coûts de déplacement et main d'œuvre du service vous seront facturés.

Facultatif - Pièces pour lesquelles la réparation par le client est facultative. Ces pièces sont également conçues pour permettre au client d'effectuer lui-même la réparation. Toutefois, si vous demandez à HP de remplacer ces pièces, l'intervention peut ou non vous être facturée, selon le type de garantie applicable à votre produit.

REMARQUE: Certaines pièces HP ne sont pas conçues pour permettre au client d'effectuer lui-même la réparation. Pour que la garantie puisse s'appliquer, HP exige que le remplacement de la pièce soit effectué par un Mainteneur Agréé. Ces pièces sont identifiées par la mention "Non" dans le Catalogue illustré.

Les pièces CSR sont livrées le jour ouvré suivant, dans la limite des stocks disponibles et selon votre situation géographique. Si votre situation géographique le permet et que vous demandez une livraison le jour même ou dans les 4 heures, celle-ci vous sera facturée. Pour bénéficier d'une assistance téléphonique, appelez le Centre d'assistance technique HP. Dans les documents envoyés avec la pièce de rechange CSR, HP précise s'il est nécessaire de lui retourner la pièce défectueuse. Si c'est le cas, vous devez le faire dans le délai indiqué, généralement cinq (5) jours ouvrés. La pièce et sa documentation doivent être retournées dans l'emballage fourni. Si vous ne retournez pas la pièce défectueuse, HP se réserve le droit de vous facturer les coûts de remplacement. Dans le cas d'une pièce CSR, HP supporte l'ensemble des frais d'expédition et de retour, et détermine la société de courses ou le transporteur à utiliser.

Pour plus d'informations sur le programme CSR de HP, contactez votre Mainteneur Agréé local. Pour plus d'informations sur ce programme en Amérique du Nord, consultez le site Web HP (<http://www.hp.com/go/selfrepair>).

Riparazione da parte del cliente

Per abbreviare i tempi di riparazione e garantire una maggiore flessibilità nella sostituzione di parti difettose, i prodotti HP sono realizzati con numerosi componenti che possono essere riparati direttamente dal cliente (CSR, Customer Self Repair). Se in fase di diagnostica HP (o un centro di servizi o di assistenza HP) identifica il guasto come riparabile mediante un ricambio CSR, HP lo spedisce direttamente al cliente per la sostituzione. Vi sono due categorie di parti CSR:

Obbligatorie – Parti che devono essere necessariamente riparate dal cliente. Se il cliente ne affida la riparazione ad HP, deve sostenere le spese di spedizione e di manodopera per il servizio.

Opzionali – Parti la cui riparazione da parte del cliente è facoltativa. Si tratta comunque di componenti progettati per questo scopo. Se tuttavia il cliente ne richiede la sostituzione ad HP, potrebbe dover sostenere spese aggiuntive a seconda del tipo di garanzia previsto per il prodotto.

NOTA: alcuni componenti HP non sono progettati per la riparazione da parte del cliente. Per rispettare la garanzia, HP richiede che queste parti siano sostituite da un centro di assistenza autorizzato. Tali parti sono identificate da un "No" nel Catalogo illustrato dei componenti.

In base alla disponibilità e alla località geografica, le parti CSR vengono spedite con consegna entro il giorno lavorativo seguente. La consegna nel giorno stesso o entro quattro ore è offerta con un supplemento di costo solo in alcune zone. In caso di necessità si può richiedere l'assistenza telefonica di un addetto del centro di supporto tecnico HP. Nel materiale fornito con una parte di ricambio CSR, HP specifica se il cliente deve restituire dei componenti. Qualora sia richiesta la resa ad HP del componente difettoso, lo si deve spedire ad HP entro un determinato periodo di tempo, generalmente cinque (5) giorni lavorativi. Il componente difettoso deve essere restituito con la documentazione associata nell'imballo di spedizione fornito. La mancata restituzione del componente può comportare la fatturazione del ricambio da parte di HP. Nel caso di riparazione da parte del cliente, HP sostiene tutte le spese di spedizione e resa e sceglie il corriere/vettore da utilizzare.

Per ulteriori informazioni sul programma CSR di HP contattare il centro di assistenza di zona. Per il programma in Nord America fare riferimento al sito Web HP (<http://www.hp.com/go/selfrepair>).

Customer Self Repair

HP Produkte enthalten viele CSR-Teile (Customer Self Repair), um Reparaturzeiten zu minimieren und höhere Flexibilität beim Austausch defekter Bauteile zu ermöglichen. Wenn HP (oder ein HP Servicepartner) bei der Diagnose feststellt, dass das Produkt mithilfe eines CSR-Teils repariert werden kann, sendet Ihnen HP dieses Bauteil zum Austausch direkt zu. CSR-Teile werden in zwei Kategorien unterteilt:

Zwingend – Teile, für die das Customer Self Repair-Verfahren zwingend vorgegeben ist. Wenn Sie den Austausch dieser Teile von HP vornehmen lassen, werden Ihnen die Anfahrt- und Arbeitskosten für diesen Service berechnet.

Optional – Teile, für die das Customer Self Repair-Verfahren optional ist. Diese Teile sind auch für Customer Self Repair ausgelegt. Wenn Sie jedoch den Austausch dieser Teile von HP vornehmen lassen möchten, können bei diesem Service je nach den für Ihr Produkt vorgesehenen Garantiebedingungen zusätzliche Kosten anfallen.

HINWEIS: Einige Teile sind nicht für Customer Self Repair ausgelegt. Um den Garantieanspruch des Kunden zu erfüllen, muss das Teil von einem HP Servicepartner ersetzt werden. Im illustrierten Teilekatalog sind diese Teile mit „No“ bzw. „Nein“ gekennzeichnet.

CSR-Teile werden abhängig von der Verfügbarkeit und vom Lieferziel am folgenden Geschäftstag geliefert. Für bestimmte Standorte ist eine Lieferung am selben Tag oder innerhalb von vier Stunden gegen einen Aufpreis verfügbar. Wenn Sie Hilfe benötigen, können Sie das HP technische Support Center anrufen und sich von einem Mitarbeiter per Telefon helfen lassen. Den Materialien, die mit einem CSR-Ersatzteil geliefert werden, können Sie entnehmen, ob das defekte Teil an HP zurückgeschickt werden muss. Wenn es erforderlich ist, das defekte Teil an HP zurückzuschicken, müssen Sie dies innerhalb eines vorgegebenen Zeitraums tun, in der Regel innerhalb von fünf (5) Geschäftstagen. Das defekte Teil muss mit der zugehörigen Dokumentation in der Verpackung zurückgeschickt werden, die im Lieferumfang enthalten ist. Wenn Sie das defekte Teil nicht zurückschicken, kann HP Ihnen das Ersatzteil in Rechnung stellen. Im Falle von Customer Self Repair kommt HP für alle Kosten für die Lieferung und Rücksendung auf und bestimmt den Kurier-/Frachtdienst.

Weitere Informationen über das HP Customer Self Repair Programm erhalten Sie von Ihrem Servicepartner vor Ort. Informationen über das CSR-Programm in Nordamerika finden Sie auf der HP Website unter (<http://www.hp.com/go/selfrepair>).

Reparaciones del propio cliente

Los productos de HP incluyen muchos componentes que el propio usuario puede reemplazar (*Customer Self Repair*, CSR) para minimizar el tiempo de reparación y ofrecer una mayor flexibilidad a la hora de realizar sustituciones de componentes defectuosos. Si, durante la fase de diagnóstico, HP (o los proveedores o socios de servicio de HP) identifica que una reparación puede llevarse a cabo mediante el uso de un componente CSR, HP le enviará dicho componente directamente para que realice su sustitución. Los componentes CSR se clasifican en dos categorías:

- **Obligatorio:** componentes para los que la reparación por parte del usuario es obligatoria. Si solicita a HP que realice la sustitución de estos componentes, tendrá que hacerse cargo de los gastos de desplazamiento y de mano de obra de dicho servicio.
- **Opcional:** componentes para los que la reparación por parte del usuario es opcional. Estos componentes también están diseñados para que puedan ser reparados por el usuario. Sin embargo, si precisa que HP realice su sustitución, puede o no conllevar costes adicionales, dependiendo del tipo de servicio de garantía correspondiente al producto.

NOTA: Algunos componentes no están diseñados para que puedan ser reparados por el usuario. Para que el usuario haga valer su garantía, HP pone como condición que un proveedor de servicios autorizado realice la sustitución de estos componentes. Dichos componentes se identifican con la palabra "No" en el catálogo ilustrado de componentes.

Según la disponibilidad y la situación geográfica, los componentes CSR se enviarán para que lleguen a su destino al siguiente día laborable. Si la situación geográfica lo permite, se puede solicitar la entrega en el mismo día o en cuatro horas con un coste adicional. Si precisa asistencia técnica, puede llamar al Centro de asistencia técnica de HP y recibirá ayuda telefónica por parte de un técnico. Con el envío de materiales para la sustitución de componentes CSR, HP especificará si los componentes defectuosos deberán devolverse a HP. En aquellos casos en los que sea necesario devolver algún componente a HP, deberá hacerlo en el periodo de tiempo especificado, normalmente cinco días laborables. Los componentes defectuosos deberán devolverse con toda la documentación relacionada y con el embalaje de envío. Si no enviara el componente defectuoso requerido, HP podrá cobrarle por el de sustitución. En el caso de todas

sustituciones que lleve a cabo el cliente, HP se hará cargo de todos los gastos de envío y devolución de componentes y escogerá la empresa de transporte que se utilice para dicho servicio.

Para obtener más información acerca del programa de Reparaciones del propio cliente de HP, póngase en contacto con su proveedor de servicios local. Si está interesado en el programa para Norteamérica, visite la página web de HP siguiente (<http://www.hp.com/go/selfrepair>).

Customer Self Repair

Veel onderdelen in HP producten zijn door de klant zelf te repareren, waardoor de reparatieduur tot een minimum beperkt kan blijven en de flexibiliteit in het vervangen van defecte onderdelen groter is. Deze onderdelen worden CSR-onderdelen (Customer Self Repair) genoemd. Als HP (of een HP Service Partner) bij de diagnose vaststelt dat de reparatie kan worden uitgevoerd met een CSR-onderdeel, verzendt HP dat onderdeel rechtstreeks naar u, zodat u het defecte onderdeel daarmee kunt vervangen. Er zijn twee categorieën CSR-onderdelen:

Verplicht: Onderdelen waarvoor reparatie door de klant verplicht is. Als u HP verzoekt deze onderdelen voor u te vervangen, worden u voor deze service reiskosten en arbeidsloon in rekening gebracht.

Optioneel: Onderdelen waarvoor reparatie door de klant optioneel is. Ook deze onderdelen zijn ontworpen voor reparatie door de klant. Als u echter HP verzoekt deze onderdelen voor u te vervangen, kunnen daarvoor extra kosten in rekening worden gebracht, afhankelijk van het type garantieservice voor het product.

OPMERKING: Sommige HP onderdelen zijn niet ontwikkeld voor reparatie door de klant. In verband met de garantievoorwaarden moet het onderdeel door een geautoriseerde Service Partner worden vervangen. Deze onderdelen worden in de geïllustreerde onderdelencatalogus aangemerkt met "Nee".

Afhankelijk van de leverbaarheid en de locatie worden CSR-onderdelen verzonden voor levering op de eerstvolgende werkdag. Levering op dezelfde dag of binnen vier uur kan tegen meerkosten worden aangeboden, indien dit mogelijk is gezien de locatie. Indien assistentie gewenst is, belt u een HP Service Partner om via de telefoon technische ondersteuning te ontvangen. HP vermeldt in de documentatie bij het vervangende CSR-onderdeel of het defecte onderdeel aan HP moet worden geretourneerd. Als het defecte onderdeel aan HP moet worden teruggezonden, moet u het defecte onderdeel binnen een bepaalde periode, gewoonlijk vijf (5) werkdagen, retourneren aan HP. Het defecte onderdeel moet met de bijbehorende documentatie worden geretourneerd in het meegeleverde verpakkingsmateriaal. Als u het defecte onderdeel niet terugzendt, kan HP u voor het vervangende onderdeel kosten in rekening brengen. Bij reparatie door de klant betaalt HP alle verzendkosten voor het vervangende en geretourneerde onderdeel en kiest HP zelf welke koerier/transportonderneming hiervoor wordt gebruikt.

Neem contact op met een Service Partner voor meer informatie over het Customer Self Repair programma van HP. Informatie over Service Partners vindt u op de HP website (<http://www.hp.com/go/selfrepair>).

Reparo feito pelo cliente

Os produtos da HP são projetados com muitas peças para reparo feito pelo cliente (CSR) de modo a minimizar o tempo de reparo e permitir maior flexibilidade na substituição de peças com defeito. Se, durante o período de diagnóstico, a HP (ou fornecedores/parceiros de serviço da HP) concluir que o reparo pode ser efetuado pelo uso de uma peça CSR, a peça de reposição será enviada diretamente ao cliente. Existem duas categorias de peças CSR:

Obrigatória – Peças cujo reparo feito pelo cliente é obrigatório. Se desejar que a HP substitua essas peças, serão cobradas as despesas de transporte e mão-de-obra do serviço.

Opcional – Peças cujo reparo feito pelo cliente é opcional. Essas peças também são projetadas para o reparo feito pelo cliente. No entanto, se desejar que a HP as substitua, pode haver ou não a cobrança de taxa adicional, dependendo do tipo de serviço de garantia destinado ao produto.

OBSERVAÇÃO: Algumas peças da HP não são projetadas para o reparo feito pelo cliente. A fim de cumprir a garantia do cliente, a HP exige que um técnico autorizado substitua a peça. Essas peças estão identificadas com a marca "No" (Não), no catálogo de peças ilustrado.

Conforme a disponibilidade e o local geográfico, as peças CSR serão enviadas no primeiro dia útil após o pedido. Onde as condições geográficas permitirem, a entrega no mesmo dia ou em quatro horas pode ser feita mediante uma taxa adicional. Se precisar de auxílio, entre em contato com o Centro de suporte técnico da HP para que um técnico o ajude por telefone. A HP especifica nos materiais fornecidos com a peça CSR de reposição se a peça com defeito deve ser devolvida à HP. Nos casos em que isso for necessário, é preciso enviar a peça com defeito à HP dentro do período determinado, normalmente cinco (5) dias úteis. A peça com defeito deve ser enviada com a documentação correspondente no material de transporte fornecido. Caso não o faça, a HP poderá cobrar a reposição. Para as peças de reparo feito pelo cliente, a HP paga todas as despesas de transporte e de devolução da peça e determina a transportadora/serviço postal a ser utilizado.

Para obter mais informações sobre o programa de reparo feito pelo cliente da HP, entre em contato com o fornecedor de serviços local. Para o programa norte-americano, visite o site da HP (<http://www.hp.com/go/selfrepair>).

カスタマーセルフリペア

修理時間を短縮し、故障部品の交換における高い柔軟性を確保するために、HP製品には多数のCSR部品があります。診断の際に、CSR部品を使用すれば修理ができるとHP（HPまたはHP正規保守代理店）が判断した場合、HPはその部品を直接、お客様に発送し、お客様に交換していただきます。CSR部品には以下の2通りがあります。

- **必須** - カスタマーセルフリペアが必須の部品。当該部品について、もしもお客様がHPに交換作業を依頼される場合には、その修理サービスに関する交通費および人件費がお客様に請求されます。
- **任意** - カスタマーセルフリペアが任意である部品。この部品もカスタマーセルフリペア用です。当該部品について、もしもお客様がHPに交換作業を依頼される場合には、お買い上げの製品に適用される保証サービス内容の範囲内においては、別途費用を負担していただくことなく保証サービスを受けることができます。

注： HP製品の一部の部品は、カスタマーセルフリペア用ではありません。製品の保証を継続するためには、HPまたはHP正規保守代理店による交換作業が必須となります。部品カタログには、当該部品がカスタマーセルフリペア除外品である旨が記載されています。

部品供給が可能な場合、地域によっては、CSR部品を翌営業日に届くように発送します。また、地域によっては、追加費用を負担いただくことにより同日または4時間以内に届くように発送することも可能な場合があります。サポートが必要なときは、HPの修理受付窓口へ電話していただければ、技術者が電話でアドバイスします。交換用のCSR部品または同梱物には、故障部品をHPに返送する必要があるかどうかが表示されています。故障部品をHPに返送する必要がある場合は、指定期限内（通常は5営業日以内）に故障部品をHPに返送してください。故障部品を返送する場合は、届いた時の梱包箱に関連書類とともに入れてください。故障部品を返送しない場合、HPから部品費用が請求されます。カスタマーセルフリペアの際には、HPは送料および部品返送費を全額負担し、使用する宅配便会社や運送会社を指定します。

客戶自行維修

HP 產品提供許多客戶自行維修 (CSR) 部件，以尽可能縮短維修時間和在更換缺陷部件方面提供更大的靈活性。如果在診斷期間 HP (或 HP 服務提供商或服務合作夥伴) 確定可以通過使用 CSR 部件完成維修，HP 將直接把該部件發送給您進行更換。有兩類 CSR 部件：

- **強制性的** — 要求客戶必須自行維修的部件。如果您請求 HP 更換這些部件，則必須為該服務支付差旅費和人工費用。
- **可選的** — 客戶可以選擇是否自行維修的部件。這些部件也是為客戶自行維修設計的。不過，如果您要求 HP 為您更換這些部件，則根據為您的產品指定的保修服務類型，HP 可能收取或不再收取任何附加費用。

注：某些 HP 部件的設計並未考慮客戶自行維修。為了滿足客戶保修的需要，HP 要求授權服務提供商更換相關部件。這些部件在部件圖解目錄中標記為“否”。

CSR 部件將在下一個工作日發運（取決於備貨情況和允許的地理範圍）。在允許的地理範圍內，可在當天或四小時內發運，但要收取額外費用。如果需要幫助，您可以致電 HP 技術支持中心，將會有技術人員通過電話為您提供幫助。HP 會在隨更換的 CSR 部件發運的材料中指明是否必須將有缺陷的部件返還給 HP。如果要求您將有缺陷的部件返還給 HP，那麼您必須在規定期限內（通常是五 (5) 個工作日）將缺陷部件發給 HP。有缺陷的部件必須隨所提供的發運材料中的相關文件一起返還。如果未能送還有缺陷的部件，HP 可能會要求您支付更換費用。客戶自行維修時，HP 將承擔所有相關運輸和部件返回費用，並指定快遞商/承運商。

有關 HP 客戶自行維修計劃的詳細信息，請與您當地的服務提供商聯繫。有關北美地區的計劃，請訪問 HP 網站 (<http://www.hp.com/go/selfrepair>)。

客戶自行維修

HP 產品設計了許多「客戶自行維修」(CSR) 的零件以減少維修時間，並且使得更換瑕疵零件時能有更大的彈性。如果在診斷期間 HP (或 HP 服務供應商或維修夥伴) 辨認出此項維修工作可以藉由使用 CSR 零件來完成，則 HP 將直接寄送該零件給您作更換。CSR 零件分為兩種類別：

- **強制的** — 客戶自行維修所使用的零件是強制的。如果您要求 HP 更換這些零件，HP 將會向您收取此服務所需的外出費用與勞動成本。
- **選購的** — 客戶自行維修所使用的零件是選購的。這些零件也設計用於客戶自行維修之用。不過，如果您要求 HP 為您更換，則可能需要也可能不需要負擔額外的費用，端視針對此產品指定的保固服務類型而定。

備註：某些 HP 零件沒有消費者可自行維修的設計。為符合客戶保固，HP 需要授權的服務供應商更換零件。這些零件在圖示的零件目錄中，被標示為「否」。

基於材料取得及環境允許的情況下，CSR 零件將於下一個工作日以快遞寄送。在環境的允許下當天或四小時內送達，則可能需要額外的費用。若您需要協助，可致電「HP 技術支援中心」，會有一位技術人員透過電話來協助您。不論損壞的零件是否必須退回，HP 皆會在與 CSR 替換零件一起運送的材料中註明。若要將損壞的零件退回 HP，您必須在指定的一段時間內（通常為五 (5) 個工作天），將損壞的零件寄回 HP。損壞的零件必須與寄送資料中隨附的相關技術文件一併退還。如果無法退還損壞的零件，HP 可能要向您收取替換費用。針對客戶自行維修情形，HP 將負責所有運費及零件退還費用並指定使用何家快遞/貨運公司。

如需 HP 的「客戶自行維修」方案詳細資訊，請連絡您當地的服務供應商。至於北美方案，請參閱 HP 網站 (<http://www.hp.com/go/selfrepair>)。

고객 셀프 수리

HP 제품은 수리 시간을 최소화하고 결함이 있는 부품 교체 시 더욱 용통성을 발휘할 수 있도록 하기 위해 고객 셀프 수리(CSR) 부품을 다량 사용하여 설계되었습니다. 진단 기간 동안 HP(또는 HP 서비스 공급업체 또는 서비스 협력업체)에서 CSR 부품을 사용하여 수리가 가능하다고 판단되면 HP는 해당 부품을 바로 사용자에게 보내어 사용자가 교체할 수 있도록 합니다. CSR 부품에는 두 가지 종류가 있습니다.

- **고객 셀프 수리가 의무 사항인 필수 부품.** 사용자가 HP에 이 부품의 교체를 요청할 경우 이 서비스에 대한 출장비 및 작업비가 청구됩니다.
- **고객 셀프 수리가 선택 사항인 부품.** 이 부품들도 고객 셀프 수리가 가능하도록 설계되었습니다. 하지만 사용자가 HP에 이 부품의 교체를 요청할 경우 사용자가 구입한 제품에 해당하는 보증 서비스 유형에 따라 추가 비용 없이 교체가 가능할 수 있습니다.

참고: 일부 HP 부품은 고객 셀프 수리가 불가능하도록 설계되었습니다. HP는 만족스러운 고객 보증을 위해 공인 서비스 제공업체를 통해 부품을 교체하도록 하고 있습니다. 이러한 부품들은 Illustrated Parts Catalog에 "No"라고 표시되어 있습니다.

CSR 부품은 재고 상태와 지리적 조건이 허용하는 경우 다음 영업일 납품이 가능하도록 배송이 이루어집니다. 지리적 조건이 허용하는 경우 추가 비용이 청구되는 조건으로 당일 또는 4시간 배송이 가능할 수도 있습니다. 도움이 필요하시면 HP 기술 지원 센터로 전화하십시오. 전문 기술자가 전화로 도움을 줄 것입니다. HP는 결함이 발생한 부품을 HP로 반환해야 하는지 여부를 CSR 교체 부품과 함께 배송된 자료에 지정합니다. 결함이 발생한 부품을 HP로 반환해야 하는 경우에는 지정된 기간 내(통상 영업일 기준 5일)에 HP로 반환해야 합니다. 이 때 결함이 발생한 부품은 제공된 포장 재료에 넣어 관련 설명서와 함께 반환해야 합니다. 결함이 발생한 부품을 반환하지 않는 경우 HP가 교체 부품에 대해 비용을 청구할 수 있습니다. 고객 셀프 수리의 경우, HP는 모든 운송 및 부품 반환 비용을 부담하며 이용할 운송업체 및 택배 서비스를 결정합니다.

HP 고객 셀프 수리 프로그램에 대한 자세한 내용은 가까운 서비스 제공업체에 문의하십시오. 북미 지역의 프로그램에 대해서는 HP 웹 사이트(<http://www.hp.com/go/selfrepair>)를 참조하십시오.

Regulatory information

Safety and regulatory compliance

For safety, environmental, and regulatory information, see *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the HP website (<http://www.hp.com/support/Safety-Compliance-EnterpriseProducts>).

Turkey RoHS material content declaration

Türkiye Cumhuriyeti: EEE Yönetmeliğine Uygundur

Ukraine RoHS material content declaration

Обладнання відповідає вимогам Технічного регламенту щодо обмеження використання деяких небезпечних речовин в електричному та електронному обладнанні, затвердженого постановою Кабінету Міністрів України від 3 грудня 2008 № 1057

Warranty information

HP ProLiant and X86 Servers and Options (<http://www.hp.com/support/ProLiantServers-Warranties>)

HP Enterprise Servers (<http://www.hp.com/support/EnterpriseServers-Warranties>)

HP Storage Products (<http://www.hp.com/support/Storage-Warranties>)

HP Networking Products (<http://www.hp.com/support/Networking-Warranties>)

Acronyms and abbreviations

ACPI

Advanced Configuration and Power Interface

DVFS

dynamic voltage and frequency scaling

FAT

file allocation table

IOPS

input/output operations per second

LEB

Logical Erase Block

LFM

linear feet per minute

MIB

management information base

MSI

Message Signaled Interrupt

NAND

Not AND

NTFS

NT File System

NUMA

Non-Uniform Memory Architecture

PCIe

peripheral component interconnect express

SMI-S

Storage Management Initiative Specification

SSD

solid-state drive

VM

Virtual Machine

VSL

virtual storage layer

Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (<mailto:docsfeedback@hp.com>). Include the document title and part number, version number, or the URL when submitting your feedback.

Index

A

about this guide 6
Adding a file system 28
authorized reseller 71
automated logging 35

B

battery replacement notice 79
before you contact HP 71

C

checking for outdated firmware 26
configuring IO Accelerator paging support 60
configuring RAID 29
contact information 71
contacting HP 71
contents summary 6
controlling Trim aggressiveness 58
creating a patch install log 34
creating an installation log 34
creating an uninstall log 34
CSR (customer self repair) 71
customer self repair (CSR) 71

D

defragmentation 26
disabling auto attach 37
disabling DVFS 65
disabling paging support 61
documentation feedback 82

E

enabling auto attach 38
enabling paging support 61
Enabling PCIe power 27
enabling PCIe power override 38
enabling Trim 58
error messages 69
European Union notice 79

F

FIO_AFFINITY parameters 44, 67
fio-attach utility 41
fio-beacon utility 42
fio-bugreport utility 42
fio-config options 44
fio-config utility 43
fio-detach utility 45
fio-format utility 46
fio-pci-check utility 47
fio-status utility 47
fio-sure-erase utility 50
fio-trim-config utility 51
fio-update-iodrive utility 52
firmware version 23
firmware, upgrading 19, 20
flashback substitution events 55
For more information 70
formatting a file system 28

G

guaranteeing minimum committable memory 63

H

health metrics 54
health monitoring techniques 54
HP Subscriber's Choice for Business 70
HP technical support 71
HP, contacting 71

I

installing software 13
introduction 7

L

limiting APCI C-states 65

M

maintenance 23, 24, 26, 37, 38
maintenance tools 37

manual installation on Windows Server 21
monitoring IO Accelerator health 54

N

naming 27
NAND flash and component failure 54
non-paged memory pool 61
NUMA architecture 67
NUMA configuration 67, 68

O

operating environment 9
outdated firmware 26
overview 7

P

page files, setting up 62
PCIe power override, enabling 38
PCIe power, enabling 27
performance and tuning 65
performance and tuning, introduction 65
performance attributes 8
performing a multi-disk configuration 28
phone numbers 71

R

RAID configuration 29
RAM consumption for page files 60
regulatory compliance identification numbers 79
regulatory compliance notices 79
required information 71
resources 70, 71

S

series number 79
setting interrupt handler affinity 66
setting NUMA affinity 66
setting up SNMP 29
setup wizard, using 14
silent installation 26
SNMP MIB support 31
SNMP, Windows details 29
software installation 13
software RAID and health monitoring 55
starting Trim 58
stopping Trim 58
support 71
supported firmware revisions 9

supported hardware 9
system drive paging file configuration 63
system log errors 69

T

Taiwan battery recycling notice 79
technical support 71
telephone numbers 71
Trim configurations 59
Trim support 57
Trim, introduction 57
Trim, platforms 57
troubleshooting 35

U

uninstalling a previously-installed drive 12
uninstalling the driver 37
unmanaged shutdown issues 38
upgrade procedures 23
upgrading driver software 24, 25
upgrading firmware 19, 20, 23
using installation logs 33
using test mode registry values 30
using the IO Accelerator as a swap 29
using Trim 57
utilities 41
utilities reference 41

V

verifying page file operation 64
Virtual Memory performance 64

W

Windows Installer logging options 33
Windows page file management 62
Windows page files 60
Windows page files, introduction 60
Windows Server 2003 22
Windows Server 2008 23
Windows Server 2012 23
Windows Server environments 12