



IBM ioMemory VSL 3.1.1 User Guide for VMware ESX and ESXi	ł
Introduction	;
About ioMemory Devices	5
About the ioMemory Virtual Storage Layer (VSL)	5
About Flashback Protection Technology	ó
System Requirements	7
Hardware Requirements	7
Supported VMware Operating Systems 7	7
Software Installation	3
Downloading the Software)
Installing the ioMemory VSL on ESXi 5.0 1	1
Installing the ioMemory VSL on ESX(i) 4.x 1	13
Installing the Optional SMI-S Provider 1	4
Upgrading the Firmware	16
Configuring the Device to Support VM Disks 1	17
Modifying a VMware Resource Pool to Reserve Memory 1	17
Maintenance	9
Device LED Indicators	9
Management Utilities	9
Common Maintenance Tasks	21
Unmanaged Shutdown Issues	25
Performance and Tuning	26
Disabling DVFS	26
Limiting ACPI C-States	26
Appendix A- Command-Line Utilities	28
fio-attach	28
fio-beacon 2	29
fio-bugreport	30
fio-detach	31
fio-format	31
fio-pci-check	33
fio-status	33
fio-update-iodrive	36
Appendix B- Using Python SMI-S Management Scripts 3	38
fio-attach.py	10
fio-beacon.py	1 1



fio-detach.py	42
fio-format.py	43
fio-status.py	45
fio-update-iodrive.py	46
Appendix C- Using a CIM Browser for SMI-S Management	49
SMI-S Interface Background	49
Description	50
Implementation	52
Indications	58
Appendix D- Monitoring the Health of ioMemory Devices	63
NAND Flash and Component Failure	63
Health Metrics	63
Health Monitoring Techniques	64
Appendix E- Using Module Parameters	65
Appendix F- Working with ioMemory Devices and VMDirectPathIO	66
Appendix G- Documentation Permissions	67
AVR Bootloader	67
tree.h	67
IBM Support	69



IBM ioMemory VSL 3.1.1 User Guide for VMware ESX and ESXi

Copyright Notice

© Copyright International Business Machines 2012. All rights reserved.

© Copyright 2006-2012 Fusion-io, Inc. All rights reserved. Fusion-io is a trademark of Fusion-io, Inc.

Part Number: D0004698-001_1 **Published**: August 20, 2012

Introduction

Overview

Congratulations on your purchase of a IBM solid-state storage device. This guide explains how to install, troubleshoot, and maintain the software for your IBM High IOPS Adapter(s).

- NOTE References to "ESX(i)" refers to ESX or ESXi.
- NOTE Throughout this manual, when you see a reference to an **IBM High IOPS Adapter**, you may substitute your particular device(s), such as an IBM High IOPS Adapter or each of the two IBM High IOPS Adapters of an IBM High IOPS Duo Adapter.
- Attention Products with Multiple Devices: Some products, such as IBM High IOPS Duo Adapters, are actually comprised of multiple IBM High IOPS Adapters. If your product consists of multiple IBM High IOPS Adapters, you will manage each IBM High IOPS Adapter as an independent device.

For example, if you have an IBM High IOPS Duo Adapter, you can independently attach, detach, and/or format each of the two IBM High IOPS Adapters. Each of the two devices will be presented as an individual device to your system.

About ioMemory Devices

Designed around a revolutionary silicon-based storage architecture, IBM High IOPS Adapters are the world's most advanced NAND flash storage devices, with performance comparable to DRAM and storage capacity on par with today's hard disks – giving you the power to improve storage performance by orders of magnitude. IBM High IOPS Adapters allow every computer to exceed the I/O performance of an enterprise SAN.

IBM High IOPS Adapters are data accelerators designed specifically to improve the bandwidth for I/O-bound applications. They are no-compromise solutions for the toughest computational challenges faced by data centers today, putting them in a league of their own.

About the ioMemory Virtual Storage Layer (VSL)

More than just a hardware driver, the ioMemory® Virtual Storage LayerTM (VSL) is the "secret sauce" that gives IBM High IOPS Adapters their amazing performance. The VSLTM is a hybrid of the RAM virtualization subsystem and the disk I/O subsystem, combining the best of both worlds. It appears like a disk to interface well with block-based applications and software. At the same time, it runs like RAM underneath to maximize performance. This provides the following game-changing 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 interrupts and context switching
- Extensibility: The VSL enables flash-optimized software development, making each ioMemory module a flexible building block for building a flash-optimized data center.

About Flashback Protection Technology

High IOPS memory devices have a finite life span of writes operations. Over the life of the product, NAND flash will eventually wear out as part of it's normal life cycle. To extend the life of the adapter, FlashbackTM redundancy is designed to mitigate the loss of failed memory location without interrupting normal operation.

This real-time NAND flash redundancy works at the chip-level so these losses are handled without sacrificing user capacity, performance, and provides additional protection above and beyond ECC (Error Correction Code) for soft failures.

IBM's Flashback ProtectionTM technology, with self-healing properties, ensures higher performance, minimal failure, and longer endurance than all other flash solutions.



System Requirements

Please read the IBM ioMemory VSL Release Notes for more information on this release.

Hardware Requirements

- Hardware Requirements: These depend on your device (including device capacity, generation, and configuration). Please see the *IBM High IOPS Hardware Installation Guide* for requirements on the following:
 - PCIe Slot
 - Cooling
 - Power
- **Supported Devices**: Also see the *IBM High IOPS Hardware Installation Guide* for a list of supported IBM High IOPS Adapters.
 - <u>Attention</u> Second generation adapters are not compatible with this version of the ioMemory VSL. If you have a newer IBM High IOPS Adapter that is not on this list, you may need to upgrade your system and all of your devices to ioMemory VSL 3.x software and firmware versions.
- **RAM Requirements**: The *IBM ioMemory VSL Release Notes* contains memory (RAM) requirements for this version of the software.

Supported VMware Operating Systems

- ESX 4.0 Update 4
- ESX 4.1 Update 2
- ESXi 4.1 Update 2
- ESXi 5.0



Software Installation

VMDirectPathI0

The ESX(i) ioMemory VSL software is only required if you plan to use the device as a VMFS Datastore. If, however, you are passing the device(s) through (using VMDirectPathIO – also known as PCI passthrough), you do not need to install the ioMemory VSL on the ESX(i) system. Instead, install the VSL on the guest operating system. For example, you would pass the device through to a Windows VM and then install the Windows ioMemory VSL on that VM (consult the Windows ioMemory VSL User Guide for installation and user instructions).

There are special considerations when passing through an IBM High IOPS Adapter, for more information, see <u>Appendix F- Working with ioMemory Devices and VMDirectPathIO</u> in this user guide before proceeding with passing through the device.

Command-Line Installation

Unless you use VUM, you will need to use a command-line interface to install the ioMemory VSL (driver). In order to manage the ioMemory devices and VSL, you **must** use a Command-Line Interface (CLI).

ESX Command Line

ESX includes the Console Operating System (COS). This CLI is available on the host, or through an SSH connection. The VMware **vCLI** (vSphere Command-Line Interface) also works with ESX.

ESXi Command Line

We recommend installing VMware **vCLI** (vSphere Command-Line Interface) to run against your ESXi system. You should install a vCLI package on a physical machine running Linux or Windows. For more information on VMware's vCLI, see <u>http://www.vmware.com/support/developer/vcli/</u>

We do not recommend using the vCLI on a virtual machine that is hosted on your ESXi system. The ioMemory VSL installation and configuration processes involve putting the ESXi host into maintenance mode and rebooting the host.

You may choose to use the **TSM** (Tech Support Mode), also known as **Shell** or **SSH** (when used remotely), instead of the vCLI to install the ioMemory VSL. The TSM/Shell may be required for managing/troubleshooting your device with the command-line utilities.

Attention VMware suggests that the TSM only be used "for the purposes of troubleshooting and remediation." VMware recommends using the vSphere Client or any other VMware Administration Automation Product to perform

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi

IBM

8



routine ESXi host configuration tasks that do not involve a troubleshooting scenario. For more information vi VMware's Knowledge Base article on using this mode:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=101

Installation Overview

- 1. If needed, uninstall previous versions of the ioMemory VSL and the utilities. Instructions for uninstalling current and previous versions of the driver are available in the <u>Common Maintenance Tasks</u> section. Once you have uninstalled the software, return to this page.
- 2. Install the latest version of the ioMemory VSL and the command-line utilities.
- 3. Reboot the ESX(i) system; this will load the driver and attach the IBM High IOPS Adapter(s).
- 4. Optional: Install and setup an SMI-S remote management solution. See <u>Installing the Optional SMI-S Provider</u> for details.
- 5. <u>Upgrade the Firmware</u> to the latest version, if needed (recommended).
- 6. Configure the Device to Support VM Disks.
 - Attention The IBM High IOPS Adapter is meant to be used as a data storage disk or caching device. Installing an ESX(i) operating system and booting from the IBM High IOPS Adapter is not supported.

ESX(i) installers may label an IBM High IOPS Adapters as a "VMware Block Device," and it will will permit you to install the ESX(i) OS on an IBM High IOPS Adapter. This is not supported, and the installation will fail on reboot.

Downloading the Software

Download the installation packages to a remote machine (preferably one that has the vCLI and/or vSphere client installed).

The ioMemory VSL software is available as an offline bundle from

<u>http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5083174</u>. Navigate to the appropriate folder for your operating system. Example files:

- iomemory-vsl_<version>.offline-bundle.zip
- cross_vmware-esx-drivers-block-iomemory-vsl_<version>-offline-bundle.zip

Attention The offline bundle may be within a .zip archive: iomemory-vsl-<version>.zip

NOTE The .iso image, if available, simply contains the same offline bundle that is available as a stand-alone download. In most cases, .iso images are used to create CD-ROMs so the software may be installed as part of an ESX 4.x installation. However, the ioMemory VSL should not be installed as part of an ESX installation. You should just download the offline bundle; unless you wish to transfer the files to the ESX host using a CD.



Also download the following support files that are appropriate for your version:

Package Example	Installation Instructions
fio-util- <version>.x86_64.rpm</version>	Installing the ioMemory VSL on ESX(i) <u>4.x</u> (ESX 4.x only) – This installs the ioMemory VSL utilities on ESX 4.x
iodrive_ <version>.fff</version>	Upgrading the Firmware
fusionio-cimprovider-esxi <version>-bundle-<version>.zip</version></version>	Installing the Optional SMI-S Provider (Optional)
fio-remote-util- <version>.noarch.rpm</version>	Installing Python WBEM Packages (Optional)

Attention The fio-remote-util-<version>.noarch.rpm (or .deb) package is available in the Linux download folders. Download this package from the folder for your Linux distribution (for the remote machine). In other words, if you plan to install these remote utilities on a Linux system running RHEL 5, go to the Linux rbel-5 download folder to find this package.

Transferring the ioMemory VSL Files to the ESX(i) Server

You will need to transfer the firmware file to the ESX(i) host. Also, depending on your ESX(i) version and your preferred installation method, you may need to transfer the two bundle installation files to the host as well. We recommend transferring all the files at this point, and then choosing the installation method later.

NOTE The fio-remote-util-<version>.noarch.rpm is not installed on the ESX¹ host, and therefore should not be transferred. These optional SMI-S scripts are installed on a remote Linux machine.

Whichever method you choose for transferring the file(s), we recommend saving the file(s) to a datastore on the host. The example paths to the bundles and firmware in this guide will show them located in a bundles directory on a datastore:

/vmfs/volumes/<datastore>/bundles/

Where <datatstore> is the name of the datastore. In the vSphere GUI, the folder will appear as:

[datastore] /bundle

Transfer Methods

You may transfer the file(s) using one of many methods, including:

- vSphere Client
- vCLI vifs command

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



• SCP (using SSH)

The file(s) can be copied to the host from your remote machine, or from an NFS share.

vCLI Example

Described below are the steps for transferring files to the ESX(i) host using vCLI

- 1. On your remote machine, make sure you have downloaded the appropriate files, and take note of their location.
- 2. Choose an available datastore (with at least 200MB of available storage) on the hypervisor that you will use to temporarily store the bundles.
- 3. Create a directory in the datastore named bundles using the vifs remote command:

vifs --server <servername> --mkdir "[<datastore>]bundles"

The brackets ([]) and quotes ("") are required. Substitute your datastore name for the <datastore> variable.

NOTE You will be prompted to enter the username and password for the ESXi host. For convenience, you can add the following options to each command:

--username <username> --password <password>

Attention vCLI in Windows: When using the vCLI in Windows, many of the commands are slightly different. Most of the commands end with .pl. Throughout this document, when you run the vCLI in windows, be sure to include the .pl to the command. This command would be:

vifs.pl --server <servername> --mkdir "[<datastore>]bundles"

4. Use the following example command line to transfer the file(s) one by one to the bundles directory of the datastore:

vifs --server <servername> --put "<path-on-local-machine>/<filename>"
"[<datastore>]bundles/<filename>"

Where <filename> is the full filename, for example:

- iodrive_<version>.fff
- cross_vmware-esx-drivers-block-iomemory-vsl_<version>.offline-bundle.zip
- iomemory-vsl_<version>.offline-bundle.zip.



Installing the ioMemory VSL on ESXi 5.0

- NOTE **VUM Installation**: These instructions describe how to install the ioMemory VSL on a single hypervisor. However, if you are familiar with, and use, the VMware Update Manager (VUM) plugin for the Virtual Center Server (vCenter Server), you can use that to install the ioMemory VSL on multiple hosts. Please see the vCenter Server documentation for more details on VUM.
- <u>Attention</u> Uninstall: An update/upgrade installation is not recommended. Instead, uninstall the previous version of the ioMemory VSL software before you install this version. See <u>Common Maintenance Tasks</u> for more information on uninstalling the software.

Before you install the ioMemory VSL, stop all Virtual Machines and put the ESX(i) host into maintenance mode (using the vSphere client or the vCLI).

You may choose to install the software using the vCLI or Shell (SSH). Whether you use the Shell or vCLI, you must first transfer the files to a datastore on the ESX(i) host.

Attention The offline bundle may be within a .zip archive: iomemory-vsl-<version>.zip. Unpack the offline bundle for installation.

vCLI Installation

1. Install the bundle by running the following command against your ESXi 5.0 system using the vCLI:

esxcli --server <servername> software vib install -d <offline-bundle>

Where <offline-bundle> is the absolute path to the offline bundle on the hypervisor host. For example, if the offline bundle is in the bundles directory of a datastore with the name of datastorel, the (local) path would be:

```
/vmfs/volumes/datastore1/bundles/iomemory-vsl_<version>.offline-bundle.zip
```

<u>Attention</u> This absolute path must begin with a forward slash (/) or ESXi will return an error message.

2. Reboot your ESXi system.

The ioMemory VSL and command-line utilities are installed on the host.





Command-line Installation

1. Install the bundle by running the following command against your ESXi 5.0 system using TSM/SSH:

```
esxcli software vib install -d <offline-bundle>
```

Where <offline-bundle> is the absolute path to the offline bundle on the hypervisor host. For example, if the offline bundle is in the bundles directory of a datastore with the name of datastorel, an example (local) path would be:

```
/vmfs/volumes/datastore1/bundles/iomemory-vsl_<version>.offline-bundle.zip
```

2. Reboot your ESXi system

The ioMemory VSL and command-line utilities are installed on the host.

You can now continue to <u>Installing the Optional SMI-S Provider</u>. If you wish to skip that optional step, continue on to the <u>Upgrading the Firmware</u> section.

Installing the ioMemory VSL on ESX(i) 4.x

<u>Attention</u> Uninstall: An update/upgrade installation is not recommended. Instead, uninstall the previous version of the ioMemory VSL software before you install this version. See <u>Common Maintenance Tasks</u> for more information on uninstalling the previous versions of the software.

Before you install the ioMemory VSL, stop all Virtual Machines and put the ESX(i) host into maintenance mode (using the vSphere client or the vCLI). Choose your preferred installation method:

vCLI Installation

The installation bundles shouldn't reside on the ESX(i) 4.x host when you using the vCLI. Instead, they will need to be on the remote machine.

- 1. On your remote machine, navigate to the directory that contains the downloaded files.
- 2. Install the bundle by running the following command against your ESX(i) 4.x system using the vCLI:

```
vihostupdate --server <server-name> --install --bundle <offline-bundle>
```

Where <offline-bundle> is the full path of the offline bundle **on your remote machine**. For example, on a Linux machine an example local path would be

./cross_vmware-esx-drivers-block-iomemory-vsl_<version>-offline-bundle.zip

3. Reboot your host system

The ioMemory VSL and command-line utilities are installed on the host.



Command-line Installation

You may use the COS (on ESX 4.x) or the TSM/SSH (on ESXi 4.1) to install the software. In both cases, you must first transfer the files to the host.

- 1. Navigate to the directory where you have transferred offline bundle.
- 2. Run the essupdate command to install the ioMemory VSL using the offline bundle.

\$ esxupdate --bundle <offline-bundle> update

Where <offline-bundle> is the full name of the offline bundle that you downloaded. For example, cross_vmware-esx-drivers-block-iomemory-vsl_<version>-offline-bundle.zip

3. Reboot the host system.

The ioMemory VSL and command-line utilities are installed on the host.

You may now follow the instructions on <u>Installing the Optional SMI-S Provider</u>. If you wish to skip that optional step, continue on to the <u>Upgrading the Firmware</u> section.

Installing the Optional SMI-S Provider

To manage the ioMemory VSL, you must use the provided management utilities. There are two options available for managing the VSL:

- **COS/Shell/TSM command-line utilities**: These utilities are installed with the ioMemory VSL software. In order to use these utilities on ESXi, the Shell/TSM (Tech Support Mode) must be enabled.
 - The fio-bugreport troubleshooting utility is only available as a COS/Shell/TSM command-line utility.
 - For more information about these utilities, see <u>Appendix A- Command-Line Utilities</u>
- **Remote SMI-S Scripts**: These provide remote management of the software and devices without enabling Tech Support Mode (TSM) or logging in to the COS.
 - To use the SMI-S interface, you must install the CIM (SMI-S) provider on the ESX(i) host and the Python SMI-S Management Scripts on a remote machine.
 - This section explains how to install the CIM provider.

IBM's SMI-S interface allows you to remotely manage the ioMemory VSL software on your ESX(i) system. The IBM SMI-S provider works with popular CIM servers, including SFCB. SFCB is part of a typical ESX(i) installation, and it is used by vSphere software to manage the ESX(i) system.



Installing the SMI-S Provider on ESXi 5.0

Be sure to transfer the CIM (SMI-S) provider offline bundle to the host (hypervisor) machine's local storage. For more information on transferring the offline bundle, see the <u>Installing the ioMemory VSL on ESXi 5.0</u> section.

- 1. Stop all VMs and put the host in Maintenance Mode.
- 2. Install the CIM provider while in Maintenance Mode by running the following command:

esxcli --server <servername> software vib install -d <offline-bundle>

Where <offline-bundle> is the absolute path to the offline bundle on the hypervisor host. For example, if the offline bundle is in the bundles directory of a datastore with the name of datastorel, the path would be: vmfs/volumes/datastorel/bundles/<offline-bundle> --no-sig-check

- NOTE **Command-line Installation**: You can install the CIM provider on the ESXi 5.0 host using the SSH/TSM. Simply use the same esxcli command without the --server option.
- 3. Reboot your ESXi system.

This installs the SMI-S provider and registers it with the SFCB server. You are now able to connect to the SMI-S provider.

Installing the SMI-S Provider on ESX(i) 4.x Using the vCLI

- 1. Stop all VMs and put the host in Maintenance Mode.
- 2. Navigate to the folder (on the remote machine) with the downloaded files.
- 3. Install the SMI-S provider.

```
vihostupdate --server <server-name> --install --bundle
./fusionio-cimprovider-<esx-version>-bundle-<version>.zip
--no-sig-check
```

4. Reboot the ESXi host.

This installs the SMI-S provider and registers it with the SFCB server. You are now able to connect to the SMI-S provider.

Installing the SMI-S Provider on ESX(i) 4.x using the Command-line Interface

To install the ioMemory VSL on an existing ESX(i) host using esxupdate:

1. Turn on the host and log in as administrator.



- 2. Stop all VMs and enter maintenance mode.
- 3. Navigate to the directory where you have transferred offline bundle.
- 4. Run the essupdate command to install drivers using the offline bundle.

```
$ esxupdate --bundle
fusionio-cimprovider-<esx-version>-bundle-<version>.zip update
--no-sig-check
```

5. Reboot the host system.

This installs the SMI-S provider and registers it with the SFCB server. You are now able to connect to the SMI-S provider.

Interfacing with the SMI-S Provider

There are two standard methods for managing your IBM High IOPS Adapters through the SMI-S provider. These are:

- **Python Management Scripts** (recommended): We provide Python scripts that can be implemented remotely on a Linux machine with the proper Python packages installed. For more information, see <u>Appendix B- Using</u> Python SMI-S Management Scripts.
- **CIM Browsers**: If you are familiar with Common Information Models, and are comfortable using CIM browsers (such as YAWN), then you can connect to the SMI-S provider using your preferred browser. For more information, including the Fusion-io SMI-S CIM model, see <u>Appendix C- Using a CIM Browser for SMI-S Management</u>.

Upgrading the Firmware

With the ioMemory VSL loaded, you need to check to ensure that the IBM High IOPS Adapter's firmware is up-to-date. To do this, run the <u>fio-status</u> command-line utility from Shell or the <u>fio-status.py</u> SMI-S remote script.

If the output shows that the device is running in minimal mode, use the <u>fio-update-iodrive</u> utility (in Tech Support Mode/Shell) or the <u>fio-update-iodrive.py</u> remote script to upgrade the firmware.

- NOTE When using VMDirectPathIO, if you upgrade the firmware on an IBM High IOPS Adapter, you must cycle the power to have the change take place. Just restarting the virtual machine won't apply the change.
- Attention Your IBM High IOPS Adapter may have a minimum firmware label affixed (for example, "MIN FW: XXXXXX"). This label indicates the minimum version of the firmware that is compatible with your device.
- Attention Do not attempt to downgrade the firmware on any IBM High IOPS Adapter, doing so may void your warranty.
 - NOTE When installing a new IBM High IOPS Adapter along with existing devices, it is best to upgrade all of the devices to the latest available versions of the firmware and ioMemory VSL. The latest versions are available at http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5083174.



For more information regarding firmware and ioMemory VSL versions and compatibility, contact support at <u>http://www.ibm.com/systems/support</u>.

Configuring the Device to Support VM Disks

Attention ESX(i) requires 512B sector sizes. New IBM High IOPS Adapters come pre-formatted with 512B sector sizes from the factory. If yours is a new device, there is no need to format it.

However, if your IBM High IOPS Adapter was previously used in a system that allowed for larger sector sizes (such as Linux and 4KB sectors), then you must perform a format using the <u>fio-format</u> utility or <u>fio-format.py</u> SMI-S remote script. Follow formatting instructions careully, including disabling and re-enabling autoattach.

Within the vSphere Client, select the **Configuration** tab. Under **Hardware** click **Storage**, then click **Add Storage** located on the top right corner. The Add Storage wizard will appear. Use this wizard to configure the device.

For more information, and an explanation of options (including setting the VM File System Block Size), consult your vSphere documentation.

- NOTE You can also create a VMFS datastore using fdisk and vmkfstools in the Tech Support Mode (directly on the ESXi host), however this method is not supported by VMware.
- <u>Attention</u> The preferred type of virtual disk is "eagerzeroedthick." We do not recommend "thin" provisioning, as it will degrade performance significantly.

It is now possible to store virtual machines on the IBM High IOPS Adapter(s).

Modifying a VMware Resource Pool to Reserve Memory

Under certain circumstances, the ESX(i) operating system may temporarily require most, if not all, of the RAM available on the system, leaving no memory for the ioMemory VSL.

NOTE For example, a host running VMware View may need to rapidly provision multiple VDI images. This may happen so quickly that the host memory is temporarily exhausted.

If the VMs starve the ioMemory VSL of RAM, the IBM High IOPS Adapter(s) may go offline or stop processing requests. To address this use case, follow the procedure and guidelines below for limiting memory consumed by the VMs.

We recommend limiting RAM available to the VMs equal to: Total Host RAM - RAM equivalent to 0.5% of the total IBM High IOPS Adapter capacity (see the **Example Scenario** below for more information on this calculation). The easiest way to set this limit is by modifying the user pool.

The exact amount to limit is workload dependent, and will require tuning for specific use cases.



To modify the user pool, follow the steps below, using the vSphere client:

- 1. Click the **Summary** tab in the vSphere client to view the current memory usage and capacity.
 - Also visible is the total IBM High IOPS Adapter datastore capacity, make note of that capacity.
- 2. Navigate to the user Resource Allocation window:
 - a. Select the host -> Configuration tab -> Software pane -> System Resource Allocation link -> Advanced link
 - b. The System Resource Pools appear.
 - c. Select the **user** node under the host tree.
 - d. The details for the user appear below, click the Edit settings link.
 - e. The user Resource Allocation window appears.
- 3. Limit the Memory allocated to the VMs.
 - a. Under **Memory Resources**, clear the **Unlimited** checkbox so you can set the limit for memory resource allocation.
 - b. You can now set the limit on VM memory consumption.

Example Scenario:

An ESXi host has

- Memory capacity of 36852MB
- Total IBM High IOPS Adapter datastore capacity of 320GB (or approximately 32000MB).

320000MB device capacity * 0.5% of device capacity ~ 1600MB of RAM equivalent.

36852MB total memory capacity - 1600MB free = 35252MB of memory limited to the host. The new value under **Limit** in **Memory Resources** would be 35252MB.





Maintenance

The IBM High IOPS Adapter includes both software utilities for maintaining the device as well as external LED indicators to display its status.

Device LED Indicators

The IBM High IOPS Adapter includes three LEDs showing drive activity or error conditions. The LEDs on your device should be similar to this configuration:



This table explains the information that these LEDs convey:

Green	Yellow	Amber	Indicates	Notes
0	0	0	Power off	
0	0	•	Power on (Driver not loaded and card not attached)	Load driver and attach card
	0	0	Power on (Driver loaded but card not attached)	Attach card
	(Flashing)	0	Writing (Rate indicates volume of writes)	Can appear in combination with the Read LED
(Flashing)	0	0	Reading (Rate indicates volume of reads)	Can appear in combination with the Write LED
	0	•	Location beacon	



Management Utilities

There are two options available for managing the VSL:

- **COS/Shell/TSM command-line utilities**: These utilities are installed with the ioMemory VSL software. In order to use these utilities on ESXi, the Shell/TSM (Tech Support Mode) must be enabled.
 - The fio-bugreport troubleshooting utility is only available as a COS/Shell/TSM command-line utility.
 - For more information about these utilities, see <u>Appendix A- Command-Line Utilities</u>
- **Remote SMI-S Scripts**: These provide remote management of the software and devices without enabling Tech Support Mode (TSM) or logging in to the COS.
 - To use the SMI-S interface, you must install the CIM (SMI-S) provider on the ESX(i) host and the Python SMI-S Management Scripts on a remote machine.

Command-Line Utilities for Tech Support Mode and COS

Several command-line utilities are included in the installation packages, these command-line utilities are only accessible through VMware's **Tech Support Mode** (also known as Shell/SSH) in ESXi and the COS in ESX.

- AttentionVMware suggests that the TSM only be used "for the purposes of troubleshooting and remediation."VMware recommends using the vSphere Client or any other VMware Administration Automation Product
to perform routine ESXi host configuration tasks that do not involve a troubleshooting scenario. For more
information visit VMware's Knowledge Base article on using this mode.
 - fio-attach
 - fio-beacon
 - fio-bugreport
 - fio-detach
 - fio-format
 - fio-pci-check (ESX only)
 - fio-status
 - fio-update-iodrive

For more information on command-line utilities, see Appendix A- Command-Line Utilities.



SMI-S Remote Management

The SMI-S interface is a powerful remote management option. The SMI-S provider runs on the ESX(i) system and it provides remote access to the ioMemory VSL software on that system. IBM provides Python SMI-S management scripts to remotely interface with the SMI-S provider.

The Python scripts create a management experience similar to running the command-line utilities on a local host. The following scripts are available:

- fio-attach.py
- fio-beacon.py
- fio-detach.py
- fio-format.py
- fio-status.py
- fio-update-iodrive.py

For more information on command-line utilities, see <u>Installing the Optional SMI-S Provider</u> and <u>Appendix B- Using</u> Python SMI-S Management Scripts.

Common Maintenance Tasks

In ESX, these task require the COS. In ESXi, some of these maintenance tasks are only accessible through VMware's **Tech Support Mode** (also known as Shell/SSH).

AttentionVMware suggests that the TSM only be used "for the purposes of troubleshooting and remediation."VMware recommends using the vSphere Client or any other VMware Administration Automation Product
to perform routine ESXi host configuration tasks that do not involve a troubleshooting scenario. For more
information visit VMware's Knowledge Base article on using this mode.

Disabling the ioMemory VSL (driver)

The ioMemory VSL (driver) automatically loads by default when the operating system starts. IBM does not support the driver being unloaded, so it is necessary to disable driver auto-load for diagnostic or troubleshooting purposes.

To disable driver auto-load, run these commands in COS/TSM then reboot the system:

```
$ esxcfg-module --disable iomemory-vsl
```

NOTE In ESX 4.x, you must also run the following command before you reboot:



esxcfg-boot -b

This prevents the ioMemory VSL driver from loading on boot, so the device won't be available to users. However, all other services and applications are now available.

Enabling the ioMemory VSL

To enable the ioMemory VSL Driver (on boot) after maintenance, run these commands in TSM/COS and reboot the system.

\$ esxcfg-module --enable iomemory-vsl

NOTE In ESX 4.x, you must also run the following command before you reboot:

esxcfg-boot -b

After a reboot, if the driver is enabled, then it will appear in the modules listed when this command is run:

\$ esxcfg-module --query

Detaching an IBM High IOPS Adapter

We do not recommend detaching IBM High IOPS Adapters that are used as datastores. The best practice is to:

- 1. Disable auto-attach (see the next sub-section)
- 2. Reboot
- 3. Perform the necessary maintenance operations
- 4. Re-enable auto-attach
- 5. Then reboot again.

If you must detach an IBM High IOPS Adapter, carefully read all of the warnings in the <u>fio-detach</u> or <u>fio-detach.py</u> section of this guide before running the detach utility. Failure to follow the instructions may cause errors, data loss and/or corruption.





Disabling Auto-Attach

Disabling Auto-Attach in ESX

To load the ioMemory VSL with auto-attach disabled, run the following command and then restart:

esxcfg-module -s 'auto_attach=0' iomemory-vsl

This will not be enforced until you reboot the system. To enable auto-attach, set the parameter back to 1.

Disabling Auto-Attach in ESXi

To load the ioMemory VSL on boot with auto-attach disabled, set the auto_attach parameter equal to 0 using the vCLI:

\$ vicfg-module --server <server-name> iomemory-vsl -s 'auto_attach=0'

This will not be enforced until you reboot the system. To enable auto-attach, set the parameter back to 1.

Uninstalling the ioMemory VSL Package and Command-line Utilities

Uninstalling the Software in ESX 4.x

To uninstall the ioMemory VSL package, run these commands:

1. Find the ioMemory VSL Bulletin ID:

\$ esxupdate query

Sample Output

-----Bulletin ID----- ----Installed----iomemory-vsl-3.0.6.360 2012-01-16T03:49:33 iomemory-vsl: block driver for ESX/ESXi 4.X

2. Remove ioMemory VSL using its Bulletin ID:

\$ esxupdate -b <Bulletin-ID> remove

To uninstall the CIM provider, follow the instructions above, but remove the fio-smis bulletin.

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



Uninstalling the Software in ESXi 4.1

To uninstall the ioMemory VSL package, run this command using the vCLI (from a remote machine):

1. Determine the bundle "bulletin" name:

vihostupdate --server <server-name> --query

Sample output:

```
-----Bulletin ID------
iomemory-vsl-2.2.0.7601742
-----Installed-----
2011-02-08T10:37:05
------Summary------
iomemory-vsl: block driver for ESXi 4.1.X
```

2. Remove the "bulletin" containing the driver & utilities:

vihostupdate --server <server-name> --remove --bulletin
iomemory-vsl-2.2.0.7601742

To uninstall the CIM provider, follow the instructions above, but remove the fio-smis bulletin.

Uninstalling the software in ESXi 5.0

To uninstall the ioMemory VSL package, run this command using the vCLI (from a remote machine):

1. Remove the VIB containing the driver & utilities:

```
esxcli --server <servername> software vib remove -n block-iomemory-vsl
```

To uninstall the CIM provider, run this command using the vCLI (from a remote machine):

1. Remove the VIB containing the driver & utilities:

esxcli --server <servername> software vib remove -n fio-smis



Uninstalling the SMI-S Scripts

Uninstall the fio-remote-util package according to your Linux distribution.

Unmanaged Shutdown Issues

Unmanaged shutdowns due to power loss or other circumstances can force the IBM High IOPS Adapter to perform a consistency check during the restart. This may take several minutes to complete.

Attention Check fio-status after a crash to see if the devices are in an "Attaching" state.

Although data written to the IBM High IOPS Adapter is not lost due to unmanaged shutdowns, important data structures may not have been properly committed to the device. This consistency check repairs these data structures.



Performance and Tuning

IBM High IOPS Adapters provide high bandwidth, and high Input/Output Operations per Second (IOPS), and are specifically designed to achieve low latency.

As IBM High IOPS Adapters improve in IOPS and low latency, the device performance may be limited by operating system settings and BIOS configuration. These settings may need to be tuned to take advantage of the revolutionary performance of IBM High IOPS Adapters.

While IBM 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

Dynamic Voltage and Frequency Scaling, or DVFS, are power management techniques that adjust the CPU voltage and/or 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 may benefit from reducing or disabling DVSF completely, even though this may increase power consumption.

DVFS, if available, should be configurable as part of your operating systems power management features as well as within your system's BIOS interface. Within the operating system and BIOS, DVFS features are often found under the Advanced Configuration and Power Interface (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, etc.) 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. When tuning for maximum performance you may benefit from limiting the C-states or turning them off completely, even though this may increase power consumption.

If your processor has ACPI C-states available, you can typically limit/disable them in the BIOS interface (sometimes referred to as a Setup Utility). APCI C-states may be part of of the Advanced Configuration and Power Interface

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



(ACPI) menu; consult your computer documentation for details.



Appendix A- Command-Line Utilities

These command-line utilities are only accessible through VMware's **Tech Support Mode** (also known as Shell/TSM) on ESXi and the Console Operating System (COS) on ESX.

- AttentionVMware suggests that the TSM only be used "for the purposes of troubleshooting and remediation."VMware recommends using the vSphere Client or any other VMware Administration Automation Product
to perform routine ESXi host configuration tasks that do not involve a troubleshooting scenario. For more
information visit VMware's Knowledge Base article on using this mode.
 - NOTE **SMI-S Management**: You may choose to use the SMI-S remote management tools instead of TSM command-line utilities. The SMI-S remote management tools provide a management experience similar to these command-line utilities. For more information, see <u>Appendix B- Using Python SMI-S Management Scripts</u>

The ioMemory VSL installation packages include various command-line utilities, installed by default to /usr/bin. These provide a number of useful ways to access, test, and manipulate your device.

Attention There are some additional utilities installed in the /usr/bin directory that are not listed below. Those additional utilities are dependencies (used by the main VSL utilities), and you should not use them directly unless Customer Support advises you to do so.

Utility	Purpose
fio-attach	Makes an IBM High IOPS Adapter available to the OS
fio-beacon	Lights the IBM High IOPS Adapter's external LEDs
fio-bugreport	Prepares a detailed report for use in troubleshooting problems
fio-detach	Temporarily removes an IBM High IOPS Adapter from OS access
fio-format	Used to perform a low-level format of an IBM High IOPS Adapter
fio-pci-check	Checks for errors on the PCI bus tree, specifically for IBM High IOPS Adapters. This utility is only supported on ESX.
fio-status	Displays information about the device
fio-update-iodrive	Updates the IBM High IOPS Adapter's firmware

NOTE There are -h (Help) and -v (Version) options for all of the utilities. Also, -h and -v cause the utility to exit after displaying the information.



fio-attach

NOTE The fio-attach utility requires that the ioMemory VSL is loaded and that the device is already detached. Refer to <u>fio-detach</u> for details on detaching a device.

Description

Attaches the IBM High IOPS Adapter and makes it available to the operating system. This creates a block device. You can then add it to ESX(i) as a storage area. The command displays a progress bar and percentage as it operates.

Syntax

fio-attach <device> [options]

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system. Use <u>fio-status</u> to view these.

Option	Description
-q	Quiet: disables the display of the progress bar and percentage.

Notes

If a device attaches, but the claiming process hangs, then one or more of the devices may not have been properly unclaimed when they were previously detached. The improperly unclaimed device(s) are preventing other devices from being claimed.

To solve this issue, attempt to attach each of the other devices individually. This will claim the device or devices that were improperly unclaimed and allow the hung device to proceed with attaching. You may then run fio-detach, if desired, on any devices to detach them again.

fio-beacon

Description

Lights the IBM High IOPS Adapter's three LEDs to locate the device. You should first detach the IBM High IOPS Adapter and then run fio-beacon. See <u>Common Maintenance Tasks</u> for best detach practices.

Syntax

fio-beacon <device> [options]

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system. This devices node is visible

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi

using <u>fio-status</u>.

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

fio-bugreport

Description

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

Syntax

fio-bugreport

Notes

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 send the output to <u>http://www.ibm.com/systems/support</u> for assistance in troubleshooting.

The output will indicate where the bugreport is saved.

Sample Output

```
~ # fio-bugreport
VMkernel-5.0.0
Report output: /var/tmp/fio-bugreport-20111006.223733-sc07HE.tar.gz
OS: VMware-ESXi-5.0.0
...
Building tar file...
Please attach the bugreport tar file
    /var/tmp/fio-bugreport-20111006.173256-sc07HE.tar.gz
    to your support case, including steps to reproduce the problem,
    and upload or email to customer support.
```



For example, the filename for a bug report file named fio-bugreport-20111006.173256-sc07HE.tar.gz indicates the following:

- Date (20111006)
- Time (173256, or 17:32:56)
- Misc. information (sc07HE.tar.gz)

fio-detach

Description

Detaches the IBM High IOPS Adapter. By default, the command displays a progress bar and percentage as it completes the detach.

Unmounting the Device

Attention Read Carefully: Read the following instructions carefully. Detaching a device while mounted, or under use, can cause errors, data loss and/or corruption.

In most cases, we do <u>not</u> recommend using the fio-detach utility to ensure that a device is detached. Instead, as a best practice, follow the instructions in the <u>Common Maintenance Tasks</u> section on disabling auto-attach as a safe detach workaround.

Syntax

fio-detach <device> [options]

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system.

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

Notes

Attempting to detach an IBM High IOPS Adapter may fail with an error indicating that the device is busy. This typically may occur if the IBM High IOPS Adapter is in use by VM or other process, or some process has the device open.



fio-format

Attention The fio-format utility requires that the ioMemory VSL be loaded with the IBM High IOPS Adapter(s) detached. Refer to <u>fio-detach</u> for details.

Description

Performs a low-level format of the device. By default, fio-format displays a progress-percentage indicator as it runs.

Attention Use this utility with care, as it deletes all user information on the device. You will be prompted as to whether you want to proceed with the format.

NOTE VMFS (VMware File System), the filesystem employed by ESX(i), requires 512 byte sector size.

Syntax

fio-format [options] <device>

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system.

Options	Description	
-b <size B K></size 	 Set the block (sector) size, in bytes or KiBytes (base 2). The default is 512 bytes. For example: -b 512B or -b 4K (B in 512B is optional). <u>Attention</u> ESX(i) only supports 512b sector sizes for use in VMFS datastores. Do not format your IBM High IOPS Adapter with any other sector size if you plan to use VMFS. If you are passing the device through to a VM (using VMDirectPathIO), then the guest VM can use any sector size appropriate for the the guest OS. In this case, formatting is done in the guest. 	
-f	Force the format size, bypassing normal checks and warnings. This option may be needed in rare situations when fio-format does not proceed properly. (The "Are you sure?" prompt still appears unless you use the -y option.)	
-d	Quiet mode: Disable the display of the progress-percentage indicator.	
-s <size M G T %></size 	 Set the device capacity as a specific size (in TB, GB, or MB) or as a percentage of the advertised capacity, for example: 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). 	

-0 <size B K M G T %></size 	Over-format the device size (to greater than the advertised capacity), where the maximum size equals the maximum physical capacity. If a percentage is used, it corresponds to the maximum physical capacity of the device. (Size is required for the -0 option; see the -s option above for size indicator descriptions.) <u>Attention</u> Before you use this option, please discuss your use case with Customer Support by contacting <u>http://www.ibm.com/systems/support</u> .
-у	Auto-answer "yes" to all queries from the application (bypass prompts).

You must re-attach the device in order to use the IBM High IOPS Adapter. See fio-attach for details.

fio-pci-check

Attention This utility is only supported on ESX. If it is run on an ESXi host, it will print the following message:

This utility is not supported on ESXi

Description

Checks for errors on the PCI bus tree, specifically for IBM High IOPS Adapters. This utility displays the current status of each IBM High IOPS Adapter. It also prints the standard PCI Express error information and resets the state.

NOTE It is perfectly 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.

Syntax

fio-pci-check [options]		
Description		
2> 1 = Disable the link; 0 = bring the link up (Not recommended)		
Scan every device in the system.		
Print the device serial number. This option is invalid when the ioMemory VSL is loaded.		
Force the link to retrain.		
Verbose: Print extra data about the hardware.		
•		

fio-status

Description

Provides detailed information about the installed devices. This script operates on /dev/fctX nodes. The utility depends on running as root and having the ioMemory VSL loaded.

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 of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system.

If <dev> is not specified, fio-status displays information for all IBM High IOPS Adapters in the system. If the ioMemory VSL is not loaded, this parameter is ignored.

Options	Description
-C	Count: Report only the number of IBM High IOPS Adapters 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.
-fk	Format key: Force alternate <"key=value"> format output, which may be easier for scripts and parsing programs to read.
-a	Print all available information for each device.

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

- Number and type of devices installed in the system
- ioMemory VSL version

Adapter information:

- Adapter type
- Product number
- External power status
- PCIe power limit threshold (if available)
- Connected IBM High IOPS Adapters

Block device information:

- Attach status
- Product name
- Product number
- Serial number

- PCIe address and slot
- Firmware version
- Size of the device, out of total capacity
- Internal temperature (average and maximum, since ioMemory VSL load) in degrees Celsius
- 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, 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, all information is printed, which includes the following information *in addition* to basic and data volume information:

Adapter information:

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

Block device information:

- Manufacturer's code
- Manufacturing date
- Vendor and sub-vendor information
- 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
- RAM usage

Error Mode Information: These include minimal mode, read-only mode, or write-reduced mode. If the ioMemory VSL is in minimal 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." is displayed.
- No media health information is displayed.

fio-update-iodrive

Description

Attention Your IBM High IOPS Adapters must be detached before running fio-update-iodrive. See <u>fio-detach</u> for details or <u>Common Maintenance Tasks</u> for information on disabling auto-attach.

Updates the IBM High IOPS Adapter's firmware. This utility scans the PCIe bus for all IBM High IOPS Adapters and updates them. A progress bar and percentage are shown for each device as the update completes.

To update one or more specific devices:

- Make sure the ioMemory VSL is loaded.
- Use the -d option with the device number.
- <u>Attention</u> It is extremely important that the power not be turned off during a firmware upgrade, as this could cause device failure. If a UPS is not already in place, consider adding one to the system prior to performing a firmware upgrade.
- Attention Note that when running multiple upgrades in sequence, it is critical to reboot the system after each upgrade. Otherwise the on-device format will not be changed, and there will be data loss.

Attention Do not use this utility to downgrade the IBM High IOPS Adapter to an earlier version of the firmware. Doing so may result in data loss and void your warranty.

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



- Attention The default action (without using the -d option) is to upgrade all IBM High IOPS Adapters 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.
- Attention Upgrade Path: There is a specific upgrade path that you must take when upgrading IBM High IOPS Adapter. Consult the *IBM ioMemory VSL Release Notes* for this ioMemory VSL release before upgrading IBM High IOPS Adapters.
 - NOTE If you receive an error message when updating the firmware that instructs you to update the midprom information, contact Customer Support.
 - NOTE When using VMDirectPathIO, if you upgrade the firmware on an IBM High IOPS Adapter, you must cycle the server power to have the change take place. Just restarting the virtual machine won't apply the change.

Syntax

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

where <firmware_version.fff> is the path and firmware archive file provided by IBM. The firmware archive path will depend on where it is located on the ESX(i) host. For example, you could transfer the archive to a folder on a datastore and then use the path to the file in that folder.

Options	Description
-d	Updates the specified devices (by fctx, where x is the number of the device shown in fio-status). If this option is not specified, all devices are updated. <u>Attention</u> Use the -d option with care, as updating the wrong IBM High IOPS Adapter could damage your device.
-f	Force upgrade (used primarily to downgrade to an earlier firmware version). <u>Attention</u> Use the -f option with care, as it could damage your card.
-1	List the firmware available in the archive.
-p	Pretend: Shows what updates would be done. However, the actual firmware is not modified.
-у	Confirm all warning messages.
-q	Runs the update process without displaying the progress bar or percentage.

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

If you arrived at this section from the <u>Upgrading the Firmware</u> section of the installation instructions, you should return to that section.



Appendix B- Using Python SMI-S Management Scripts

This section explains how to use Python SMI-S Management Scripts to manage IBM High IOPS Adapters on your ESX(i) host(s). Before you can run these scripts on a remote system, you must first install the SMI-S provider on your ESX(i) host system. For more information, see <u>Installing the Optional SMI-S Provider</u>.

NOTE We recommend using these Python management scripts to interface with the SMI-S provider, especially if you are unfamiliar with CIM browsers. If you are familiar with CIM models and browsers (such as YAWN), and you would like to use that method, continue to <u>Appendix C- Using a CIM Browser for SMI-S</u> <u>Management</u>

Our Python scripts interface with the SMI-S provider to create a management experience similar to running the command-line utilities on a local host.

NOTE If you're using Python 2.6 without patches, you may need to install the argparse package manually.

Installing Python WBEM Packages

The Python management scripts will run on a remote Linux machine with the proper packages installed.

- Download the fio-remote-util package from <u>http://www.ibm.com/support/entry/portal/docdisplay?Indocid=MIGR-5083174</u> to your remote system (the package is available in the Linux download folder for your Linux distribution).
 <u>Attention</u> The fio-remote-util-<version>.noarch.rpm (or .deb) package is available in the Linux download folders. Download this package from the folder for your Linux distribution (for the remote machine). In other words, if you plan to install these remote utilities on a Linux system running RHEL 5, go to that download folder to find this package.
- 2. Make sure the following required packages are installed (consult your distribution documentation for instructions on how to install these packages):
 - Python 2.6
 - Python argparse package
 - Python PyWBEM 0.7 package
 - VMware vCLI
- 3. Install the fio-remote-util package
 - This installs the Python scripts in /usr/bin.



Running the Python Management Scripts

The scripts run under Python, and therefore may require the command python before the script name, for example:

```
python fio-status.py --server <ip-address>:<port> --username <username>
--password <password> -a
```

NOTE When the scripts are installed using the .rpm or .deb packages, the scripts should be marked as executable in the system. This means that you shouldn't need to use the command python before the script name.

You can manually mark the scripts as executable. To mark the scripts as executable, run the following command:

chmod +x fio-*.py

With this change, you will not have to include python into the command.

Options

Each script has its own options that are specific to that script and its functions, which are listed in the following sections.

The **Remote Options**, however, are available for all scripts. The remote options are how you configure your scripts to connect to the ESX(i) host. These include the IP address and login credentials.

For more information on individual scripts and their options, see the reference section below.

Connection Issues

```
$ ./fio-status.py --server 10.10.10.110 -a
... python stack trace ...
pywbem.cim_operations.CIMError: (0, 'Socket error: [Errno 111] Connection
refused')
```

If the the connection is refused (see above example), check the following:

- Make sure the ESX(i) host is on and properly functioning.
- Make sure the SFCB server and SMI-S provider are installed and running on the ESX(i) host system(see Installing the SMI-S Provider on ESX(i) for installation details).
- Make sure the address / host name are correct, including port number (if needed).

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



Python Scripts Reference

The Python Management Scripts allow you to remotely manage your IBM High IOPS Adapters.

These scripts are available wherever you have installed them. As a best practice, we recommend that you install them in the /usr/bin Linux directory. This will allow you to run the scripts without defining the path.

Script	Purpose
fio-attach.py	Makes an IBM High IOPS Adapter available to the OS
fio-beacon.py	Lights the IBM High IOPS Adapter's external LEDs
fio-detach.py	Temporarily removes an IBM High IOPS Adapter from OS access
<u>fio-format.py</u>	Used to perform a low-level format of an IBM High IOPS Adapter
fio-status.py	Displays information about the device
fio-update-iodrive.py	Upgrades the firmware on the device

NOTE There are -h (Help) and -v (Version) options for all of the scripts.

fio-attach.py

- NOTE These Python management scripts run on a remote system. For information on installing these scripts and their dependent software, see Appendix B- Using Python SMI-S Management Scripts.
- Attention The fio-attach.py script requires that the ioMemory VSL be loaded, but with the IBM High IOPS Adapter(s) detached. See <u>fio-detach.py</u> for more information on detaching devices.
- <u>Attention</u> For full script functionality, including remounting the device, the vCLI **MUST** be installed on the same remote machine that you run the script on. You can test to see if the vCLI is properly installed by running:

esxcli --server <server> --username <user> --password <password>

Correct installation will result in a usage menu for the relevant ESX host version, otherwise an error message will print. vCLI 5.0 is recommended.

Description

Attaches the IBM High IOPS Adapter and makes it available to the operating system. This creates a block device. You can then add it to ESX(i) as a storage area. The script will return one of these three results:

- complete
- failed
- request timed out



Syntax Example

fio-attach.py [options] [remote options] <device>

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system. You can use <u>fio-status.py</u> to display the device node(s) of installed device(s).

Option	Descripti	ion
quiet	Quiet: dis	ables the display of progress.
Remote Options	Desc	cription
timeou	timeout Seconds to wait for request to complete (default:30; a timeout value of 0 means don't wait) NOTE If you have a larger device and/or multiple devices, this process might take more than 3 seconds to complete. If the script does time out, you can run <u>fio-status.py</u> to check the status of the device(s).	
server	Remo	ote host address/dns name (required). This address can include the port, for example: server 10.10.1.1:5989
userna	me Remo	ote user (default:anonymous)
passwo	rd Remo	ote user's password (default:none)
no-ssl	Disal	ble secure HTTP on connection

fio-beacon.py

NOTE These Python management scripts run on a remote system. For information on installing these scripts and their dependent software, see <u>Appendix B- Using Python SMI-S Management Scripts</u>.

Description

Lights the IBM High IOPS Adapter's three LEDs to locate the device. You should first detach the IBM High IOPS Adapter and then run fio-beacon.py.

Syntax Example

fio-beacon.py [options] [remote options] <device>

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system. You can use <u>fio-status.py</u> to display the device node(s) of installed device(s).

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi

Options	Descrit	otion		
options				
-0	Off: Iu	rns off the three LEDs.		
-1	On: Lig	In: Lights the three LEDs.		
-p	Prints th to stand	ne PCI bus ID of the device at <device> to standard output. Usage and error information may be written ard output rather than to standard error.</device>		
Option	Descr	tiption		
quie	Quiet	Quiet: disables the display of progress.		
Remote 0	Options	Description		
serve	er Remote host address/dns name (required). This address can include the port, for example:			
		server 10.10.1.1:5989		
useri	name	Remote user (default:anonymous)		
passi	word	Remote user's password (default:none)		
no-ss	Disable secure HTTP on connection			

fio-detach.py

- NOTE These Python management scripts run on a remote system. For information on installing these scripts and their dependent software, see <u>Appendix B- Using Python SMI-S Management Scripts</u>.
- Attention For full script functionality, including unmounting/unclaiming the device, the vCLI **MUST** be installed on the same remote machine that you run the script on. You can test to see if the vCLI is properly installed by running:

esxcli --server <server> --username <user> --password <password>

Correct installation will result in a usage menu for the relevant ESX host version, otherwise an error message will print. vCLI 5.0 is recommended.

Attention Detaching a device while mounted, or under use, can cause errors, data loss and/or corruption. Make sure the vCLI tools are properly installed to make sure the device is properly unmounted.

Description

Detaches the IBM High IOPS Adapter and removes the corresponding block device. The script will return one of these three results:

- complete
- failed



request timed out.

Syntax Example

fio-detach.py [options] [remote options] <device>

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system. You can use <u>fio-status.py</u> to display the device node(s) of installed device(s).

Options	Desc	ription
-d	Quiet	: Will not display progress.
Remote Options		Description
timed	out	Seconds to wait for request to complete (default:30; a timeout value of 0 means don't wait) NOTE If you have a larger device and/or multiple devices, this process might take more than 30 seconds to complete. If the script does time out, you can run <u>fio-status.py</u> to check the status of the device(s).
serve	er	Remote host address/dns name (required). This address can include the port, for example:server 10.10.1.1:5989
userr	name	Remote user (default:anonymous)
passi	word	Remote user's password (default:none)
no-sa	sl	Disable secure HTTP on connection

Notes

If the device continues to fail to detach, it may be because the IBM High IOPS Adapter is mounted, or some process has the device open.

fio-format.py

- NOTE These Python management scripts run on a remote system. For information on installing these scripts and their dependent software, see Appendix B- Using Python SMI-S Management Scripts.
- Attention The fio-format.py script requires that the ioMemory VSL be loaded with the IBM High IOPS Adapter(s) detached. Refer to <u>fio-detach.py</u> for details.

Description

Performs a low-level format of the board. The script will return one of these three results:

- complete
- failed
- request timed out.

Attention Use this utility with care, as it deletes all user information on the card. You will be prompted as to whether you want to proceed with the format.

NOTE VMFS, the default filesystem employed by ESX(i), requires 512 byte sector size.

Syntax Example

fio-format.py [options] [remote options] <device>

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system. You can use <u>fio-status.py</u> to display the device node(s) of installed device(s).

Options	Description
-b <size B K></size 	 Set the block (sector) size, in bytes or KiBytes (base 2). The default is 512 bytes. For example: -b 512B or -b 4K (B in 512B is optional). <u>Attention</u> ESX(i) only supports 512b sector sizes for use in VMDKs. Do not format your ioMemory device with any other sector size if you plan to use VMDKs. If you are passing the device through to a VM (using VMDirectPathIO), then the guest VM can use any sector size appropriate for the the guest OS. In this case, formatting is done in the guest.
-d	Quiet mode: Disable the display of the progress.
-s <size M G T %></size 	 Set the device capacity as a specific size (in TB, GB, or MB) or as a percentage of the advertised capacity, for example: 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).
-у	Auto-answer "yes" to all queries from the application (bypass prompts).
Remote Options	Description
timeout	 Seconds to wait for request to complete (default:60; a timeout value of 0 means don't wait) NOTE If you have a larger device, this process might take more than 60 seconds to complete. If the script does time out, you can run <u>fio-status.py</u> to check the status of the

device(s).

	IBM
server	Remote host address/dns name (required). This address can include the port, for example:server 10.10.1.1:5989
username	Remote user (default:anonymous)
password	Remote user's password (default:none)
no-ssl	Disable secure HT*TP on connection

You must re-attach the device in order to use the IBM High IOPS Adapter. See <u>fio-attach.py</u> for details.

fio-status.py

NOTE These Python management scripts run on a remote system. For information on installing these scripts and their dependent software, see <u>Appendix B- Using Python SMI-S Management Scripts</u>.

Description

Provides detailed information about the installed devices. This script operates on /dev/fctX nodes. The script depends on having the ioMemory VSL loaded.

Syntax

--password

fio-status.py [options] [remote options] [<device>]

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM High IOPS Adapter installed on the system.

If <device> is not specified, fio-status.py displays information for all cards in the system. If the ioMemory VSL is not loaded, this parameter is ignored.

Options	Descrip	otion
-C	Count:]	Report only the number of IBM High IOPS Adapters installed.
-a	Print all	available information for each device.
-1	List: ret	arns the output in a format that reflects the CIM class hierarchy.
Remote	Options	Description
serve	er	Remote host address/dns name (required). This address can include the port, for example:
		server 10.10.1.1:5989
useri	name	Remote user (default:anonymous)

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi

Remote user's password (default:none)



--no-ssl

Disable secure HTTP on connection

Output

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

- Number and type of devices installed in the system
- ioMemory VSL version

Adapter information:

- Adapter type
- Product number
- External power status
- PCIe power limit threshold (if available)
- Connected IBM High IOPS Adapters

Block device information:

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

fio-update-iodrive.py

If you arrived at this section from the <u>Upgrading the Firmware</u> section of the installation instructions, follow the link to return to that section.



Description

Attention Your IBM High IOPS Adapters must be detached before running fio-update-iodrive. See <u>fio-detach.py</u> for details or <u>Common Maintenance Tasks</u> for information on disabling auto-attach.

- <u>Attention</u> It is extremely important that the power not be turned off during a firmware upgrade, as this could cause device failure. If a UPS is not already in place, consider adding one to the system prior to performing a firmware upgrade.
- Attention Note that when running multiple upgrades in sequence, it is critical to reboot the system after each upgrade. Otherwise the on-device format will not be changed, and there will be data loss.
- Attention Do not use this utility to downgrade the IBM High IOPS Adapter to an earlier version of the firmware. Doing so may result in data loss and void your warranty.
- Attention Upgrade Path: There is a specific upgrade path that you must take when upgrading IBM High IOPS Adapter. Consult the *IBM ioMemory VSL Release Notes* for this ioMemory VSL release before upgrading IBM High IOPS Adapters.
 - NOTE If you receive an error message when updating the firmware that instructs you to update the midprom information, contact Customer Support.

Syntax

fio-update-iodrive.py -d <directory-path> -u <firmware-file.fff> [options]
[remote-options] <device>

Required	Parameters	Description		
-d <directory-path></directory-path>		Directory: Where <directory-path> the directory where the firmware file resides. You can transfer the file to a datastore. Example datastore path: /vmfs/volumes/datastorel/directory-name/</directory-path>		
-u <firmwa< td=""><td>are-file.fff></td><td>Use File: where <firmware-file.fff> is the firmware filename. E iodrive-firmware.fff</firmware-file.fff></td><td>xample:</td></firmwa<>	are-file.fff>	Use File: where <firmware-file.fff> is the firmware filename. E iodrive-firmware.fff</firmware-file.fff>	xample:	
<device< td=""><td>2></td><td>Where <device> is the name of the device node (/dev/fctx), when number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM H installed on the system. You can use <u>fio-status.py</u> to display the d installed device(s).</device></td><td>re x indicates the card High IOPS Adapter evice node(s) of</td></device<>	2>	Where <device> is the name of the device node (/dev/fctx), when number: 0, 1, 2, etc. For example, /dev/fct0 indicates the first IBM H installed on the system. You can use <u>fio-status.py</u> to display the d installed device(s).</device>	re x indicates the card High IOPS Adapter evice node(s) of	
Options	Description			
-f	Force upgrade (user <u>Attention</u> Use the	d primarily to downgrade to an earlier firmware version). -f option with care, as it could damage your card.		
-1	List the firmware av	vailable in the archive.		
-p	Pretend: Shows wh	at updates would be done. However, the actual firmware is not modified.		
-s	Show: Display the o	current device software version and exit.		
-n	No prompt, don't c	confirm before committing.		
-d	Runs the update pr	ocess without displaying the progress bar or percentage.		

-	-		-	-
-	-		-	_
-	-	-		_
=	-			
		_		

Remote Options	Description	
timeout	Seconds to wait for request to complete (a timeout value of 0 means don't wait). Default time is 30 minutes. NOTE If the script does time out, you can run <u>fio-status.py</u> to check the status of the device(s).	
server	Remote host address/dns name (required). This address can include the port, for example:server 10.10.1.1:5989	
username	Remote user (default:anonymous)	
password	Remote user's password (default:none)	
no-ssl	Disable secure HTTP on connection	

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

If you arrived at this section from the <u>Upgrading the Firmware</u> section of the installation instructions, you should return to that section.



Appendix C- Using a CIM Browser for SMI-S Management

This section outlines our SMI-S CIM model, including the instances and associations within that model. You can use this model along with a CIM browser to interface with the SMI-S provider installed on an ESX(i) host. This will allow you to manage your IBM High IOPS Adapter(s).

Before you can use a CIM browser to interface with the SMI-S provider, you must first install the SMI-S provider on your ESX(i) host system. For more information, see <u>Installing the SMI-S Provider on ESX(i)</u>.

Attention We recommend using our Python WBEM scripts to interface with the SMI-S provider instead of a CIM browser, especially if you are unfamiliar with CIM browsers. This section is meant for users who are versed in WBEM, SMI-S and DMTF standards. For more information on using the Python SMI-S management scripts, see <u>Appendix B- Using Python SMI-S Management Scripts</u>.

SMI-S Interface Background

The SMI-S interface is based on Web-Based Enterprise Management (WBEM) and provides a Common Information Model (CIM) model that represents the IBM High IOPS Adapter and associated software, in accordance with existing Distributed Management Task Force (DMTF), Storage Networking Industry Association (SNIA), and Storage Management Initiative Specification (SMI-S) standards. This model permits backward-compatible extension, accommodating new hardware and software features developed by IBM.

References

CIM Schema v2.26 http://www.dmtf.org/standards/cim/cim_schema_v2260

DMTF DSP1011, Physical Asset Profile http://www.dmtf.org/standards/published_documents/DSP1011_1.0.2.pdf

DMTF DSP1023, Software Inventory Profile http://www.dmtf.org/standards/published_documents/DSP1023_1.0.1.pdf

DMTF DSP1033, Profile Registration http://www.dmtf.org/standards/published_documents/DSP1033_1.0.0.pdf

DMTF DSP1075 PCI Device Profile http://www.dmtf.org/standards/published_documents/DSP1075_1.0.0.pdf



DMTF DSP1002, Diagnostics Profile http://www.dmtf.org/standards/published_documents/DSP1002_2.0.0.pdf

SMI-S v1.4 Architecture http://www.snia.org/sites/default/files/SMI-Sv1.4r6_Architecture.book_.pdf

SMI-S v1.4 Common Profiles http://www.snia.org/sites/default/files/SMI-Sv1.4r6_CommonProfiles.book_.pdf

SMI-S v1.4 Host Profiles http://www.snia.org/sites/default/files/SMI-Sv1.4r6_Host.book_.pdf

SMI-S v1.4 Common Diagnostic Model http://www.dmtf.org/standards/mgmt/cdm/

Description

SMI-S is a collection of specifications that traditionally focus on Storage Area Network (SAN) systems based on the SCSI command set, such as Fibre Channel, iSCSI, and SAS. However, the general pattern used to model these storage systems can be applied to solid-state, direct-attached storage systems such as those provided by IBM.

IBM High IOPS Adapters are modeled using the SMI-S patterns established in the Storage HBA, Direct Attached (DA) Ports, and Host Discovered Resources Profiles. The physical aspects of the IBM High IOPS Adapter and all firmware and ioMemory VSL software are modeled using published DMTF specifications, including the Physical Asset, Software Inventory, PCI Device Profiles, and Common Diagnostic Model Profile.

The following chart describes the IBM SMI-S CIM model, with IBM High IOPS Adapters and their associated firmware and software. For simplicity, the prefix FIO_ has been removed from the class names.





A: IOMemoryPort Class

The central instance of the model is of the IOMemoryPort class (A in the figure), a logical representation of the IBM High IOPS Adapter. It supports the extrinsic methods necessary to provision the drive. An instance of PCIDevice (B) and IOMemoryPort exist for each installed IBM High IOPS Adapter, and they are associated with instances of ConcreteIdentity (1). An instance of SSDStatistics (C), which contains important performance and capacity data for the device, is associated by an ElementStatisticalData association (2) to each IOMemoryPort. IOMemoryPort is scoped by an instance of the ComputerSystem class. The SystemDevice (3) aggregation aggregates IOMemoryPort within the containing ComputerSystem.

E: IOMemoryPortController Class

An instance of IOMemoryPortController (E) represents the ioMemory VSL used to control the installed IBM High IOPS Adapters. IOMemoryPortController specializes CIM_PortController, and it aggregates IoMemoryPort with the ControlledBy (4) aggregation. The software version and vendor information are represented by the SoftwareIdentity (F) instance that is associated to IOMemoryPortController (E) via ElementSoftwareIdentity (5). The SoftwareIdentity that represents the installed ioMemory VSL software is associated to the scoping ComputerSystem using the InstalledSoftwareIdentity association (6).

An instance of the ProtocolEndpoint class (G) represents both ends of the logical data path between the IOMemoryPort and the solid-state storage. This aspect of the model is derived from the pattern in the DA Ports Profile, where the port is both an initiator and target. ProtocolEndpoint is associated to the IOMemoryPort by DeviceSAPImplementation (7) and to the ComputerSystem by HostedAccessPoint (8).



H: LogicalSSD Class (Block Device)

The block device exposed to applications (file systems, database, and logical volume manager) is modeled using an instance of LogicalSSD (H), a subclass of CIM_DiskDrive. It is associated with a StorageExtent (J) using the MediaPresent association (9), but the StorageExtent will always be present. It is also associated to the ProtocolEndpoint (G) representing the IOMemoryPort using SAPAvailableForElement (10) and to the scoping ComputerSystem using SystemDevice (3).

IBM High IOPS Adapters, being PCIe devices, are also represented by an instance of the PCIDevice class (B). IOMemoryPort is an alternate representation of the PCIDevice and its associated control device. It is associated to it by the ConcreteIdentity association.

K: SoftwareIdentity

The ioMemory VSL software is also represented with SoftwareIdentity, which is associated to the PCIDevice by the ElementSoftwareIdentity association (11). The SoftwareIdentity (firmware) is associated to the scoping ComputerSystem by the InstalledSoftwareIdentity association (12). An instance of SoftwareInstallationService (L) is associated with each PCIDevice, which can be used to update device firmware.

M: Physical Aspects

The physical aspects of IBM High IOPS Adapters are represented by an instance of the PhysicalPackage class (M), which is associated to the PCIDevice by Realizes (13) and to the scoping ComputerSystem by SystemPackaging (14). The temperature sensors on IBM High IOPS Adapters are represented by an instance of TemperatureSensor (N) and is associated to the PhysicalPackage by AssociatedSensor.

Implementation

This section describes the arrangement of instances and associations for the IBM device CIM model. Not all class properties are described in detail. Consult the CIM schema for detailed description of all properties.

A WBEM CIM provider based on this model will be developed in the future. IBM intends to support popular CIMOMs, including OpenPegasus, OpenWBEM, SFCB, and Windows WMI.

The device health is indicated by the value of the HealthLevel property. Values include: Healthy, Warning, Reduced Write, and Read Only. These values are mapped to standardHealthState values – OK, Degraded/Warning, and Critical Failure – as appropriate.

Extrinsic methods for device provisioning include attach, detach, format, and update. The attach method creates a block device for the IBM High IOPS Adapter. Detach disables the block device. A format option enables users to specify the device size in either megabytes or a percentage. The update method allows users to upgrade the firmware on the device.

Device longevity is indicated by the value of the HealthPercentage property. FlashbackAvailability indicates whether or not this feature of the IBM High IOPS Adapter is online.

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



IOMemoryPorts are aggregated by IOMemoryPortController via the ControlledBy aggregation. Instances of IOMemoryPort are associated to their corresponding PCIDevice with the ConcreteIdentity association. The IOMemoryPort is a logical device of the scoping ComputerSystem and is indicated as such by the SystemDevice aggregation.

Products with two or more IBM High IOPS Adapters, such as the IBM High IOPS Duo Adapter do appear like two separate IBM High IOPS Adapters. For products with multiple devices, the IOMemoryPort class is extended to include information about the carrier card type, serial number, and external power connection for the product as a whole.

IOMemoryPort

One instance of IOMemoryPort exists for each IBM High IOPS Adapter installed in the ComputerSystem.

The LocationIndicator property reflects the state of the device indicator beacon (e.g., all LEDs on solid). Reading the value gives the current state of the indicator. Writing the value with "On" or "Off" turns the indicator on or off and can be used to determine the device's physical location.

SSDStatistics

One instance of SSDStatistics exists for each IOMemoryPort instance. Properties of this object provide performance and capacity information. Some of this information is only available when the drive is attached (i.e., the state of the associated IOMemoryPort is "Attached").

IOMemoryPortController

Only one instance of IOMemoryPortContoller exists, representing the ioMemory VSL software used to control IOMemoryPorts. The IOMemoryPortController specializes the CIM_PortController.

IOMemoryPortController is aggregated to the scoping ComputerSystem using the SystemDevice aggregation. IOMemoryPortController is associated with a SoftwareInventory instance representing the ioMemory VSL software properties via the ElementSoftwareIdentity association.

ProtocolEndpoint

One instance of ProtocolEndpoint exists for each instance of IOMemoryPort. It is associated to the IOMemoryPort using DeviceSAPImplementation and to LogicalSSD using SAPAvailableForElement. Because an IOMemoryPort represents both the initiator and target ports, only one ProtocolEndpoint per IOMemoryPort is needed to model the connection between IOMemoryPort and LogicalSSD.

LogicalSSD

One instance of LogicalSSD, a subclass of CIM_DiskDrive, exists for each block device (/dev/fioX) exposed by an IBM High IOPS Adapter. Correlatable IDs are used, based on operating system device names. This enables client applications to associate block devices discovered through this model with resources discovered from other SMI-S models instrumented on the host system.

ComputerSystem aggregates LogicalSSDs via SystemDevice. The LogicalSSD instances are associated to their

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



ProtocolEndpoints via SAPAvailableForElement. If the IOMemoryPort associated to the endpoint is not attached, then the Availability property is set to "Off Line," and the DeviceID property value is "Unknown."

StorageExtent

One instance of StorageExtent is associated with each LogicalSSD and represents the logical storage of the associated device.

SoftwareIdentity

One instance of SoftwareIdentity exists to represent the ioMemory VSL software. The firmware is also modeled using SoftwareIdentity but requires an instance for each IBM High IOPS Adapter installed. The IsEntity property has a value of True, indicating that the SoftwareIdentity instance corresponds to a discrete copy of the ioMemory VSL software or firmware. The MajorVersion, MinorVersion, RevisionNumber, and BuildNumber properties convey the driver/firmware version information. The Manufacturer property can be used to identify Fusion-io.

Another option for the firmware is to omit the InstalledSoftwareIdentity association with ComputerSystem, because the firmware is not really installed on ComputerSystem. This option would depend on how users want to model the firmware.

SoftwareInstallationService

An instance of SoftwareInstallationService exists for each PCIDevice and can be used to update the associated device's firmware.

PCIDevice

An instance of PCIDevice is instantiated for each IBM High IOPS Adapter (PCIe card) in the computer. Properties are set as follows:

- BusNumber bus number where the PCIe device exists
- DeviceNumber device number assigned to the PCI device for this bus.
- FunctionNumber set to the function number for the PCI device.
- SubsystemID, SubsystemVendorID, PCIDeviceID, VendorID, and RevisionID are optional but can be populated if values can be extracted from the configuration registers of the PCI device.

PCIDevice is associated with IOMemoryPort, its alternate logical representation, using ConcreteIdentity. The PCIDevice is also associated with PhysicalPackage, representing the physical aspects of the IBM High IOPS Adapter, via Realizes.



PhysicalPackage

One instance of PhysicalPackage exists for each discrete, physical IBM High IOPS Adapter installed in the computer system. The Manufacturer, Model, SKU, SerialNumber, Version, and PartNumber properties can be used to describe these aspects of the physical card. PhysicalPackage is associated with PCIDevice via Realizes and the scoping ComputerSystem via SystemPackaging.

TemperatureSensor

One instance of TemperatureSensor exists for each PhysicalPackage. Temperature information for the drive is stored in the properties of this object.

Diagnostic Test

One instance of DiagnosticTest will exist. The RunDiagnostic() method will trigger a snapshot of device status for the specified ManagedElement which must be an instance of IoMemoryPort. The diagnostic run is synchronous and runs instantaneously. The resulting ConcreteJob object will associate to the originating DiagnosticTest instance and the respective IoMemoryPort instance that was specified (see Figure 2). At this time, RunDiagnostic() can only be used with the default DiagnosticSettingData provided.

Each run will add a single entry of DiagnosticSettingDataRecord and associated DiagnosticCompletionRecord in the DiagnosticLog. The RecordData property of the DiagnosticCompletionRecord will record critical device status at the time of the run. The format of the RecordData string can be found in the RecordFormat property.

The format is a series of status strings, each of which can hold one of the following values delimited by an asterisk (*) character: "Unknown", "OK", "Warning", or "Error". Currently, seven status values are recorded: WearoutStatus, WritabilityStatus, FlashbackStatus, TemperatureStatus, MinimalModeStatus, PciStatus and InternalErrorStatus. All of these should report "OK" under normal operating conditions.

WearoutStatus will be set to "Warning" when less than 10% reserve space is left on the device. It will be set to "Error" when there is no more reserved space.

- WritabilityStatus will be set to "Error" whenever the device is write throttling or in read-only mode. This can happen due to a variety of conditions including device wearout and insufficient power.
- FlashbackStatus will report "Warning" if a catastrophic error causes Flashback protection to be degraded.
- TemperatureStatus will report "Warning" when the device temperature is nearing the maximum safe temperature and "Error" when the maximum safe temperature is reached or surpassed.
- MinimalModeStatus will report either "Warning" or "Error" whenever the device is in minimal mode.
- PciStatus will report "Warning" or "Error" if there are compatibility problems with the host PCIe bus.
- InternalErrorStatus will report "Error" if there are any internal problems with the ioMemory VSL.

The CompletionState property will summarize the results and may be set to Unknown, OK, Warning or Failed. If any status is in error the state will report as Failed. Otherwise, if there is any warning status the state will report Warning. The Message property will be set to indicate the appropriate action if there are any warnings or errors.

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



DiagnosticSetting Data

There will be an instance of DiagnosticSettingData associated with the DiagnosticTest instance (see Figure 2). It records the default settings for each call to RunDiagnostic.

DiagnosticServiceCapabilities

There is an instance of DiagnosticServiceCapabilities associated with the DiagnosticTest instance which records the capabilities of the DiagnosticTest service.

DiagnosticLog

An instance of DiagnosticLog is associated with the DiagnosticTest instance and will store the results of each run.

DiagnosticSettingRecord

A copy of the default DiagnosticSettingData will be stored in a DiagnosticSettingDataRecord each time a diagnostic is run and will be associated with an instance of DiagnosticCompletionRecord.

DiagnosticCompletionRecord

An instance of DiagnosticCompletionRecord will store the results of each RunDiagnostic execution. The details are explained in DiagnosticTest.

RegisteredDiskDriveLiteProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Disk Drive Lite Profile. The following properties are set as follows:

- InstanceID set to "SNIA:DiskDriveLiteProfile-1.4.0"
- RegisteredOrganization set to "11" (SNIA)
- RegisteredName set to "DirectAccess Ports Profile"
- RegisteredVersion set to "1.4.0"

RegisteredDAPortsProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the DA Ports Profile. The properties are set as follows:

- InstanceID set to "SNIA:DAPortsProfile-1.4.0"
- RegisteredOrganization set to "11" (SNIA)
- RegisteredName set to "DirectAccess Ports Profile"



• RegisteredVersion - set to "1.4.0"

RegisteredStorageHBAProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Storage HBA Profile. The properties are set as follows:

- InstanceID set to "SNIA: StorageHBAProfile-1.4.0"
- RegisteredOrganization set to "11" (SNIA)
- RegisteredName set to "Storage HBA Profile"
- RegisteredVersion set to "1.4.0"

RegisteredHostDiscoveredResourcesProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Host Discovered Resources Profile. The properties are set as follows:

- InstanceID set to "SNIA: HostDiscoveredResourcesProfile-1.2.0"
- RegisteredOrganization set to "11" (SNIA)
- RegisteredName set to "Host Discovered Resources Profile"
- RegisteredVersion set to "1.2.0"

RegisteredPCIDeviceProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the PCI Device Profile. The properties are set as follows:

- InstanceID set to "DMTF:DSP1075-PCIDevice-1.0.0a"
- RegisteredOrganization set to "2" (DMTF)
- RegisteredName set to "PCIDevice Profile"
- RegisteredVersion set to "1.0.0a"

RegisteredSoftwareInventoryProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Software Inventory Profile. The properties are set as follows:

- InstanceID set to "DMTF:DSP1023-SoftwareInventory-1.0.1"
- RegisteredOrganization set to "2" (DMTF)
- RegisteredName set to "Software Inventory Profile"



• RegisteredVersion - set to "1.0.1"

RegisteredSoftwareUpdateProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Software Update Profile. The properties are set as follows:

- InstanceID set to "DMTF:DSP1023-SoftwareUpdate-1.0.0"
- RegisteredOrganization set to "2" (DMTF)
- RegisteredName set to "Software Update Profile"
- RegisteredVersion set to "1.0.0"

RegisteredPhysicalAssetProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Physical Asset Profile. The properties are set as follows:

- InstanceID set to "DMTF: PhysicalAssetProfile-1.0.2"
- RegisteredOrganization set to "2" (DMTF)
- RegisteredName set to "PhysicalAsset Profile"
- RegisteredVersion set to "1.0.2"

RegisteredSensorsProfile

Only one instance of this class is needed. It resides in the /root/interop namespace and indicates the implementation of the Sensors Profile. The properties are set as follows:

- InstanceID set to "SNIA: SensorsProfile-1.0.0"
- RegisteredOrganization set to "11" (SNIA)
- RegisteredName set to "Sensors Profile"
- RegisteredVersion set to "1.0.0"

RegisteredCommonDiagnosticProfile

Only one instance of this class is needed. It will reside in the /root/interop namespace and indicate the implementation of the Common Diagnostic Model Profile. The InstanceID property will be set to a value of "DMTF:DiagnosticsProfile-2.0.0a". The RegisteredOrganization property will be set to a value of "2" (DMTF). The RegisteredName property will be set to a value of "Diagnostics Profile". The RegisteredVersion property will be set to a value of "2.0.0a".



Indications

An indication will be generated periodically when a serious condition exists for a particular IBM High IOPS Adapter. The WBEM provider currently supports six types of indications. They alert users of the SMI-S provider to conditions such as imminent wearout, degradation of writability, degradation of the flashback feature, higher temperature, and internal error states.

The indications will be instances of the FIO_AlertIndication class which simply specializes the CIM_AlertIndication class.

The values for the properties of the FIO_AlertIndication instances are under development and may change as testing proceeds and feedback is received. Additional indication types may also be added as necessary.



FIO_AlertIndication

Property	Value
IndicationIdentifier	See below for each type
IndicationTime	Timestamp when sent
AlertingManagedElement	IoMemoryPort.DeviceID= <device id=""></device>
AlertingElementFormat	CIMObjectPath (2)
AlertType	Device Alert (5)
PerceivedSeverity	See below for each type
ProbableCause	See below for each type
SystemCreationClassName	"FIO_AlertIndication"
SystemName	<hostname></hostname>
ProviderName	"fiosmis"
CorrelatedIndications	Not used
Description	Class description
OtherAlertType	Not used
OtherSeverity	Not used
ProbableCauseDescription	Not used
EventID	Same as IndicationIdentifier
OwningEntity	<vendor></vendor>
MessageID	TBD
Message	TBD
MessageArguments	TBD

Reduced Writability Indication

The ioMemory VSL can dramatically reduce write throughput to manage device conditions such as excessive wear, high temperature, and insufficient power. The reduced writability indication is generated while the drive is in this mode. If the triggering condition is excessive wear, the IoMemoryPort health percentage will report 0% health.

Property	Value
IndicationIdentifier	<mfr>":"<hostname>":write"</hostname></mfr>
PerceivedSeverity	Degraded/Warning (3)



Read-only Indication

When the drive has reached the end-of-life, it can no longer be written to and can only be read from. The read-only indication will be sent when this occurs. The IoMemoryPort health percentage will continue to report 0% health when this happens.

Property	Value
IndicationIdentifier	<mfr>":"<hostname>":read_only"</hostname></mfr>
PerceivedSeverity	Degraded/Warning (3)
ProbableCause	Threshold Crossed (52)

Wearout Indication

As the drive wears out, this indication is generated as a warning when the drive health percentage drops below 10%, before write throughput is reduced.

Property	Value
IndicationIdentifier	<mfr>":"<hostname>":wearout"</hostname></mfr>
PerceivedSeverity	Degraded/Warning (3)
ProbableCause	Threshold Crossed (52)

Flashback Indication

If a catastrophic part failure degrades the effectiveness of the flashback feature, this indication will be sent.

Property	Value
IndicationIdentifier	<mfr>":"<hostname>":flashback"</hostname></mfr>
PerceivedSeverity	Degraded/Warning (3)
ProbableCause	Loss of Redundancy (88)

High Temperature Indication

This indication will be sent when the temperature of the card becomes excessive.

Property	Value
IndicationIdentifier	<pre><mfr>":"<hostname>":temperature"</hostname></mfr></pre>
PerceivedSeverity	Critical (6)



ProbableCause	
---------------	--

Temperature Unacceptable (51)

Error Indication

If the ioMemory VSL is in an error state the error indication will be sent.

Property	Value
IndicationIdentifier	<mfr>":"<hostname>":error"</hostname></mfr>
PerceivedSeverity	Major (6)
ProbableCause	Other (1)





Appendix D- Monitoring the Health of ioMemory Devices

This section describes how the health of IBM High IOPS Adapters can be measured and monitored in order to safeguard data and prolong device lifetime.

NAND Flash and Component Failure

An IBM High IOPS Adapter is a highly fault-tolerant storage subsystem that provides many levels of protection against component failure and the loss nature of solid-state storage. As in all storage subsystems, component failures may occur.

By proactively monitoring device age and health, you can ensure reliable performance over the intended product life.

Health Metrics

The ioMemory VSL manages block retirement using pre-determined retirement thresholds. fio-status or fio-status.py shows 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. See the <u>Health Monitoring Techniques</u> section below for methods for capturing this alarm event.

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

After the 0% threshold, the device will soon enter *read-only* mode – any attempt to write to the IBM High IOPS Adapter causes an error. Some filesystems may require special mount options in order to mount a read-only block device in addition to 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 filesystem to not try and replay the journal.

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

The IBM High IOPS Adapter may 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 wearout.

- NOTE For service or warranty-related questions, contact the company from which you purchased the device.
- NOTE For products with multiple IBM High IOPS Adapters, these modes are maintained independently for each device.

Health Monitoring Techniques

fio-status: Output from the fio-status utility shows the health percentage and device state. These items are referenced as "Media status" in the sample output below.

```
Found 1 ioDrive in this system
Fusion-io driver version: 2.2.3 build 240
Adapter: ioDrive
Fusion-io ioDrive 160GB, Product Number:FS1-002-161-ES
...
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
```

SNMP/SMI-S: On Windows or Linux, see the corresponding appendix for details on how to configure SNMP or SMI-S health indicators.

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

- Healthy
- Read-only
- Reduced-write
- Unknown



Appendix E- Using Module Parameters

The following table describes the module parameters you can set using the esxcfg-module command.

<u>Attention</u> The the remote option (--server) is only required for the vCLI. Sample Command:

esxcfg-module --server <server-name> iomemory-vsl -s '<parameter>=<value>'

NOTE You must reboot the ESX(i) system to enforce any parameter changes.

Module Parameter	Default Value (min/max)	Description
auto_attach	1	1 = Always attach the device on driver load.0 = Don't attach the device on driver load.
force_minimal_mode	0	1 = Force minimal mode on the device.0 = Do not force minimal mode on the device.
parallel_attach	1	1 = Enable parallel attach of multiple devices.0 = Disable parallel attach of multiple devices.
tintr_hw_wait	0 (0, 255)	Interval (microseconds) to wait between hardware interrupts. Also known as interrupt coalescing. 0 is off.

NOTE Module parameters are global—they apply to all IBM High IOPS Adapters in the computer.

To query the current module parameters, use the following command:

esxcfg-module --server <server-name> iomemory-vsl -g



Appendix F- Working with ioMemory Devices and VMDirectPathIO

Each IBM High IOPS Adapter can either be used as a VMFS datastore in ESX(i), or they can be passed through directly to a virtual machine. In VMware documentation this is often referred to as VMDirectPathIO and is commonly referred to as PCI passthrough.

- Attention If you are passing the device(s) through, you do not need to install the ioMemory VSL on the ESX(i) system. Instead, install the software on the guest system. Only install the driver if you plan on creating a VMFS on the device(s).
 - NOTE When using VMDirectPathIO, if you upgrade the firmware on an IBM High IOPS Adapter, you must cycle the power to have the change take place. Just restarting the virtual machine won't apply the change.

Using Products with Multiple Devices

Some products contain multiple devices, such as the IBM High IOPS Duo Adapter. The ioMemory VSL does not support splitting the modules between two functions or virtual machines. The following scenarios are supported:

- Both IBM High IOPS Adapters are used as a VMFS datastore in ESX(i).
- Both IBM High IOPS Adapters are passed through to the same virtual machine (using VMDirectPathIO).



Appendix G- Documentation Permissions

The AVR bootloader and the tree.h file, which ship in binary form with the driver, contain the following documentation copyright requirements.

AVR Bootloader

• Copyright © 2002, 2003, 2004, 2005, 2006, 2007 Eric B. Weddington

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holders nor the names of contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. *

tree.h

Copyright 2002 Niels Provos <provos@citi.umich.edu> All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi



following conditions are met:

- 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



IBM Support

IBM High IOPS Adapter software and documentation are available on the web at the following address:

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5083174.

IBM part number 81Y1046

IBM ioMemory VSL 2.3.10 User Guide for VMware ESX and ESXi