IBM ioMemory VSL 3.2.6

USER GUIDE FOR VMWARE ESX AND ESXI

DECEMBER 11, 2013

IBM

Table of Contents

IBM ioMemory VSL 3.2.6 User Guide for VMware ESX and ESXi

Legal Notices

Part Number: D0004698-004_1
Published: December 11, 2013

Introduction

Overview

Congratulations on your purchase of an IBM solid-state storage device. This guide explains how to install, troubleshoot, and maintain the software for your IBM High IOPS Adapters.

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 Second Generation 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 an IBM High IOPS Duo Adapter, 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.

NOTE-

References to "ESX(i)" refers to ESX or ESXi.

About the IBM High IOPS Platform

The IBM High IOPS platform combines ioMemory VSL software (VSL stand for Virtual Storage Layer) with IBM High IOPS hardware to take enterprise applications and databases to the next level.

Performance

The IBM High IOPS platform provides consistent microsecond latency access for mixed workloads, multiple gigabytes per second access and hundreds of thousands of IOPS from a single product. The sophisticated IBM High IOPS architecture allows for nearly symmetrical read and write performance with best-in-class low queue depth performance, making the IBM High IOPS platform ideal across a wide variety of real world, high-performance enterprise environments.

The IBM High IOPS platform integrates with host system CPUs as flash memory to give multiple (and mostly idle) processor cores, direct and parallel access to the flash. The platform's cut-through architecture gives systems more work per unit of processing, and continues to deliver performance increases as CPU power increases.

Endurance

The IBM High IOPS platform offers best-in-class endurance in all capacities, which is crucial for caching and write-heavy databases and applications.

Reliability

The IBM High IOPS platform eliminates concerns about reliability like NAND failures and excessive wear. The all-new intelligent, self-healing feature called Adaptive Flashback provides complete, chip-level fault tolerance. Adaptive Flashback technology enables an IBM High IOPS product to repair itself after a single chip or a multi-chip failure without interrupting business continuity.

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 Adapter Hardware Installation Guide* for requirements on the following:
  - PCIe Slot
  - Cooling
  - Power
- **Supported Devices**: Also see the *IBM High IOPS Adapter Hardware Installation Guide* for a list of supported IBM High IOPS Adapters.
- **RAM Requirements**: The *IBM ioMemory VSL Release Notes* contains memory (RAM) requirements for this version of the software.

For specific IBM High IOPS System x server configuration information and requirements, refer to the following URL: http://www.ibm.com/support/entry/portal/docdisplay?lndocid=SERV-IOPS

Supported VMware Operating Systems

- ESX 4.0 Update 4
- ESX 4.1 Update 3
- ESXi 4.0 Update 4
- ESXi 4.1 Update 3
- ESXi 5.0 Update 3*
- ESXi 5.1 Update 1*
- ESXi 5.1*
- ESXi 5.5**

* The `scsi` version of the ioMemory VSL software is only supported on these ESXi hypervisor versions.

** ESXi 5.5 is only supported with the `scsi` version of the ioMemory VSL software.

Upgrading Legacy Adapters (IMPORTANT)

Please read these IBM High IOPS Adapter compatibility considerations.

Multiple High IOPS adapters are installed in a single system:

When multiple High IOPS Adapters are installed in the same server, all devices must operate with the same version of software. High IOPS adapters require matching firmware, drivers and utilities. This is a very important

consideration when adding a new Second Generation High IOPS Adapter in a server where Legacy Adapters are deployed.

When Upgrading Legacy Adapters operating with a previous generation of software (1.2.x or v2.x), you must back up the data on the adapter before upgrading to prevent data loss. After upgrading the ioMemory VSL to version 3.x, the legacy adapters will not logically attach to the system until the firmware is also updated. Detailed instructions for upgrading software is provided in Upgrading Devices from VSL 2.x to 3.x on page 54.

Upgrading from version 1.2.x or 2.x software to 3.x:

Upgrading Legacy adapters from 1.2.x software to version 3.1.1 offers a number of significant changes and improvements, however there are some important considerations

When performing an upgrade from 1.2.x to 3.x, you must perform a staged upgrade (upgrade to the 2.x software and firmware before upgrading to 3.x). The device driver name has also changed from fio-driver (version 1.2.x) to iomemory-vsl (2.x and above).

The upgrade process from 2.x to 3.x will require the adapter to be formatted. Formatting will remove all existing data from the card and the data must be restored after the update completes. Users must back up their data before proceeding with the upgrade process to version 3.x.

The firmware upgrade process updates and modifies important hardware settings that are not compatible with 1.2.x or 2.2.3 versions of software. Once updated, the card cannot be black-leveled to the previous versions of software. Please see the "change history" documentation for a complete list of new features, enhancements, and fixes.

Replacing a failed legacy High IOPS card and "mandatory" update requirements:

As the supply of legacy adapters diminishes from inventory, it becomes more likely that warranty replacement cards will transition to the newer versions of the High IOPS adapters. Replacement High IOPS cards and may require firmware updates to support the new or existing cards in the server.

Any situation when mixing the flash NAND technology occurs, the minimum version of software supported by the latest generation of hardware prevails. A mandatory upgrade of software is required to support the latest generation of hardware with backward compatibility to legacy cards in the server.

Change History's Update Recommendations:

Change histories files provide an ongoing list of changes to a series of software compatible with a family of hardware. Please review the change histories using the following guidelines as to how IBM recommends or suggests updates to code levels at the website below:

http://www.ibm.com/support/entry/portal/docdisplay?brand=5000008&lndocid=HELP-FIX

Software Installation

Attention!

Every IBM High IOPS Adapter in a system must be upgraded to the appropriate firmware.

For example, if you have a system running ioMemory VSL software version 2.2.3 with IBM High IOPS Adapters previously installed, and you want to install new Second Generation IBM High IOPS Adapters (that require the latest version of the firmware), then you will need to upgrade all of the existing devices with firmware that supports this version of the ioMemory VSL software. Follow the upgrade path in the *IBM ioMemory VSL Release Notes* to determine the upgrade sequence.

Attention!

Upgrade Previous Devices First

If you have IBM High IOPS Adapters configured for ioMemory VSL software version 2.x or earlier, you must upgrade their firmware before installing new devices in the system. See Upgrading Devices from VSL 2.x to 3.x on page 54 for the upgrade instructions.

Attention!

IBM High IOPS Adapters cannot be installed as part of an ESX(i) 4.x installation.

VMDirectPathIO

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 software on the ESX(i) system. Instead, install the ioMemory VSL software on the guest operating system. For example, you would pass the device through to a Windows VM and then install the Windows ioMemory VSL software on that VM (consult the *IBM ioMemory VSL User Guidefor Windows* for installation and user instructions).

There are special considerations when passing through an IBM High IOPS Adapter, for more information, see Working with ioMemory Devices and VMDirectPathIO on page 53 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 software. In order to manage the IBM High IOPS Adapters and ioMemory VSL software, 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

VMware provides the **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 http://www.vmware.com/support/developer/vcli/

Attention!

We do not recommend using the vCLI within a virtual machine that is hosted on your ESXi system. The ioMemory VSL software installation and configuration processes involve 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 software. 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 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: http://kb.vmware.com/kb/1017910.

Installation Overview

1. Download the latest version of the software at http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**).

2. If you are installing this version of ioMemory VSL software on a system with Legacy IBM High IOPS Adapters configured with firmware for ioMemory VSL software version 2.x, you must carefully follow the instructions in the Upgrading Devices from VSL 2.x to 3.x on page 54. (Follow those instructions instead of the normal installation instructions.)

3. If you have a previous version of the ioMemory VSL software installed, you will need to uninstall the ioMemory VSL software and the utilities.

4. Install the latest version of the ioMemory VSL software and command-line utilities.

5. Reboot the ESX(i) system; this will load the ioMemory VSL software driver and attach the IBM High IOPS Adapter(s).

6. Determine if you need to upgrade the firmware to the latest version, see Upgrading the Firmware on page 17.

7. Follow the instructions in Configuring the Device to Support VM Disks on page 19.

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 ioMemory device is not supported. ESX(i) installers may permit you to install the ESX(i) OS on an ioMemory device. 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 http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**). Navigate to the appropriate folder for your operating system.

Attention!

See **SCSI vs. Block Versions** below for more information on the ioMemory VSL software options.

Example offline bundles to download:

- `block-iomemory-vsl_<version>.offline-bundle.zip`
- `scsi-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`

Firmware

Also download the firmware archive file, for example:

`fio-firmware-highiops-<version>.<date>.fff`

libvsl Package

We recommend downloading and installing the VSL management library. This library is required if you decide to use a remote management tool based on the SMI-S interface. For more information on these tools see Monitoring and Managing Devices on page 23. Example filename:

- `libvsl-1.0.0-5X-offline-bundle.<version>.zip`

SCSI vs. Block Versions

If you are downloading the ioMemory VSL software for ESXi, then you may have the option to download a `scsi` device version of the software or a raw `block` device version of the software. Both of these versions present IBM High IOPS Adapters as block devices, but the `scsi` version attaches the device to the system *specifically* as a SCSI device, supporting SCSI commands, and the Block version attaches the device as a raw block device.

Both versions will work with all IBM High IOPS Adapters that are supported with this release and both are installed in the same way, ***but you can only have one version installed at a time***. You may consider using the `scsi` device version of the ioMemory VSL software to be able to:

- Use devices that are greater than 2TB in capacity.
- Use more than 16 IBM High IOPS Adapters in a host system.

Attention!

Upgrading to SCSI

Read the *Upgrading to SCSI Version on VMware Hypervisors* section in the *IBM ioMemory VSL Release Notes* for this release if you are upgrading from a raw-block version of the ioMemory VSL software to the SCSI version. Please follow the steps provided in that section, as inadvertent selections during the upgrade process may clear data from the device.

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.

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.

Transfer Methods

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

- vSphere Client
- vCLI `vifs` command
- 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:

- `fio-firmware-highiops-<version>.<date>.fff`
- `iomemory-vsl_<version>.offline-bundle.zip`.

Installing the ioMemory VSL on ESXi 5.x

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.

Attention!

Uninstall

An update/upgrade installation is not recommended. Instead, uninstall the previous version of the ioMemory VSL software before you install this version. See Uninstalling the Software on page 28 for more information on uninstalling the software.

Attention!

Upgrading to SCSI

Read the *Upgrading to SCSI Version on VMware Hypervisors* section in the *IBM ioMemory VSL Release Notes* for this release if you are upgrading from a raw-block version of the ioMemory VSL software to the SCSI version. Please follow the steps provided in that section, as inadvertent selections during the upgrade process may clear data from the device.

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.x 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 `datastore1`, the (local) path would be: `/vmfs/volumes/datastore1/bundles/iomemory-vsl_<version>.offline-bundle.zip`

Attention!

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

a. If you also wish to install the `libvsl` offline bundle, repeat the above step using the path to the `libvsl` offline bundle.

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.x 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 `datastore1`, an example (local) path would be: `/vmfs/volumes/datastore1/bundles/iomemory-vsl_<version>.offline-bundle.zip`

a. If you also wish to install the `libvsl` offline bundle, repeat the above step using the path to the `libvsl` offline bundle.

2. Reboot your ESXi system

The ioMemory VSL and command-line utilities are installed on the host. Continue on to Upgrading the Firmware on page 17.

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

Attention!

Uninstall

An update/upgrade installation is not recommended. Instead, uninstall the previous version of the ioMemory VSL software before you install this version. See Uninstalling the Software on page 28 for more information on uninstalling the previous versions of the software.

vCLI Installation

The installation bundles shouldn't reside on the ESX(i) 4.x host when 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`

a. If you also wish to install the `libvsl` offline bundle, repeat the above step using the path to the `libvsl` offline bundle.

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.x) 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 `esxupdate` 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`

a. If you also wish to install the `libvsl` offline bundle, repeat the above step using the path to the `libvsl` offline bundle.

3. Reboot the host system.

The ioMemory VSL and command-line utilities are installed on the host. Continue on to Upgrading the Firmware on page 17.

Upgrading the Firmware

With the ioMemory VSL software loaded, you need to check to ensure that the IBM High IOPS Adapter's firmware is up-to-date and then update the firmware if needed. You can do this with the command-line utilities.

NOTE-

Make sure you have downloaded the firmware archive file that goes with this version of the ioMemory VSL software.

Attention!

There is a specific upgrade path that you must take when upgrading an IBM High IOPS Adapter. Consult the *IBM ioMemory VSL Release Notes* for this ioMemory VSL software release before upgrading IBM High IOPS Adapters.

Attention!

Do not attempt to downgrade the firmware on any IBM High IOPS Adapter, doing so may void your warranty.

When installing a new IBM High IOPS Adapter along with existing devices, you must upgrade all of the currently installed devices to the latest available versions of the firmware and ioMemory VSL software before installing the new devices. Consult the *IBM ioMemory VSL Release Notes* for this ioMemory VSL software release for any upgrade considerations.

Attention!

Upgrading VMware Guest OS

If you are using your IBM High IOPS Adapter with a VMware guest OS (using VMDirectPathIO), you must cycle the power on the VMware host server after you upgrade the device(s). Just restarting the virtual machine won't apply the change.

More information on these command-line utilities is available in Command-line Utilities Reference on page 37. All command-line utilities require TSM/Shell or COS enabled.

1. Run the `fio-status` utility and examine the output. See fio-status on page 43 for usage information.
    - If any device is in minimal mode and the reason is outdated firmware.
    - If the a device is not in minimal mode, but the firmware listed for that device is a lower number than the latest firmware version available with this version of the ioMemory VSL software, then the firmware is old, but not outdated.
2. If the firmware is old or outdated, update it using the `fio-update-iodrive` utility. See fio-update-iodrive on page 46 for complete information and warnings.

Configuration

Once you have your IBM High IOPS Adapter and ioMemory VSL software installed and loaded, and the firmware on the device is current, you may need to configure the device and/or software. This section outlines some of the common configurations that you may need to consider.

Configuring the Device to Support VM Disks

Attention!

512B Sector Sizes

ESX(i) VMFS requires 512B physical sector sizes. Depending on your device type and how it has been used, you many need to format it to 512B sectors. You may want to run `fio-format` utility against all of the devices in the system.

- If needed, format the device to 512B physical sector sizes.

  Attention!

  Formatting the device will erase all user data, see fio-format on page 41 for full information.

- Run the following command in the COS/Shell/TSM:

```
fio-format /dev/fct*
```

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**.

Attention!

The type of virtual disk we recommend is "thick." A "thin" provisioning will conserve space, but may degrade performance significantly.

Once the storage has been added and configured, 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 software.

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 software 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. If 0.5% is not enough, consult the RAM requirements in the *IBM ioMemory VSL Release Notes* for an idea of how much RAM the VSL may use in worst-case scenarios.

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 320000MB).

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

Performance and Tuning

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

As IBM High IOPS Adapters improve 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 High IOPS Adapters generally perform well out of the box, this section describes some of the common areas where tuning may help achieve optimal performance.

Disable CPU Frequency Scaling

Dynamic Voltage and Frequency Scaling (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 IOPS. When tuning for performance, you may benefit from reducing or disabling DVFS completely, even though this may increase power consumption.

DVFS, if available, is often 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 can 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.

Setting ACPI C-State Options

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

Monitoring and Managing Devices

IBM provides many tools for managing your IBM High IOPS Adapters. These tools will allow you to monitor the devices for errors, warnings, and potential problems. They will also allow you to manage the devices including performing the following functions:

- Firmware upgrades
- Low-level formatting
- Attach and detach actions
- Device status and performance information
- Configuring Swap and Paging
- Generating bug reports

Management Tools

IBM has provided several tools for monitoring and managing IBM High IOPS Adapters. These include stand-alone tools that require no additional software and data-source tools that can be integrated with other applications.

Consider the descriptions of each tool to decide which tool (or combination of tools) best fits your needs.

Attention!

The ioMemory VSL software does print some error messages to the system logs, and while these messages are very useful for troubleshooting purposes, the ioMemory VSL software log messages are not designed for continual monitoring purposes (as each is based on a variety of factors that could produce different log messages depending on environment and use case). For best results, use the tools described in this section to regularly monitor your devices.

Stand-alone Tools

- **COS/Shell/TSM Command-line Utilities**: These utilities are installed with the ioMemory VSL software and are run manually in a terminal. In order to use these utilities on ESXi, the Shell/TSM (Tech Support Mode) must be enabled. The `fio-status` utility provides status for all devices within a host. The other utilities allow you to perform other management functions. See Command-line Utilities Reference on page 37 for full details.

Data-source Tools

These data-source tools provide comprehensive data, just like the stand-alone tools, but they do require integration with additional software. At a minimum, some tools can interface with a browser. ***However, the benefit of these tools is that they can be integrated into existing management software that is customized for your organization.***

These tool packages and documentation are also available as separate downloads (separate from the ioMemory VSL software packages).

- **SMI-S CIM Provider**: The CIM provider allows you to monitor and manage your devices using the Common Information Model. You can use a normal CIM browser, or customize your existing application to interface with the CIM provider.
    - **Optional Remote SMI-S Scripts**: These optional scripts provide remote management of the software and devices without enabling Tech Support Mode (TSM) or logging in to the COS. These scripts integrate with the CIM provider and together the two tools make a complete solution.

Example Conditions to Monitor

This section gives examples of conditions you can monitor. It is intended as an introduction and not as a comprehensive reference. These conditions will have slightly different names, states, and values, depending on the tool you choose. For example, an SNMP MIB may have a different name than a SMI-S object or an API function.

In order to properly monitor these conditions, you should become familiar with the tool you choose to implement and read the documentation for that tool. You may also discover additional conditions that you wish to frequently monitor.

For quick reference, the possible states/values of these conditions are described as Normal (**GREEN**), Caution/Alert (**YELLOW**), or Error/Warning (**RED**). You may implement your own ranges of acceptable states/values, especially if you use a data-source tool.

Device Status

All of the monitoring tools return information on the status of the IBM High IOPS Adapters, including the following states:

GREEN	Attached
YELLOW	Detached, Busy (including: Detaching, Attaching, Scanning, Formatting, and Updating)
RED	Minimal Mode, Powerloss Protect Disabled

If the device is in Minimal Mode, the monitoring tool can display the reason for the Minimal Mode status.

Required Actions

If the device is in Minimal Mode, the action will depend on the reason. For example, if the reason is outdated firmware, then you will need to update the firmware.

Temperature

IBM High IOPS Adapters require adequate cooling. In order to prevent thermal damage, the ioMemory VSL software will start throttling write performance once the on-board controller reaches a specified temperature. If the controller temperature continues to rise, the software will shut down the device once the controller temperature reaches the maximum operating temperature.

These temperatures depend on the device. Newer IBM High IOPS Adapters have higher thermal tolerances. Consult the *IBM High IOPS Adapter Hardware Installation Guide* to determine the thermal tolerances of all devices you will monitor. **This table uses the controller thermal tolerances for newer devices** (93°C throttling, 100°C shutdown).

GREEN	<93°C
YELLOW	93-99°C
RED	100°C

You may wish to shift the conditions by a few degrees so the **YELLOW** condition exists before throttling occurs. For example:

GREEN	<90°C
YELLOW	90-96°C
RED	97°C

Attention!

NAND Board Temperature

Newer IBM High IOPS Adapters also report the temperature of the NAND Boards. This is also a critical temperature to monitor. Consult the *IBM High IOPS Adapter Hardware Installation Guide* to see if your device reports this temperature and to see the temperature thresholds.

Required Actions

If the temperature is at or approaching the **YELLOW** condition, thermal mitigation steps may be necessary. Evaluate the server environment and system requirements necessary to operate the High IOPS adapters. Server operating conditions are documented in the user guides for the server and the requirement to operate High IOPS adapter is at the following website, which may include updates to uEFI and IMM code levels:

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=SERV-IOPS

Health Reserves Percentage

IBM High IOPS Adapters are highly fault-tolerant storage subsystem with 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 pro-actively monitoring device age and health, you can ensure reliable performance over the intended product life. The following table describes the Health Reserve conditions.

GREEN	>10%
YELLOW	0-10%
RED	0%

At the 10% healthy threshold, a one-time warning is issued. At 0%, the device is considered unhealthy. It enters *write-reduced* mode. After the 0% threshold, the device will soon enter *read-only* mode.

For complete information on Health Reserve conditions and their impact on performance, see Monitoring the Health of Devices on page 49.

Required Actions

The device needs close monitoring as it approaches 0% reserves and goes into write-reduced mode, which will result in reduced write performance. Prepare to replace the device soon.

Write (Health Reserves) Status

In correlation with the Health Reserves Percentage, the management tools will return write states similar to these:

GREEN	Device is healthy
YELLOW	Device is getting close to entering reduced write mode.
RED	Device has entered reduced-write or read-only mode to preserve the flash from further wearout.

Required Actions

The device needs close monitoring as it approaches 0% reserves and goes into write-reduced mode, which will result in reduced write performance. Prepare to replace the device soon.

Device LED Indicators

If you have physical access to the devices and depending on your device configuration, you can use the LED indicator(s) on the bracket to monitor their status.

Attention!

IBM High IOPS Adapters may have an additional LEDs (that are not on the bracket, as shown below. You can ignore those other LEDs, as they are not meant for monitoring device and software functionality.

The LEDs on your device should be similar to one of these configurations:

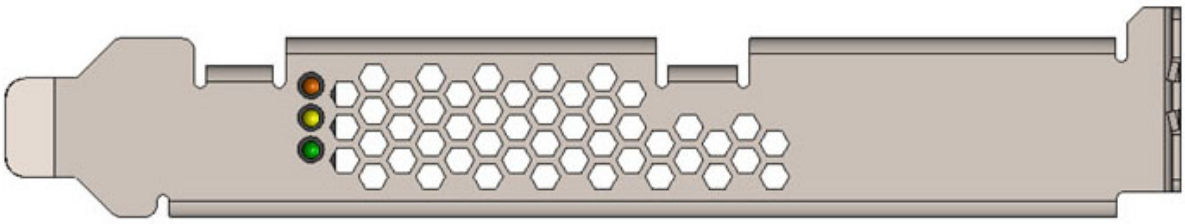

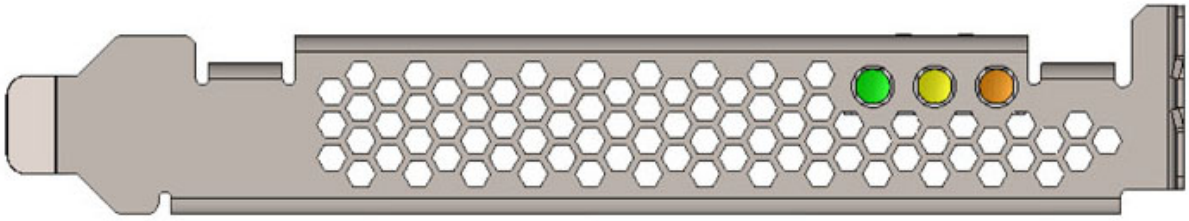

This table explains the information that these LEDs convey:

Green	Yellow	Amber	Indication	Notes
OFF	**OFF**	**OFF**	Power is off.	
OFF	**OFF**	**LIT**	Power is on. Problem with device, or driver not loaded (and device unattached).	Use `fio-status` to view problem, or load driver (and attach device).
LIT	**OFF**	**OFF**	Power is on. Driver loaded (device may or may not be attached).	You may need to attach the device.
LIT	**FLASHING**	**OFF**	Writing (Rate indicates volume of writes).	Can appear in combination with the Read LED indication.
FLASHING	**OFF**	**OFF**	Read (rate indicated volume of reads).	Can appear in combination with the Write LED indication.
LIT	**LIT**	**LIT**	Location Beacon.	Use the `fio-beacon` utility to initiate this behavior.

Maintenance

This section explains additional software maintenance functions not covered in the sections Configuration on page 19 and Monitoring and Managing Devices on page 23.

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).

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 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.

Uninstalling the Software

Uninstalling the software in ESXi 5.x

To uninstall the ioMemory VSL package, run this command using the vCLI from a **remote machine** (remove `--server <servername>` if you are on the host CLI):

1. Remove the VIB containing the driver & utilities:

   Block version command:

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

   SCSI version command:

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

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------
iomemory-vsl-3.0.6.360

-----Installed-----
2012-01-16T03:49:33

------------------Summary-------------------
iomemory-vsl: block driver for ESX/ESXi 4.X
```

2. Remove ioMemory VSL using its `Bulletin ID`:

```
$ esxupdate -b <Bulletin-ID> remove
```

Uninstalling the Software in ESXi 4.x

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
```

Unloading or Disabling the Software Driver

If you need to diagnose or troubleshoot a problem, you may need to unload or disable the ioMemory VSL software. Both methods will take the ioMemory VSL software offline, however we recommend disabling the ioMemory VSL software autoload and rebooting rather than unloading the ioMemory VSL driver.

Attention!

You **must** properly unmount and detach all IBM High IOPS Adapters before unloading the ioMemory VSL driver. However if you disable autoload and then reboot, the devices will safely unmount and detach on shutdown and then not auto attach or mount on boot.

If you must unload the driver (and detach an IBM High IOPS Adapter), carefully read all of the warnings in fio-detach on page 40 (or the optional remote management documentation) before running the detach utility. Failure to follow the instructions may cause errors, data loss and/or corruption.

Again, we recommend disabling autoload rather than directly unloading the driver.

Disabling the ioMemory VSL software Autoload

The ioMemory VSL driver automatically loads by default when the operating system starts. By disabling the autoload and rebooting, the ioMemory VSL software will be offline.

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, before you reboot you must also run the following command to save the boot configuration:

```
esxcfg-boot -b
```

NOTE-

If you disable the driver autoload in ESXi 5.1 or newer, you also disable the ability to directly load the driver. You will need to re-enable the driver to load the driver.

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 available.

Unloading the ioMemory VSL Driver

If you need to unload the driver for diagnostic or troubleshooting purposes, you **must** properly unmount and detach all IBM High IOPS Adapters.

To unload the ioMemory VSL driver, run this commands in COS/TSM:

```
vmkload_mod -u iomemory-vsl
```

Loading or Enabling the Driver

Loading or Enabling the ioMemory VSL Driver

If you have previously disabled driver autoload, you may want to re-enable it and then reboot. Otherwise, you may bring the driver online by loading directly. If you load the driver, you will need to re-attach the IBM High IOPS Adapters in your system. See fio-detach on page 40 for more information.

NOTE-

If you disable the driver autoload in ESXi 5.1 or newer, you also disable the ability to directly load the driver. You will need to re-enable the driver to load the driver.

If you enable driver autoload and then reboot, the IBM High IOPS Adapters should attach automatically (unless you have disabled auto attach, see below).

NOTE-

Depending on your situation, you may want to both re-enable driver autoload (to ensure that the ioMemory VSL driver will load on the next boot) and directly load the driver to bring it online immediately.

Enabling ioMemory VSL Driver Autoload

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, before you reboot you must also run the following command to save the boot configuration:

```
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
```

Loading the ioMemory VSL Driver

To immediately load the ioMemory VSL software, run the following command:

```
vmkload_mod iomemory-vsl
```

This command loads the driver with default parameters (even if you have modified the parameters). You can force the driver to load with modified perameters, for example:

```
vmkload_mod iomemory-vsl auto_attach=0
```

See Using Module Parameters on page 51 for a list of these parameters.

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 fio-detach on page 40 (or the optional remote management documentation) before running the detach utility. Failure to follow the instructions may cause errors, data loss and/or corruption.

Disabling Auto-Attach

When the ioMemory VSL software is installed, it is configured to automatically attach any devices when the ioMemory VSL software is loaded. Sometimes you may want to disable the auto-attach feature (to assist in troubleshooting or diagnostics).

Disabling Auto-Attach in ESX

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

```
$ 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`.

Disabling Auto-Attach in ESXi

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

```
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`.

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 (also called a rescan) repairs these data structures.

Improving Rescan Times

The rescan of the device (also called a consistency check) the VSL performs after an unmanaged shutdown may take an extended period of time depending on the total capacity of the device(s) that the ioMemory VSL software needs to scan.

Default Fast Rescan

By default, all IBM High IOPS Adapters formatted with the fio-format utility are formatted to have improved rescan times. You can disable this default fast rescan by reformatting the device and using the `-R` option. Disabling this feature will reclaim some reserve capacity that is normally set aside to help improve rescan times.

If you leave the default fast rescan feature in place you can also take further steps to improve rescan times by implementing one of the following module parameters.

Faster Rescans Using Module Parameters

These two module parameters require the default fast rescan formatting structure, and they also use system memory (RAM) to help improve rescan times. The extra memory enables the rescan process to complete faster, which reduces downtime after a hard shutdown. This memory allocation is only temporary and is freed up after the rescan process is complete.

If you decide to use one of these parameters, you will need to set the upper limit of RAM used by that parameter. To do this, you will need to determine how much RAM each parameter may use in your scenario, how much system RAM is available, and (therefore) which parameter is more suited for your use case.

For more information on setting module parameters, see Using Module Parameters on page 51.

Here is a quick comparison of the two parameters:

- **RMAP Parameter**
  - **Fastest**: This improvement results in the fastest rescan times.
  - **Less Scalable**: (All or nothing.) This parameter requires enough RAM to function. If the RAM limit is set too low, then the ioMemory VSL software will not use RMAP at all, and it will revert back to the default fast rescan process.
  - **Target Scenario**: This parameter will improve any use case if there is enough RAM available for the parameter. It is more suited for smaller capacity IBM High IOPS Adapters and/or systems with fewer IBM High IOPS Adapters installed. We also recommend it for devices that have been used for many small random writes.
- **RSORT Parameter**

- **Faster**: This improves rescan times over the default fast rescan process.
- **Scalable**: With this parameter, the ioMemory VSL softwareworks with the system RAM to improve rescan times until it reaches the RAM limit set in the parameter. At that point, the software reverts back to the default fast rescan process.
- **Target Scenario**: This parameter will improve rescan times in any use scenario. It is especially useful in systems with multiple IBM High IOPS Adapters and/or larger-capacity IBM High IOPS Adapters. We also recommend it when IBM High IOPS Adapters are used to store databases.

RMAP Parameter

The `rmap_memory_limit_MiB` parameter sets the upper memory (RAM) limit (in mebibytes) used by the ioMemory VSL software to perform the RMAP rescan process. You should only use this option if you have enough memory for all of your IBM High IOPS Adapters in the system. If you do not have enough memory to use this option, use the RSORT parameter instead.

Because this parameter requires a set amount of memory, it often works best with fewer IBM High IOPS Adapters and/or smaller-capacity IBM High IOPS Adapters in a system, but the determining factor is how much memory is in the system and whether there is enough to set the appropriate memory limit.

This parameter requires `4.008` bytes of RAM per block of IBM High IOPS Adapter capacity.

1. First determine the number of blocks that are formatted for each device.
   a. This information is visible when you format the device using the `fio-format` utility.
   b. Or you can estimate the number of block using the device capacity and the formatted sector size.

      This example shows a quick estimation of the number of blocks on a `400GB` device with `512B` size sectors (2 sectors per KB):

      ```
      400GB * 1000MB/GB * 1000KB/MB * 2 Blocks/kB = 800,000,000 Blocks
      ```

2. Multiply the number of blocks by `4.008` bytes of RAM per block (and translate that into MiB) to determine the memory limit that is required for this parameter to function.
   a. In the example above there were `800 million` blocks:

      ```
      800,000,000 Blocks * 4.008B/Block * 1KiB/1024B * 1MiB/1024KiB =
      ~3058MiB of RAM
      ```

   b. In this example, you would need about 3100 MiB of RAM available in your system for a 400GB IBM High IOPS Adapter formatted for 512B sectors, and you would need to set the RMAP parameter to `3100`.

NOTE-

Default Value

The RMAP parameter is, by default, set to `3100`. It is set to this low default value so the rescan process does not use all of the RAM in systems that have less available memory.

- If the RMAP value is too low for the number of IBM High IOPS Adapter blocks in the system, then the ioMemory VSL software will not use the RMAP process to improve rescan times, it will just use the default fast rescan process. (RMAP is an all-or-nothing setting.)
- If you don't have enough system memory to use the RMAP parameter, consider using the RSORT parameter. The RSORT parameter will use its RAM limit to improve the rescan process, and then the ioMemory VSL softwarerevert to the default fast rescan process to finish the consistency check.

3. Set the module parameter to the value you have determined. See Using Module Parameters on page 51 for more information on setting parameters.

RSORT Parameter

The `rsort_memory_limit_MiB` parameter sets the memory (RAM) limit used by the ioMemory VSL software to perform the RSORT rescan process. The RSORT rescan process is faster than the default rescan process and we recommend using it to rescan devices that are used datastores for databases.

If this parameter is given any memory limit, the ioMemory VSL software will use the RSORT process until either the rescan is done or it consumes the memory limit. If the process runs out of memory, it will revert to the default fast rescan process. However, in order to optimize the use of this process, you can calculate the target RAM usage and set the limit based on that target. There is no penalty for setting a high limit, the RSORT process will only use the RAM it needs (up to the limit that is set).

This target is based on `32` bytes per write extent. For example, if your database writes 16kB at a time, there is one write extent per 16kB of IBM High IOPS Adapter capacity.

NOTE-

Blocks per Write Extent

One measure of the the benefits of the RSORT process is to see how many blocks are written per write extent. The RSORT process improves rescan times over the default fast rescan process on when a device has `8` or more blocks written per extent. For example, if your IBM High IOPS Adapter is formatted to `512B` sector sizes (2 sectors per KB), and your database writes in `8KB` chunks, then your database writes `16` blocks per write extent and RSORT would improve the rescan times.

1. First determine the number of blocks that are formatted for each device.
   a. This information is visible when you format the device using the `fio-format` utility.
   b. Or you can estimate the number of block using the total device capacities and their formatted sector sizes.

This example shows a quick estimation of the number of blocks on 1200GB of IBM High IOPS Adapter capacity with 512B size sectors (2 sectors per KB):

```
1200GB * 1000MB/GB * 1000KB/MB * 2 Blocks/kB = 2,400,000,000
Blocks
```

2. Divide the number of blocks by the write extents per block to determine the total possible number of write extents on the device(s).

    a. In the example above there were 2.4 billion blocks. We will assume 16KB write extents (32 blocks per write on 512B sectors):

```
2,400,000,000 Blocks * 1 Write Extent/32 Blocks = 150,000,000
Writes
```

3. Multiply the number of writes by 32 bytes of RAM per write (and translate that into MiB) to determine the memory target for this parameter.

    a. In the example above there were 150 million write extents:

```
150,000,000 Writes * 32B/Write * 1KiB/1024B * 1MiB/1024KiB =
~4578MiB of RAM
```

    b. In this example, you would want to set the RSORT limit to about 4600 MiB of RAM available in your system for 1200GB of IBM High IOPS Adapter capacity formatted for 512B sectors.

NOTE-

Default Value

The RMAP parameter is, by default, set to 0m and it has a maximim of 100000 (100GB).

---

4. Set the module parameter to the value you have determined. See Using Module Parameters on page 51 for more information on setting parameters.

Appendix A - Command-line Utilities Reference

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.

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 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 Management Tools on page 23

The ioMemory VSL software installation packages include various command-line utilities, installed by default in `/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 ioMemory 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.
`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

Description

Attaches the IBM High IOPS Adapter and makes it available to the operating system. The command displays a progress bar and percentage as it operates.

NOTE-

In most cases, the ioMemory VSL software automatically attaches the device on load and does a scan. You only need to run fio-attach if you ran fio-detach or if you set the ioMemory VSL software's auto_attach parameter to 0.

NOTE-

If the IBM High IOPS Adapter is in minimal mode, then auto-attach is disabled until the cause of the device being in minimal mode is fixed.

Syntax

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

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

Option	Description
`-r`	Force a metadata rescan. This may take an extended period of time, and is not normally required. **Attention!** Only use this option when directed by Customer Support.
`-c`	Attach only if clean.
`-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 LED(s) to locate the device. You should first detach the IBM High IOPS Adapter and then run `fio-beacon`.

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 detected on the system. The device numbers are visible using `fio-status`.

Option	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.

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 contact Customer Support at http://www.ibm.com/systems/support for assistance in troubleshooting.

The output will indicate where the bugreport is saved.

Attention!

The `fio-bugreport` utility uses `vm-support` to gather information including system logs. If `vm-support` fails, this utility will fail to gather the system logs.

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-20090921.173256-sdv9ko.tar.bz2
  to your support case, including steps to reproduce the problem.
  If you do not have an open support case for this issue, please open a
support
  case with a problem description and then attach this file to your new
case.
```

For example, the filename for a bug report file named `fiobugreport-20090921.173256-sdv9ko.tar.bz2` indicates the following:

- Date (20090921)
- Time (173256, or 17:32:56)
- Misc. information (`sdv9ko.tar.bz2`)

fio-detach

Description

Detaches the IBM High IOPS Adapter. By default, the command also 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 **not** recommend using the `fio-detach` utility to ensure that a device is detached. Instead, as a best practice, follow the instructions in Disabling Auto-Attach on page 32 as a safe detach workaround.

Before you use this utility, you must make sure each device is unmounted, and we recommend that you put the host system in maintenance mode.

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 detected on the system.

Option	Description
-q	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 some process has the device open.

fio-format

Description

Attention!

The fio-format utility requires that the ioMemory VSL software be loaded with the IBM High IOPS Adapter(s) detached. Refer to fio-detach on page 40 for details.

Performs a low-level format of the IBM High IOPS Adapter. 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.

NOTE-

If you do not include the -s or -o options, the device size defaults to the advertised capacity. If used, the -s and -o options must include the size or percentage indicators.

Attention!

Do not interrupt the formatting! We recommend adding power backup to your system to prevent power failures during formatting. If formatting is interrupted, please contact Customer Support.

Syntax

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

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

Options	Description
-b <size B|K>	Set the block (sector) size, in bytes or kibibytes (base 2). The default is 512 bytes. For example: -b 512B or -b 4K (B in 512B is optional). **Attention!** 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 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.)
-q	Quiet mode: Disable the display of the progress-percentage indicator.
-s <size M|G|T|%>	Set the device capacity as a specific size (in TB, GB, or MB) or as a percentage of the advertised capacity: • 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)
-o <size B|K|M|G|T|%>	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 -o option; see the -s option above for size indicator descriptions.) **Attention!** Before you use this option, please discuss your use case with Customer Support.
-R	Disable fast rescan on unclean shutdown to reclaim some reserve capacity.
-y	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 on page 38 for details.

fio-pci-check

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 correctable errors 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]
```

Option	Description
`-d <value>`	`1` = Disable the link; `0` = bring the link up (Not recommended).
`-e`	Enable PCI-e error reporting.
`-f`	Scan every device in the system.
`-n`	Do not perform any writes to config space. Will prevent errors from being cleared.
`-o`	Optimize the IBM High IOPS Adapter PCIe link settings by increasing the maximum read request size if it is too low.
`-r`	Force the link to retrain.
`-v`	Verbose: Print extra data about the hardware.

fio-status

Description

Provides detailed information about the installed devices. This utility operates on either `fctx` or `fiox` devices. The utility depends on running as root and having the ioMemory VSL driver loaded. If no driver is loaded, a smaller set of status information is returned.

`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 detected on the system.

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

Option	Description
-a	Report all available information for each device.
-e	Show all errors and warnings for each device. This option is for diagnosing issues, and it hides other information such as format sizes.
-c	Count: Report only the number of 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.
-fj	Format JSON: creates the output in JSON format.
-fx	Format XML: creates the output in XML format.
-u	Show unavailable fields. Only valid with -fj or -fx.
-U	Show unavailable fields and details why. Only valid with -fj or -fx. **NOTE-** Some fio-status fields are unavailable depending on the operating system or device. For example, some legacy fields are unavailable on newer IBM High IOPS Adapters.
-F<field>	Print the value for a single field (see the next option for field names). Requires that a device be specified. Multiple -F options may be specified.
-l	List the fields that can be individually accessed with -F.

Attention!

Output Change

The standard formatting of fio-status ouput has changed compared to the output from ioMemory VSL software version 2.x. This will affect any custom management tools that used the output of this utility.

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 software 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 software load) in degrees Centigrade
- Health status: healthy, nearing wearout, write-reduced or read-only
- Reserve capacity (percentage)
- Warning capacity threshold (percentage)

Data Volume Information: If the `-d` option is used, 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: If the ioMemory VSL software is in minimal mode, read-only mode, or write-reduced mode when `fio-status` is run, the following differences occur in the output:

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

fio-update-iodrive

Attention!
You should back up the data on the IBM High IOPS Adapter prior to any upgrade as a precaution.

Description

Attention!
Your IBM High IOPS Adapters must be detached before running `fio-update-iodrive`. See fio-detach on page 40 for details or Disabling Auto-Attach on page 32 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.

Attention!
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 firmware upgrades in sequence, it is critical to load the ioMemory VSL driver after each firmware upgrade step. 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. Contact Customer Support at at http://www.ibm.com/systems/support if you have issues with your upgrade.

NOTE-

The default action (without using the `-d` option) is to upgrade all IBM High IOPS Adapters with the firmware contained in the `fio-firmware-highiops-<version>.<date>.fff` firmware archive 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 software 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.

To update one or more specific devices:

- If the ioMemory VSL driver is loaded, use the `-d` option with the device number.

Syntax

```
fio-update-iodrive [options] <firmware-path>
```

where `<firmware-path>` is the full path to the firmware archive file `fio-firmware-highiops-<version>.<date>.fff` available at

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**). 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.

Option	Description
`-d`	Updates the specified devices (by `fctx`, where {i}x{i} is the number of the device shown in `fio-status`). If this option is not specified, all devices are updated. **Attention!** Use the `-d` option with care, as updating the wrong IBM High IOPS Adapter could damage your device.
`-f`	Force upgrade (used when directed by Customer Support). **Attention!** Use the `-f` option with care, as it could damage your card.
`-l`	List the firmware available in the archive.
`-p`	Pretend: Shows what updates would be done. However, the actual firmware is not modified.
`-c`	Clears locks placed on a device.
`-q`	Runs the update process without displaying the progress bar or percentage.
`-y`	Confirm all warning messages.

If you arrived at this section from Upgrading the Firmware on page 17, you should return to that section.

Appendix B - Monitoring the Health of 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.

Health Metrics

The ioMemory VSL software manages block retirement using pre-determined retirement thresholds. The IBM Flash Management Console and the `fio-status` utilities show a health indicator that starts at 100 and counts down to 0. As certain thresholds are crossed, various actions are taken.

At the 10% healthy threshold, a one-time warning is issued. See Health Monitoring Techniques on page 49 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.

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 form 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 -a`: Output from the `fio-status` utility (using the `-a` option) shows the health percentage and device state. These items are referenced as "`Media status`" in the sample output below.

```
Found 3 ioMemory devices in this system
Fusion-io driver version: 3.x.x build xxxx

Adapter: Single Adapter
```

```
        Fusion-io ioDrive 1.30TB, Product Number:F00-001-1T30-CS-0001,
SN:1133D0248, FIO SN:1134D9565
...
Media status: Healthy; Reserves: 100.00%, warn at 10.00%; Data: 99.12%
Lifetime data volumes:
   Physical bytes written: 6,423,563,326,064
   Physical bytes read   : 5,509,006,756,312
```

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

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

IBM Flash Management Console: In the Device Report tab, look for the Reserve Space percentage in the right column. The higher the percentage, the healthier the drive is likely to be.

Software RAID and Health Monitoring

Software RAID stacks are typically designed to detect and mitigate the failure modes of traditional storage media. The IBM High IOPS Adapter attempts to fail as gracefully as possible, and these new failure mechanisms are compatible with existing software RAID stacks. An IBM High IOPS Adapter in a RAID group will fail to receive data at a sufficient rate if: a) the device is in a write-reduced state, and b) it is participating in a write-heavy workload. In this case, the device will be evicted from the RAID group. A device in read-only mode will be evicted when write I/Os are returned from the device as failed. Catastrophic failures are detected and handled just as though they are on traditional storage devices.

Appendix C - Using Module Parameters

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

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 (min/max)	Description
`auto_attach`	`1 (0, 1)`	`1` (default) = Always attach the device on driver load. `0` = Don't attach the device on driver load.
`background_attach`	`1 (0, 1)`	`1` (default) = The host scans and attaches IBM High IOPS Adapters in the background while the system continues to boot. **Attention!** The default setting saves time on startup. However, in some instances the IBM High IOPS Adapters may not be ready when the host system scans for datastores. If this is occurring, set this parameter to `0` and reboot. `0` = Force the host server to wait until all IBM High IOPS Adapters are ready and attached before continuing to fully boot and become operational.
`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-

Other than , 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 D - 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 software on the ESX(i) system. Instead, install the software on the guest system. Only install the ioMemory VSL software on the host 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 on the VMware host server 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 E - Upgrading Devices from VSL 2.x to 3.x

This version of the ioMemory VSL software supports new features, including the latest generation of High IOPS architecture and improved Flashback protection. These features require the latest version of the IBM High IOPS Adapter firmware. Every IBM High IOPS Adapter in a system running 3.1.x or later must be upgraded to the latest version of the firmware.

For example, if you have a system running 2.x ioMemory VSL software with Legacy IBM High IOPS Adapters previously installed, and you want to install new Second Generation IBM High IOPS Adapters (that require the latest version of the firmware), then you will need to upgrade all of the existing devices to the latest firmware version.

Attention!

You cannot revert a device's firmware to an earlier version once you have upgraded the device (without voiding your warranty). If you experience problems with your upgrade, please contact Customer Support at http://www.ibm.com/systems/support.

Attention!

Upgrading devices (previously configured for VSL 2.x.x) to work with VSL 3.x.x will require a low-level media format of the device. No user data will be maintained during the process. Be sure to backup all data as instructed.

Attention!

Upgrade Path

Depending on the current firmware version of your devices, you may need to upgrade your device's firmware multiple times in order to preserve internal structures. Consult the ioMemory VSL software for the upgrade path. Visit http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**) for all of the required software and firmware versions.

For more information on upgrading from one version to the next, see the *IBM ioMemory VSL Release Notes* (available at http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**)) for the version you will upgrade the device to. Then follow the upgrade instructions in that version's user guide for your operating system (including the firmware update instructions).

Upgrade Procedure

Be sure to follow the upgrade path in the *IBM ioMemory VSL Release Notes*. Make sure that all previously installed Legacy IBM High IOPS Adapters are updated with the appropriate firmware.

Attention!

If you plan to use Legacy IBM High IOPS Adapters and Second Generation IBM High IOPS Adapters in the same host, perform this upgrade on all existing Legacy IBM High IOPS Adapters **before** installing the new Second Generation IBM High IOPS Adapters.

1. Prepare each existing Legacy IBM High IOPS Adapter for upgrade.

    a. Backup user data on each device.

    Attention!

    The upgrade process will require a low-level media format of the device. No user data will be maintained during the process; be sure to make a complete backup.

    Use a backup method of your choice. For best results, use software and backup devices that have proven effective in the past. Do not backup the data onto another IBM High IOPS Adapter on the same system. The back up must be to a local disk or to an externally attached volume.

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

    ```
    fio-bugreport
    ```

    c. Detach Legacy IBM High IOPS Adapters, for example:

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

    For more information, see fio-detach on page 40.

2. Uninstall the 2.x ioMemory VSL software:

    a. Stop all of the virtual machines and put the host in maintenance mode.

    b. Follow the instructions for your platform:

        i. ESXi 5.x uninstall vCLI command:

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

        ii. ESXi 4.x vCLI uninstall procedure:

            i. Determine the bundle "bulletin" ID:

            ```
            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
```

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

```
vihostupdate --server <server-name> --remove --
bulletin <Bulletin-ID>
```

iii. ESX 4.x uninstall procedure:

i. Enter the following command to determine the Bulletin ID of any previously installed versions:

```
esxupdate query
```

Sample Output

```
------Bulletin ID------
iomemory-vsl-3.0.6.360

-----Installed-----
2012-01-16T03:49:33

------------------Summary-------------------
iomemory-vsl: block driver for ESX/ESXi 4.X
```

ii. Run the following command to remove installed version(s):

```
esxupdate -b <Bulletin-ID> remove
```

3. Install the new ioMemory VSL software and related packages.

a. Download the ioMemory VSL offline bundle, firmware, and (optional) CIM provider from http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**). For more information, see the section Downloading the Software on page 12 at the beginning of this guide.

b. Install the ioMemory VSL software and utilities. Follow the instructions in one of the following sections, then return to this procedure:

NOTE-

Be sure to follow the suggestion to copy over the firmware file to the ESX(i) host, and make note of the firmware location.

- Installing the ioMemory VSL on ESXi 5.x on page 15
- Installing the ioMemory VSL on ESX(i) 4.x on page 16

c. Reboot the system.

4. Update the firmware on each device to the latest version using `fio-update-iodrive`.

Attention!

Prevent Power Loss

Take measures to prevent power loss during the update, such as a UPS. Power loss during an update may result in device failure. For all warnings, alerts, and options pertaining to this utility, see fio-update-iodrive on page 46.

Sample syntax:

```
fio-update-iodrive <firmware-path>
```

Where `<firmare-path>` is the full path to the firmware archive file (`fio-firmware-highiops-<version>.<date>.fff`) available at http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-65723 (follow that link and then select **IBM High IOPS software matrix**). This command will update all of the devices to the selected firmware. If you wish to update specific devices, consult fio-update-iodrive on page 46 for more options.

5. Reboot the system

6. **Destructive Step**

Attention!

Running `fio-format` in the next step will erase the entire device, including user data. Once this format is started, the device cannot be downgraded to the 2.x driver without voiding your warranty. If you experience problems with your upgrade, please contact Customer Support at http://www.ibm.com/systems/support.

7. Format each device using `fio-format`, for example:

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

You will be prompted to confirm you wish to erase all data on the device.

Attention!

The format may take an extended period of time, depending on the wear on the device.

8. Attach all Legacy IBM High IOPS Adapters, for example:

```
fio-attach /dev/fct*
```

9. Check the status of all devices using fio-status, for example:

```
fio-status -a
```

Your Legacy IBM High IOPS Adapters have now been successfully upgraded for this version of the ioMemory VSL software. You may now install any Second Generation IBM High IOPS Adapters.

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-65723 (follow that link and then select **IBM High IOPS software matrix**).

IBM part number 00AH233