

IBM ioMemory VSL 3.1.1

User Guide for Linux

IBM ioMemory VSL 3.1.1 User Guide for Linux	4
Introduction	5
About IBM High IOPS Adapters	5
About the ioMemory Virtual Storage Layer (VSL)	5
About Flashback Protection Technology	6
System Requirements	7
Hardware Requirements	7
Supported Linux Distributions	7
Upgrading Legacy Adapters (IMPORTANT)	7
Software Installation	9
Installing RPM Packages	9
Building the ioMemory VSL from Source	11
Loading the ioMemory VSL Facility (Driver)	13
Setting the ioMemory VSL Options	17
Upgrading the Firmware	17
Using the Device as Swap	18
Using the Logical Volume Manager	19
Configuring RAID	19
Understanding Discard (TRIM) Support	22
Maintenance	24
Device LED Indicators	24
GUI Managment	25
Command-line Utilities	25
Common Maintenance Tasks	26
Disabling Auto-Attach	27
Unmanaged Shutdown Issues	27
Disabling the ioMemory VSL	28
Performance and Tuning	29
Disabling DVFS	29
Limiting ACPI C-States	29
Setting NUMA Affinity	30
Setting the Interrupt Handler Affinity	30
Appendix A- Utilities Reference	31
fio-attach	31
fio-beacon	32
fio-bugreport	32

fio-detach	34
fio-format	35
fio-pci-check	36
fio-snmp-agentx	37
fio-status	37
fio-sure-erase	40
fio-update-iodrive	42
Appendix B- Monitoring the Health of Devices	45
NAND Flash and Component Failure	45
Health Metrics	45
Health Monitoring Techniques	46
Software RAID and Health Monitoring	46
Appendix C- Using Module Parameters	47
Appendix D- SMI-S Interface	49
Installing the SMI-S WMI Provider on Linux	50
Linux Testing	50
Description	52
Implementation	54
Indications	60
Appendix E- NUMA Configuration	64
Appendix F- Upgrading Devices from VSL 2.x to 3.x	67
Upgrade Procedure	67
Appendix G- Open Source License Acknowledgement	72
IBM Support	93

IBM ioMemory VSL 3.1.1 User Guide for Linux

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: D0001565-002_3

Published: September 20th, 2012

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 IBM High IOPS Adapter or each of the two IBM High IOPS Adapters of an IBM High IOPS Duo Adapter.

NOTE 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.

About IBM High IOPS Adapters

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 Layer™ (VSL) is the "secret sauce" that gives IBM High IOPS Adapters their amazing performance. The VSL™ 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 its normal life cycle. To extend the life of the adapter, Flashback™ 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 Protection™ 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.
- **RAM Requirements:** The *IBM ioMemory VSL Release Notes* contains memory (RAM) requirements for this version of the software.

Supported Linux Distributions

- Red Hat Enterprise Linux (RHEL) 5.6, 5.7, 6.0, 6.1, 6.2
- SUSE Linux Enterprise Server (SLES) 10 SP4, 11, 11 SP1

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 [Appendix F- Upgrading Devices from VSL 2.x to 3.x](#) of this user guide.

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 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&lnidocid=HELP-FIX>

Software Installation

NOTE All commands require administrator privileges. Use sudo or log in as "root" to run the install.

Installation Overview

1. If you are installing this version of ioMemory VSL on a system with previous versions of IBM High IOPS Duo Adapters installed, you must carefully follow the instructions in the [Appendix F- Upgrading Devices from VSL 2.x to 3.x](#) section.

NOTE If you do not need to upgrade previous versions of IBM High IOPS Duo Adapters to the 3.x.x firmware, but your system does have previous versions of the ioMemory VSL installed, you will need to uninstall the ioMemory VSL package and the utilities. See the [Common Maintenance Tasks](#) section for instructions. Once you have uninstalled the packages, return to this page.

2. Install the latest version of the ioMemory VSL. You can install the ioMemory VSL as

- A pre-compiled binary package
- A source-to-build package

NOTE Follow the instructions under [Installing RPM Packages](#) to determine whether pre-compiled binary packages are available for your kernel version or if you need to [build the driver package from source](#).

3. Install utilities and management software (included in driver installation instructions).
4. [Load the ioMemory VSL](#) and [Set the Options](#).
5. [Upgrade the Firmware](#) to the latest version, if needed (recommended). This applies to IBM High IOPS Adapters that may be using a version of the firmware that is earlier than the latest version.

Installing RPM Packages

To install the Linux ioMemory VSL and utilities:

1. You will need to install a version of the ioMemory VSL that is built for your kernel. To determine what kernel version is running on your system, use the following command at a shell prompt:

```
$ uname -r
```

2. Compare your kernel version with the binary versions of the software available at <http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5083174>.

- If there is a binary version of the software that corresponds to your kernel version, download that. **For example:**

```
iomemory-vsl-<kernel-version>-<VSL-version>.x86_64.rpm
```

- If there is no binary version of the software corresponding to your kernel, download the source package. **For example:**

```
iomemory-vsl-<VSL-version>.src.rpm
```

NOTE Exact package names may vary, depending on software and kernel version chosen.

3. Download all of the support RPM packages from <http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5083174>. These packages provide utilities, firmware, and other files.

- **Examples:**

Package	What is installed
fio-util-<VSL-version>.x86_64.rpm	ioMemory VSL utilities – Recommended
fio-firmware-<firmware-version>.noarch.rpm	Firmware archive – Recommended
libvsl-<version>.x86_64.rpm	Libraries needed for management tools – Recommended
fio-common-<VSL-version>.x86_64.rpm	Files required for the init script – Recommended
fio-sysvinit-<VSL-version>.x86_64.rpm	Init script – Recommended
fio-smis-<version>.x86_64.rpm	SMI-S CIM Provider – Optional see Appendix D- SMI-S Interface
libfio-dev-<version>.x86_64.rpm	2.x Management SDK, deprecated – Optional
libfio-doc-<version>.x86_64.rpm	2.x Management SDK, deprecated – Optional
libvsl-dev-<version>.x86_64.rpm	Current Management SDK – Optional
libvsl-doc-<version>.x86_64.rpm	Current Management SDK – Optional

4. Build the ioMemory VSL software from source: if you downloaded the software source package, you must now follow the instructions in [Building the ioMemory VSL from Source](#). If you downloaded a binary version of the software, continue.

5. Change to the directory to where you downloaded the installation packages.

6. Enter the following command to install the custom-built software package. Use the package name that you just copied/downloaded into that directory.

```
rpm -Uvh iomemory-vsl-<kernel-version>-<VSL-version>.x86_64.rpm
```

7. Enter the following commands to install the support files:

```
rpm -Uvh lib*.rpm
rpm -Uvh fio*.rpm
```

The ioMemory VSL and utilities are installed to the following locations:

Package Type	Installation Location
ioMemory VSL	/lib/modules/<kernel-version>/extra/fio/iomemory-vsl.ko
Utilities	/usr/bin
Firmware	/usr/share/fio/firmware

NOTE Under SLES 10, the `fio-sysvinit` package will install with an error:

```
Service udev has to be enabled for service iomemory-vsl
```

After the package is installed on a SLES 10 system, follow the instructions in [Loading the ioMemory VSL Facility \(Driver\)](#) to modify the `/etc/init.d/iomemory-vsl` file so the `sysvinit` script will function properly.

NOTE IBM High IOPS SSD Management Application, a free browser-based solution for managing IBM High IOPS Adapter, is available as a separate download.

Once the packages are installed, continue to [Loading the ioMemory VSL Facility \(Driver\)](#) later in the section.

Building the ioMemory VSL from Source

The ioMemory VSL is distributed as a source package. If a binary version of the software is not available, you will need to build the ioMemory VSL from source. Use the source package that is made for your distribution. Source packages from other distributions are not guaranteed to work.

1. Download current ioMemory VSL source and support packages from <http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5083174>
2. Change directory to wherever you downloaded the source package.

3. Follow the instructions below to create an installation package tailored to your needs.

Building an RPM Installation Package

To build an RPM installation package:

1. Install the prerequisite files for your kernel version.

NOTE Some of the prerequisite packages may already be in the default OS installation. If your system is not configured to get packages over the network, then you may need to mount your install CD/DVD.

- On RHEL 5/6, you need `kernel-devel`, `kernel-headers`, `rpm-build`, `GCC4`, and `rsync`.

```
$ yum install kernel-devel kernel-headers rpm-build gcc rsync
```

NOTE `yum` may not install the correct (matching) kernel versions of `kernel-devel` or `kernel-headers` packages. `yum` will download the latest version. Use the following command to force `yum` to download the exact versions:

```
yum install kernel-headers-`uname -r` kernel-devel-`uname  
-r` gcc rsync rpm-build
```

If the exact versions are no longer available in the repository, then you will have to manually search for them on the Internet.

- On SLES 10/11 you need `kernel-syms`, `make`, `rpm-build`, `GCC4`, and `rsync`.

```
$ zypper install kernel-syms make rpm gcc rsync
```

2. To build an RPM installation package for the current kernel, run this command:

```
$ rpmbuild --rebuild iomemory-vsl-<VSL-version>.src.rpm
```

When using a .rpm source package for a non-running kernel, run this command:

```
$ rpmbuild --rebuild --define 'rpm_kernel_version <kernel-version>'
iomemory-vsl-<VSL-version>.src.rpm
```

The new RPM package is located in a directory that is indicated in the output from the `rpmbuild` command. To find it, look for the "wrote" line. In the following example, the RPM packages are located in the `/usr/src/redhat/RPMS/x86_64/` directory.

```
...
Processing files: iomemory-vsl-source-2.2.3.66-1.0.x86_64.rpm
Requires(rpmlib): rpmlib(PayloadFilesHavePrefix) <= 4.0-1
rpmlib(CompressedFileNames) <= 3.0.4-1
Obsoletes: iodrive-driver-source
Checking for unpackaged file(s): /usr/lib/rpm/check-files
/var/tmp/iomemory-vsl-2.2.3.66-root
Wrote:
/usr/src/redhat/RPMS/x86_64/iomemory-vsl-2.6.18-128.el5-2.2.3.66-1.0.x86_64.rpm
/usr/src/redhat/RPMS/x86_64/iomemory-vsl-source-2.2.3.66-1.0.x86_64.rpm
```

3. Make a note of the RPM location; you will need this information later in the installation.
4. You have now built installation packages for your distribution and kernel.
5. Copy your custom-built software installation RPM package into the directory where you downloaded the installation packages.
6. Return to the [Installing RPM Packages](#) section of this guide.

Loading the ioMemory VSL Facility (Driver)

To load the ioMemory VSL driver:

1. Run this command:

```
$ modprobe iomemory-vsl
```

NOTE The ioMemory VSL automatically loads at system boot. The IBM High IOPS Adapter is now available to the OS as `/dev/fiox`, where *x* is a letter (i.e., a, b, c, etc.).

- For SLES 10 systems, the above `modprobe` command works to manually load the driver, however to automatically start the driver at boot time with UDEV, edit the `/etc/init.d/iomemory-vsl` file to modify the `init` parameter and change `udev` to `boot.udev` as specified below.

NOTE You must install the `fio-sysvinit` package before you can edit the `/etc/init.d/iomemory-vsl` file.

After the change, the file should look like this:

```
### BEGIN INIT INFO
# Provides:          iomemory-vsl
# Required-Start:    boot.udev
```

- On SLES systems, you must also allow unsupported modules for this command to work.
 - **SLES 11:** Modify the `/etc/modprobe.d/iomemory-vsl.conf` file and uncomment the appropriate line:

```
# To allow the ioMemory VSL driver to load on SLES11,
uncomment below
allow_unsupported_modules 1
```

- **SLES 10:** Modify the `/etc/sysconfig/hardware/config` file so the `LOAD_UNSUPPORTED_MODULES_AUTOMATICALLY` sysconfig variable is set to `yes`, for example:

```
LOAD_UNSUPPORTED_MODULES_AUTOMATICALLY=yes
```

2. To confirm the IBM High IOPS Adapter is attached, run the `fio-status` utility from the command line. The output lists each drive and its status (attached or not attached).

Attention If the IBM High IOPS Adapter is not automatically attaching, check the `/etc/modprobe.d` files to see if the `auto_attach` option is turned off (set to 0).

Controlling ioMemory VSL Loading

You can control driver loading either through the init script or through udev.

In newer Linux distributions, users can rely on the udev device manager to automatically find and load drivers for their installed hardware at boot time, though udev can be disabled and the init script used in nearly all cases. For older Linux distributions without this functionality, users must rely on a boot-time init script to load needed drivers. IBM provides an init script in `/etc/init.d/iomemory-vsl` to load the VSL driver in SLES10 distributions.

Using the init Script

On systems where udev loading of the driver doesn't work, or is disabled, the init script may be enabled to load the driver at boot. On some distros it may be enabled by default.

NOTE The init Script is part of the `fio-sysvinit` package, which must be installed before you can enable it.

You can disable this loading of the ioMemory VSL with the following command:

```
$ chkconfig --del iomemory-vsl
```

To re-enable the ioMemory VSL loading in the init script, use the following command:

```
$ chkconfig --add iomemory-vsl
```

For further details on the init script, see [Using the Init Script](#) later in this section.

Using udev

On systems that rely on udev to load drivers, users need to modify an ioMemory VSL's options file if they want to prevent udev from auto-loading the ioMemory VSL at boot time. To do this, locate and edit the `/etc/modprobe.d/iomemory-vsl.conf` file that already has the following line:

```
# blacklist iomemory-vsl
```

To disable loading, remove the `"#"` from the line and save the file.

With the blacklist command in place, restart Linux. The ioMemory VSL will not be loaded by udev.

To restore the udev-loading of the ioMemory VSL, replace the `"#"` to comment out the line.

On either udev or init script systems

Users can disable the loading of the ioMemory VSL at boot time, and thus prevent the auto-attach process for diagnostic or troubleshooting purposes on either udev or init script systems. Follow the steps in the [Disabling Auto-Attach](#) section to disable or re-enable the auto-attach functionality.

Alternatively, you can prevent the ioMemory VSL from loading by appending the following parameter at the kernel command line of your boot loader:

```
iodrive=0
```

However, this method is not preferred as it prevents the ioMemory VSL from functioning at all, thus limiting the amount of troubleshooting you can perform.

Using the Init Script

The ioMemory install process places an init script in `/etc/init.d/iomemory-vsl`. In turn, this script uses the setting options found in the options file in `/etc/sysconfig/iomemory-vsl`. The options file must have `ENABLED` set (non-zero) for the init script to be used:

```
ENABLED=1
```

The options file contains documentation for the various settings---two of which, `MOUNTS` and `KILL_PROCS_ON_UMOUNT`, are discussed further in the [Handling ioMemory VSL Unloads](#) section later in this document.

Mounting Filesystems

Because the ioMemory VSL does not load by the standard means (in the `initrd`, or built into the kernel), using the standard method for mounting filesystems (`/etc/fstab`) for filesystems hosted on the IBM High IOPS Adapter does not work. To set up auto-mounting of a filesystem hosted on an IBM High IOPS Adapter:

1. Add the filesystem mounting command to `/etc/fstab` as normal.
2. Add the 'noauto' option to `etc/fstab` as in the two following sample entries.

```
/dev/fioa /mnt/fioa ext3 defaults,noauto 0 0
/dev/fiob1 /mnt/ioDrive ext3 defaults,noauto 0 0
```

(where the `a` in `fioa` can be `a`, `b`, `c`, etc., depending on how many IBM High IOPS Adapters you have installed in the system).

To have the init script mount these drives after the ioMemory VSL is loaded and unmounted and before the

ioMemory VSL is unloaded, add a list of mount points to the options file using the procedure documented there.

For the filesystem mounts shown in the earlier example, the line in the options file would look like this:

```
MOUNTS="/mnt/fioa /mnt/iodrive"
```

Handling ioMemory VSL Unloads

Special consideration must be taken during ioMemory VSL unload time. By default, the init script searches for any processes holding open a mounted filesystem and kills them, thus allowing the filesystem to be unmounted. This behavior is controlled by the option `KILL_PROCS_ON_UMOUNT` in the options file. If these processes are not killed, the filesystem cannot be unmounted. This may keep the ioMemory VSL from unloading cleanly, causing a significant delay on the subsequent boot.

Setting the ioMemory VSL Options

This section explains how to set ioMemory VSL options. For more information about setting specific options, see [Appendix C- Using Module Parameters](#).

One-Time Configuration

ioMemory VSL options can be set at install time, on the command line of either `insmod` or `modprobe`. For example, to set the `auto_attach` ioMemory VSL option to 0, run the command:

```
$ modprobe iomemory-vsl auto_attach=0
```

This option takes effect only for this load of this ioMemory VSL; subsequent calls to `modprobe` or `insmod` will not have this option set.

Persistent Configuration

To maintain a persistent setting for an option, add the desired option to `/etc/modprobe.d/iomemory-vsl.conf` or a similar file. To prevent the IBM High IOPS Adapters from auto-attaching, add the following line to the `iomemory-vsl.conf` file:

```
options iomemory-vsl auto_attach=0
```

This ioMemory VSL option then takes effect for every subsequent ioMemory VSL load, as well as on autoload of the ioMemory VSL during boot time.

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 [fio-status](#) command-line utility.

If the output shows that the device is running in minimal mode, download the latest firmware from <http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5083174>, then use the IBM High IOPS SSD Management Application software or the [fio-update-iocdrive](#) utility to upgrade the firmware.

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

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.

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 before installing the new devices.

Consult the *Release Notes* for this ioMemory VSL release for any upgrade considerations.

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 host after you upgrade the device(s). Just restarting the virtual machine won't apply the change.

Using the Device as Swap

To safely use the IBM High IOPS Adapter as swap space requires passing the `preallocate_memory` kernel module parameter. The recommended method for providing this parameter is to add the following line to the `/etc/modprobe.d/iomemory-vsl.conf` file:

```
options iomemory-vsl preallocate_memory=1072,4997,6710,10345
```

- Where 1072,4997,6710,10345 are serial numbers obtained from [fio-status](#).

A 4K sector size format is required for swap—this reduces the ioMemory VSL memory footprint to reasonable levels. Use [fio-format](#) to format the IBM High IOPS Adapter with 4k sector sizes.

NOTE Be sure to provide the serial numbers for the IBM High IOPS Adapter, not an adapter, when applicable.

NOTE The `preallocate_memory` module parameter is necessary to have the device usable as swap space. See

[Appendix C- Using Module Parameters](#) for more information on setting this parameter.

Attention You must have enough RAM available to enable the IBM High IOPS Adapter with pre-allocation enabled for use as swap. Attaching an IBM High IOPS Adapter, with pre-allocation enabled, without sufficient RAM may result in the loss of user processes and system instability.

Consult the *Release Notes* for RAM requirements with this version of the ioMemory VSL.

NOTE The `preallocate_memory` parameter is recognized by the ioMemory VSL at load time, but the requested memory is not actually allocated until the specified device is attached.

Using the Logical Volume Manager

The Logical Volume Manager (LVM) volume group management application handles mass storage devices like IBM High IOPS Adapters if you add the IBM High IOPS Adapter as a supported type:

1. Locate and edit the `/etc/lvm/lvm.conf` configuration file.
2. Add an entry similar to the following to that file:

```
types = [ "fio", 16 ]
```

The parameter "16" represents the maximum number of partitions supported by the device.

NOTE If using LVM or MD, do not use `udev` to load the ioMemory VSL driver. The init script will ensure that the LVM volumes and MD devices are detached before attempting to detach the IBM High IOPS Adapter.

Configuring RAID

You can configure two or more IBM High IOPS Adapters into a RAID array using software-based RAID solutions.

NOTE If you are using RAID1/Mirrored and one device fails, be sure to run `fio-format` on the replacement device (not the existing, good device) before rebuilding the RAID. Following are some examples of some common RAID configurations using the `mdadm` utility.

Attention The Linux kernel RAID 5 implementation performs poorly at high data rates. This is an issue in the Linux kernel. Alternatives include using RAID 10, or possibly a third-party RAID stack.

RAID 0

To create a striped set, where `fioa` and `fiob` are the two IBM High IOPS Adapters you want to stripe, run this command:

```
$ mdadm --create /dev/md0 --chunk=256 --level=0 --raid-devices=2 /dev/fioa
/dev/fiob
```

Making the Array Persistent (Existing after Restart)

NOTE On some versions of Linux, the configuration file is in `/etc/mdadm/mdadm.conf`, not `/etc/mdadm.conf`.

Inspect `/etc/mdadm.conf`. If there are one or more lines declaring the devices to inspect, make sure one of those lines specifies "partitions" as an option. If it does not, add a new `DEVICE` line to the file specifying "partitions" like this:

```
DEVICE partitions
```

Also add a device specifier for the fio IBM High IOPS Adapters:

```
DEVICE /dev/fio*
```

To see if any updates are needed to `/etc/mdadm.conf`, issue the following command:

```
$ mdadm --detail --scan
```

Compare the output of this command to what currently exists in `mdadm.conf` and add any needed sections to `/etc/mdadm.conf`.

NOTE For example, if the array consists of two devices, there will be three lines in the output of the command that are not present in the `mdadm.conf` file: one line for the array, and two device lines (one line for each device). Be sure to add those lines to the `mdadm.conf` so it matches the output of the command.

For further details please see the `mdadm` and `mdadm.conf` man pages for your distribution.

With these changes, on most systems the RAID 0 array will be created automatically upon restart. However, if you have problems accessing `/dev/md0` after restart, run the following command:

```
$ mdadm --assemble --scan
```

You may also want to disable udev loading of the ioMemory VSL driver, if needed, and use the init script provided for driver loading. Please see the [Using the Init Script](#) section of this guide for further details on how to use the init script.

NOTE In SLES 11, you may need to run the following commands to make sure these services are run on boot:

```
chkconfig boot.md on
chkconfig mdadm on
```

RAID 1

To create a mirrored set using the two IBM High IOPS Adapters `fioa` and `fiob`, run this command:

```
$ mdadm --create /dev/md0 --level=1 --raid-devices=2 /dev/fioa /dev/fiob
```

RAID 10

To create a striped, mirrored array using four IBM High IOPS Adapters (`fioa`, `fiob`, `fioc`, and `fiod`), run this command:

```
$ mdadm --create /dev/md0 -v --chunk=256 --level=raid10 --raid-devices=4
/dev/fioa /dev/fiob /dev/fioc /dev/fiod
```

Building a RAID10 Across Multiple Devices

In a RAID10 configuration, sets of two disks are mirrored, and then those mirrors are striped. When setting up a RAID10 across multiple IBM High IOPS Adapters, it is best to make sure that no mirror resides solely on the two IBM High IOPS Adapters that comprise an a single product (such as an IBM High IOPS Duo Adapter).

In order to get the data to lay out properly,

- Use the `--layout=n2` option when creating the RAID10 (though it should be the default)
- Ensure that no two IBM High IOPS Adapters from the same device are listed side by side.

The following sample code shows some recommended configurations.

NOTE The following commands assume that all IBM High IOPS Adapters have been freshly formatted with the `fio-format` utility.

Attention The ordering of the fiox devices is critical.

```
# 2 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 4 /dev/md0 \
/dev/fioa /dev/fioc \
/dev/fiob /dev/fiod
# Mirror groups are: fioa,fioc and fiob,fiod

# 3 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 6 /dev/md0 \
/dev/fioa /dev/fiod \
/dev/fioc /dev/fiof \
/dev/fioe /dev/fiob

# 4 Duos RAID10
```

```
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 8 /dev/md0 \
/dev/fioa /dev/fiod \
/dev/fioc /dev/fiof \
/dev/fioe /dev/fioh \
/dev/fiog /dev/fiob

# 8 Duos RAID10
$ mdadm --create --assume-clean --level=raid10 --layout=n2 -n 16 /dev/md0 \
/dev/fioa /dev/fiod \
/dev/fioc /dev/fiof \
/dev/fioe /dev/fioh \
/dev/fiog /dev/fioj \
/dev/fioi /dev/fiol \
/dev/fiok /dev/fion \
/dev/fiom /dev/fiop \
/dev/fioo /dev/fiob
```

Understanding Discard (TRIM) Support

With this version of the ioMemory VSL, Discard (also known as TRIM) is enabled by default.

Discard addresses an issue unique to solid-state storage. When a user deletes a file, the device does not recognize that it can reclaim the space. Instead the device assumes the data is valid.

Discard is a feature on newer filesystem releases. It informs the device of logical sectors that no longer contain valid user data. This allows the wear-leveling software to reclaim that space (as reserve) to handle future write operations.

Discard (TRIM) on Linux

Discard is enabled by default in the release. However, for discard to be implemented, the Linux distribution must support this feature, and discard must be turned on.

In other words, if your Linux distribution supports discard, and discard is enabled on the system, then discard will be implemented on your IBM High IOPS Adapter.

Under Linux, discards are not limited to being created by the filesystem, discard requests can also be generated directly from userspace applications using the kernel's discard ioctl.

Attention There is a known issue that ext4 in Kernel.org 2.6.33 or earlier may silently corrupt data when discard is enabled. This has been fixed in many kernels provided by distribution vendors. Please check with your kernel provider to be sure your kernel properly supports discard. For more information, see the Errata in the *Release Notes* for this version of the software.

NOTE On Linux, MD and LVM do not currently pass discards to underlying devices. Thus any ioDrive device that is part of an MD or LVM array will not receive discards sent by the filesystem.


The LVM release included in Red Hat 6.1 supports passing discards for several targets, but not all ([RHEL 6.1 documentation](#)). Please see your distribution's documents for exact details.

Maintenance

The ioMemory VSL includes software utilities for maintaining the device. You can also install SNMP as a monitoring option.

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 one of these configurations:



This table explains the information that these LEDs convey:

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

GUI Management

IBM High IOPS SSD Management Application is a free browser-based solution for managing IBM High IOPS Adapters. It is also available from <http://www.ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5083174>, but it is located in a different download location along with documentation.

The IBM High IOPS SSD Management Application can perform:

- Firmware upgrades
- Low-level formatting
- Attach and detach actions
- Device status and performance information

Command-line Utilities

Several command-line utilities are included in the installation packages for managing your IBM High IOPS Adapter:

- `fio-attach`
- `fio-beacon`
- `fio-bugreport`
- `fio-detach`
- `fio-format`
- `fio-pci-check`
- `fio-snmp-agentx`
- `fio-status`
- `fio-sure-erase`
- `fio-update-iodrive`

For more information on command-line utilities, see [Appendix A- Utilities Reference](#)

Common Maintenance Tasks

The following are the most common tasks for maintaining your IBM High IOPS Adapter using command-line utilities.

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

NOTE If you came to this section from the [Software Installation](#) section, [return](#) to that section after you uninstall previous versions of the driver and utilities.

Unloading the ioMemory VSL driver

To unload the ioMemory VSL, run this command:

```
$ modprobe -r iomemory-vsl
```

Uninstalling the ioMemory VSL Package

Versions 1.2.x

Remove prior versions of the VSL software using this command:

```
$ rpm -e iodrive-driver
```

Versions 2.x.x

With versions 2.x.x of the ioMemory VSL, you must specify the kernel version of the package you are uninstalling. Run this command to find the installed driver packages:

```
$ rpm -qa | grep -i iomemory
```

Sample output:

```
iomemory-vsl-2.6.18-194.el5-2.2.2.82-1.0
```

Uninstall the ioMemory VSL by running a command similar to this example (specify the kernel version of the driver you wish to uninstall):

```
$ rpm -e iomemory-vsl-2.6.18-194.el5-2.2.0.82-1.0
```

Disabling Auto-Attach

When the ioMemory VSL is installed, it is configured to automatically attach any devices when the ioMemory VSL is loaded. Sometimes you may want to disable the auto-attach feature. To do so:

1. Edit the following file:

```
/etc/modprobe.d/iomemory-vsl.conf
```

2. Add the following line to that file:

```
options iomemory-vsl auto_attach=0
```

3. Save the file. To re-enable auto-attach, simply edit the file and either remove that line or change it to the following:

```
options iomemory-vsl auto_attach=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 or more to complete.

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.

Disabling the ioMemory VSL

The ioMemory VSL automatically loads by default when the operating system starts. You can disable ioMemory VSL auto-load for diagnostic or troubleshooting purposes.

To disable ioMemory VSL auto-load:

1. Append the following parameter to the kernel command line of your boot loader:

```
iodrive=0
```

The ioMemory VSL won't load, so the device won't be available to users.

NOTE You can also uninstall the ioMemory VSL to keep it from loading, or move it out of the `/lib/modules/<kernel_version>` directory.

2. Proceed with troubleshooting to correct the problem. If the problem is outdated firmware, use `iodrive=1` to place the ioMemory VSL in minimal mode. You can then use `fio-update-iodrive` or the ioManager application to update the firmware.
3. Use either the `fio-attach` utility or the ioManager application to attach the ioMemory VSL to the operating system.

Performance and Tuning

IBM High IOPS Adapters provide high bandwidth, and high Input/Output 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 DVFS 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.

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 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/disable them in the BIOS interface (sometimes referred to as a Setup Utility). ACPI C-states may be part of the Advanced Configuration and Power Interface (ACPI) menu; consult your computer documentation for details.

C-States Under Linux

Newer Linux kernels have drivers that may attempt to enable ACPI C-states even if they are disabled in the BIOS. You can limit the C-state in Linux (with or without the BIOS setting) by adding the following to the kernel boot options:

```
processor.max_cstate=1
```

In this example, the maximum C-state allowed will be C1. To completely disable C-states, use a value of 0.

Setting NUMA Affinity

Servers with a NUMA (Non-Uniform Memory Access) architecture require special installation instructions in order to maximize IBM High IOPS Adapter performance. These servers include the IBM x3850 server.

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

See [Appendix E- NUMA Configuration](#) for more information on setting this affinity.

Setting the Interrupt Handler Affinity

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

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

Appendix A- Utilities Reference

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.

Utility	Purpose
<code>fio-attach</code>	Makes an IBM High IOPS Adapter available to the OS
<code>fio-beacon</code>	Lights the IBM High IOPS Adapter's external LEDs
<code>fio-bugreport</code>	Prepares a detailed report for use in troubleshooting problems
<code>fio-detach</code>	Temporarily removes an IBM High IOPS Adapter from OS access
<code>fio-format</code>	Used to perform a low-level format of an IBM High IOPS Adapter
<code>fio-pci-check</code>	Checks for errors on the PCI bus tree, specifically for IBM High IOPS Adapters
<code>fio-snmp-agentx</code>	SNMP sub-agent that implements the SNMP FUSION-IODRV-MIB for the ioMemory VSL
<code>fio-status</code>	Displays information about the device
<code>fio-sure-erase</code>	Clears or purges data from the device
<code>fio-update-iodrive</code>	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. This creates a block device in `/dev` named `fiox` (where `x` is `a`, `b`, `c`, etc.). You can then partition or format the IBM High IOPS Adapter, or set it up as part of a RAID array. The command displays a progress bar and percentage as it operates.

NOTE In most cases, the ioMemory VSL 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's `auto_attach` parameter to 0.

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.

You can specify multiple IBM High IOPS Adapters. For example, `/dev/fct1 /dev/fct2` indicates the second and third IBM High IOPS Adapters installed on the system. You can also use a wildcard to indicate all IBM High IOPS Adapters on the system. For example, `/dev/fct*`

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

fio-beacon

Description

Lights the IBM High IOPS Adapter's LEDs 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 installed on the system.

Options	Description
<code>-0</code>	Off: (Zero) Turns off the three LEDs
<code>-1</code>	On: Lights the three LEDs
<code>-p</code>	Prints the PCI bus ID of the device at <code><device></code> to standard output. Usage and error information 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. The results are saved in the `/tmp` directory in the file that indicates the date and time the utility was run.

Example:

```
/tmp/fio-bugreport-20100121.173256-sdv9ko.tar.bz2
```


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 <http://www.ibm.com/systems/support> for assistance in troubleshooting.

You are prompted to send an e-mail describing the problem to `support@fusionio.com` with the bug report file attached.

NOTE Disregard the message about contacting Fusion-io Support; instead, contact <http://www.ibm.com/systems/support>.

Sample Output

```
-bash-3.2# fio-bugreport
Collecting fio-status -a
Collecting fio-status
Collecting fio-pci-check
Collecting fio-pci-check -v
Collecting fio-read-lebmap /dev/fct0
Collecting fio-read-lebmap -x /dev/stdout/dev/fct0
Collecting fio-read-lebmap -t /dev/fct0
Collecting fio-get-erase-count/dev/fct0
Collecting fio-get-erase-count -b /dev/fct0
Collecting lspci
Collecting lspci -vvvvv
Collecting lspci -tv
Collecting messages file(s)
Collecting procfusion file(s)
Collecting lsmod
Collecting uname -a
Collecting hostname
Collecting sar -r
Collecting sar
Collecting sar -A
Collecting syslog file(s)
Collecting proc file(s)
Collecting procirq file(s)
Collecting dmidecode
Collecting rpm -qa iodrive*
Collecting find /lib/modules
Please send the file /tmp/fio-bugreport-20090921.173256-sdv9ko.tar.bz2
along with your bug report to support@fusionio.com The file is in the /tmp
directory.
```

For example, the filename for a bug report file named
/tmp/fiobugreport-20090921.173256-sdvk0.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 and removes the corresponding `fctx` IBM High IOPS Adapter block device from the OS. The `fio-detach` utility waits until the device completes all read/write activity before executing the detach operation. By default, the command also displays a progress bar and percentage as it completes the detach.

Attention Before using this utility, ensure that the device you want to detach is **NOT** currently mounted and in use.

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.

You can specify multiple IBM High IOPS Adapters. For example, `/dev/fct1 /dev/fct2` indicates the second and third IBM High IOPS Adapters installed on the system. You can also use a wildcard to indicate all IBM High IOPS Adapters on the system. For example, `/dev/fct*`

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

Notes

With this version of ioMemory VSL, 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 part of a software RAID (0,1,5) volume, is mounted, or some process has the device open.

The tools `fuser`, `mount`, and `lsof` can be helpful in determining what is holding the device open.

fio-format

Description

NOTE IBM High IOPS Adapters ship pre-formatted, so `fio-format` is generally not required except to change the logical size or block size of a device, or to erase user data on a device. To ensure the user data is truly erased, use [fio-sure-erase](#).

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.

- NOTE** Using a larger block (sector) size, such as 4096 bytes, can significantly reduce worst-case ioMemory VSL host memory consumption. However, some applications are not compatible with non-512-byte sector sizes.
- NOTE** If you do not include the `-s` option, the device size defaults to the advertised capacity. If used, the `-s` option must include the size or percentage indicators.

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 installed on the system.

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

fio-pci-check

Description

Checks for errors on the PCI bus tree, specifically for 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]
```

Options	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.
-i	Print the device serial number. This option is invalid when the ioMemory VSL is loaded.
-r	Force the link to retrain.
-v	Verbose: Print extra data about the hardware.

fio-snmp-agentx**Description**

This utility is an SNMP sub-agent that implements the SNMP FUSION-IODRV-MIB for the ioMemory VSL. `fio-snmp-agentx` communicates with the SNMP master agent via the agentx protocol.

Syntax

```
fio-snmp-agentx [options]
```

Option	Description
-f	Force the sub-agent to run in the foreground instead of as a daemon.
-l <log file>	Log file to use.
-s	Send errors to stderr instead of to syslog.

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 loaded. If no ioMemory VSL 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 installed on the system.

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

Options	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.
-f j	Format JSON: creates the output in JSON format.
-f x	Format XML: creates the output in XML format.
-u	Show unavailable fields. Only valid with -f j or -f x.
-U	Show unavailable fields and details why. Only valid with -f j or -f x.
-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:

Starting with version 3.1.1 and later, the standard formatting of `fio-status` output has changed. 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 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 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 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-sure-erase

Attention As a best practice, do not use this utility if there are any IBM High IOPS Adapters installed in the system that you do not want to clear or purge. First remove any devices that you do not want to accidentally erase. Once the data is removed with this utility it is gone forever. **It is not recoverable.**

Attention Before you use this utility, be sure to back up any data that you wish to preserve.

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

NOTE If the device is in Read-only mode, perform a format using `fio-format` before running `fio-sure-erase`. If the device is in Minimal mode, then `fio-sure-erase` cannot erase the device. Updating the firmware may take the device out of Minimal Mode. If the device remains in Minimal mode,

contact Customer Support at <http://www.ibm.com/systems/support> for further assistance.

In order to run `fio-sure-erase`, the block device **must be detached**. See the [fio-detach](#) section for more information.

Description

The `fio-sure-erase` is a command-line utility that securely removes data from IBM High IOPS Adapters. It complies with the "Clear" and "Purge" level of destruction from the following standards:

1. DOD 5220.22-M – Comply with instructions for Flash EPROM
2. NIST SP800-88– Comply with instructions for Flash EPROM

See below for more information on Clear and Purge support.


Syntax

```
fio-sure-erase [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. Use [fio-status](#) to view this device node.

NOTE Products with Multiple Devices:

`fio-sure-erase` works on individual IBM High IOPS Adapters. For example, if you are planning to purge an IBM High IOPS Duo Adapter, you will need to perform this operation on each of the product's two IBM High IOPS Adapters.

Options	Description
<code>-p</code>	Purge instead of Clear: performs a write followed by an erase. For more information on Purge, see below. <div>  Purging the device may take hours to accomplish, depending on the size of the device that needs to be purged. </div>
<code>-y</code>	No confirmation: do not require a yes/no response to execute the utility.
<code>-q</code>	Quiet: do not display the status bar.

NOTE If you run `fio-sure-erase` with no options, a Clear is performed. For more information, see below.

Each block of memory consists of uniform 1 bits or 0 bits.

Clear Support

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

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

The following describes the steps taken in the Clear operation:

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

Purge Support

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

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

The following describes the steps taken in the Purge operation:

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

fio-update-iodrive

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

Description

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 driver after each firmware upgrade step. Otherwise the on-drive 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.

NOTE The default action (without using the `-d` or `-s` 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.

NOTE You must detach all IBM High IOPS Adapters before updating the firmware.

Attention Upgrade Path:

There is a specific upgrade path that you must take when upgrading IBM High IOPS Adapter. Consult the *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.

To update one or more specific devices:

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

Syntax

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

where `<iodrive_version.fff>` is the path and firmware archive file provided by IBM. The default path is `/usr/share/fio/firmware`. This parameter is required.

Options	Description
<code>-d</code>	Updates the specified devices (by <code>fctx</code> , where <code>x</code> is the number of the device shown in <code>fio-status</code>). If this option is not specified, all devices are updated. Attention Use the <code>-d</code> or <code>-s</code> options with care, as updating the wrong IBM High IOPS Adapter could damage your device.
<code>-f</code>	Force upgrade (used primarily to downgrade to an earlier firmware version). If the ioMemory VSL is not loaded, this option also requires the <code>-s</code> option. Attention Use the <code>-f</code> option with care, as it could damage your card.
<code>-l</code>	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.
-s	Updates the devices in the specified slots using '*' as a wildcard for devices. The slots are identified in the following PCIe format (as shown in lspci): <div style="border: 1px solid black; padding: 10px; margin: 10px 0;">[[[<domain>]:<bus>]:[<slot>][.<func>]]</div>

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

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.

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 pro-actively 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. The IBM High IOPS SSD Management Application and the `fiio-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 the [Health Monitoring Techniques](#) 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 3 ioMemory devices in this system
Fusion-io driver version: 3.0.6 build 364

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
```

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

SNMP/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

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 editing the `/usr/modprobe.d/iomemory-vsl.conf` file and changing their values.

Each module parameter in the configuration file must be preceded by options `iomemory_vsl`. The `/etc/modprobe.d/iomemory-vsl.conf` file has some example parameters that are commented out. You may use these examples as templates and/or uncomment them in order to use them.

NOTE These changes must be completed before the ioMemory VSL is loaded in order to take effect.

Module Parameter	Default (min/max)	Description
<code>auto_attach</code>	1	1 = Always attach the device on driver load. 0 = Don't attach the device on driver load.
<code>fio_dev_wait_timeout_secs</code>	30	Number of seconds to wait for <code>/dev/fio*</code> files to show up during driver load. For systems not using <code>udev</code> , this should be set to 0 to disable the timeout and avoid an unneeded pause during driver load.
<code>force_minimal_mode</code>	0	1 = Force minimal mode on the device. 0 = Do not force minimal mode on the device.
<code>numa_node_override</code>	Nothing Selected	A list of <code><affinity specification></code> couplets that specify the affinity settings of all devices in the system. Each item in the couplet is separated by a colon, and each couplet set is separated by a comma. Where each <code><affinity specification></code> couplet has the following syntax: <div style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <code><fct-number> : <node-number></code> </div> See Appendix E- NUMA Configuration for more information on using this parameter.
<code>parallel_attach</code>	1	1 = Enable parallel attach of multiple devices. 0 = Disable parallel attach of multiple devices.
<code>preallocate_memory</code>	No devices selected	For the selected devices, pre-allocate all memory necessary to have the drive usable as swap space. Where the <code><value></code> for this parameter is a comma-separated list of device serial numbers.

<code>tintr_hw_wait</code>	0 (0, 255)	Interval (microseconds) to wait between hardware interrupts. Also known as interrupt coalescing. 0 is off.
<code>use_workqueue</code>	3 (1 or 3)	Linux only: 3 = use standard OS I/O elevators; 0 = bypass.

NOTE Other than `preallocate_memory`, module parameters are global — they apply to all IBM devices in the computer.

Appendix D- SMI-S Interface

Attention The SMI-S provider has a new CIM namespace: `root/fio`

Introduction to the SMI-S Interface

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 Fusion-io.

It is assumed that the reader is versed in WBEM, SMI-S and DMTF standards. This document and associated model may change at any time as feedback is received.

References

CIM Schema v2.22

http://www.dmtf.org/standards/cim/cim_schema_v2220

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 Profile

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/>

Installing the SMI-S WMI Provider on Linux

The IBM SMI-S provider implements a standard WBEM interface based on DMTF and SNIA standards for remote management of IBM High IOPS Adapters. The provider is a CMPI-based provider and should work with popular CIMOMs including SFCB, OpenPegasus, and OpenWBEM. We also have a version that supports WMI on Windows.

In order to use the provider, a CIMOM must be installed and configured. The provider and associated MOF files must then be installed and registered with the CIMOM. The MOF files define the CIM objects available from the SMI-S provider.

Initially, the provider has been tested with SFCB on Linux and WMI on Windows.

1. Install the SFCB CIM broker (CIMOM).

- On Debian-like:

```
$ sudo apt-get install sfcbl sfcbl-test wbmcli sblim-cmpi-base  
rsync
```

- Others: Install sblim-sfcc, sblim-sfcc-devel, cim-schema-2.21.0, sblim-sfcbl, sblim-indication-helper, sblim-cmpi-base and sblim-testsuite

NOTE RPMs available for SLES, RHEL, and others on OpenSUSE Build Service.

2. Configure SFCB: Copy the file (sfcbl.cfg) to /etc/sfcbl
3. Install SMI-S provider: Install fio-smis package from IBM distribution and copy /usr/lib/fio/libfiosmis.so to /usr/lib/sfcbl (or /usr/lib64/sfcbl as appropriate).
4. Register the SMI-S provider with SFCB:

```
cd /usr/share/fio/cim  
  
sh /usr/share/sblim-cmpi-base/provider-register.sh -r  
fiosmis.registration cimv222-dmtf.mof
```

5. Restart SFCB:

```
$ /etc/init.d/sfcbl restart
```

Linux Testing

The `wbemcli` utility can be used to test the SMI-S provider.

1. Query the provider for the ioMemory VSL version and the firmware version for each IBM High IOPS Adapter in the system:

```
$ wbemcli -nl ei http://localhost:5988/root/cimv2:FIO_SoftwareIdentity
```

The output should look something like this (values may change as development continues):

```
localhost:5988/root/cimv2:FIO_SoftwareIdentity.InstanceID="Fusion-io
drive driver"
-InstanceID="Fusion-io drive driver"
-TargetTypes=
-OtherExtendedResourceTypeDescription=
-MinExtendedResourceTypeRevisionNumber=
. . .
```

2. Query the SMI-S provider for each IBM High IOPS Adapter's health:

```
wbemcli -nl ei http://localhost:5988/root/cimv2:FIO_IoMemoryPort
```

The output should look something like this (values may change as development continues):

```
localhost:5988/root/cimv2:FIO_IoMemoryPort.DeviceID="fct1",CreationClassN
. .
```

3. Query capacity and usage counters of a specific IBM High IOPS Adapter (in this case fct0):

```
$ wbemcli -nl gi
http://localhost:5988/root/cimv2:FIO_SSDStatistics.InstanceID="fct0"
```

The output should look something like this (values may change as development continues):

```
localhost:5988/root/cimv2:FIO_SSDStatistics.InstanceID="fct0"
-InstanceID="fct0"
-WriteOperations=0
-ReadOperations=6887
-PhysicalMBytesWritten=1523769
. . .
```

The Linux SMI-S provider can be tested remotely with the `wbemcli` utility by replacing `localhost` in the examples above with the hostname or IP address of the remote host. This method cannot be used to test the Windows SMI-S provider remotely, however, since (of course) Windows doesn't follow the emerging standard.

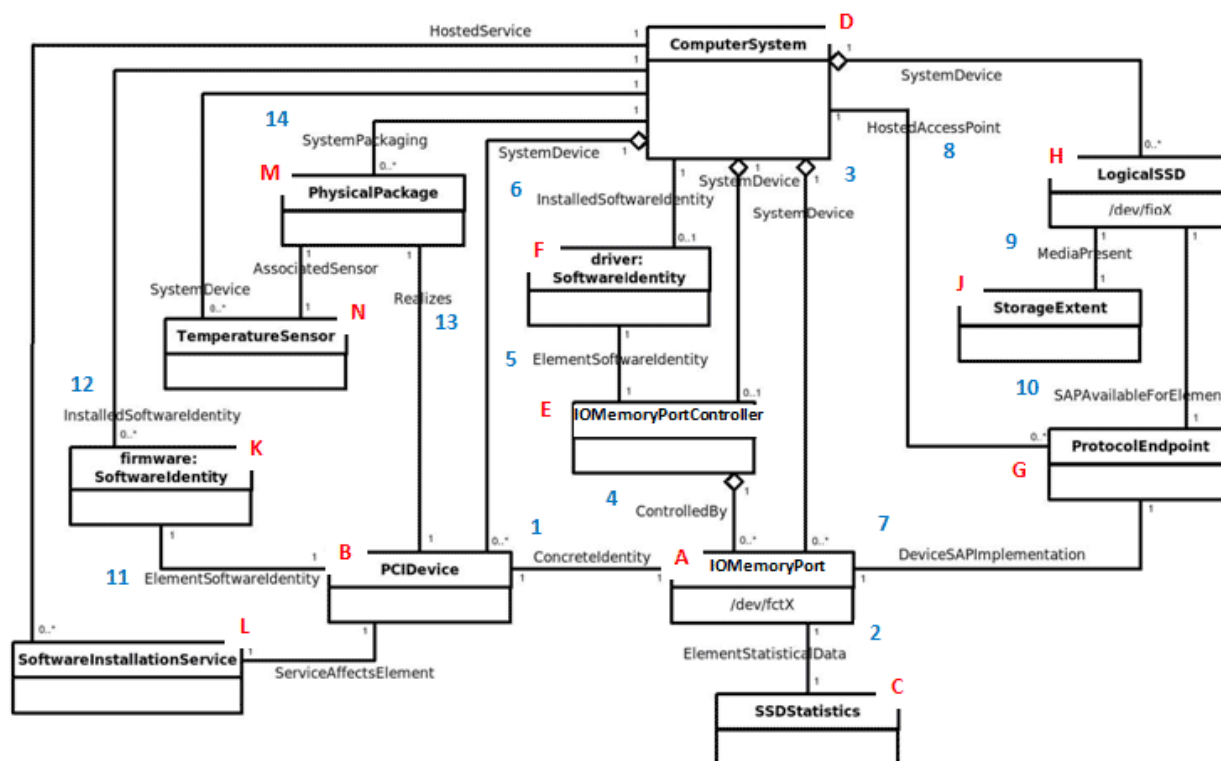
The SMI-S provider indications can be tested as well.

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.

`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 `IOMemoryPortController` 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 `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 ioDrive 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 ioDrive, via Realizes.

PhysicalPackage

One instance of PhysicalPackage exists for each discrete, physical ioDrive 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.

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.

FIO_AlertIndication

Property	Value
IndicationIdentifier	See below for each type
IndicationTime	Timestamp when sent
AlertingManagedElement	IoMemoryPort.DeviceID=<device ID>
AlertingElementFormat	CIMObjectPath (2)
AlertType	Device Alert (5)
PerceivedSeverity	See below for each type
ProbableCause	See below for each type
SystemCreationClassName	"FIO_AlertIndication"
SystemName	<hostname>
ProviderName	"fiosmis"
CorrelatedIndications	Not used
Description	Class description
OtherAlertType	Not used
OtherSeverity	Not used
ProbableCauseDescription	Not used
EventID	Same as IndicationIdentifier
OwningEntity	<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"
PerceivedSeverity	Degraded/Warning (3)

ProbableCause	Threshold Crossed (52) Temperature Unacceptable (51) Power Problem (36)
---------------	---

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"
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"
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"
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	<mfr>":"<hostname>":temperature"
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"
PerceivedSeverity	Major (6)
ProbableCause	Other (1)

Appendix E- NUMA Configuration

About NUMA Architecture

Servers with a NUMA (Non-Uniform Memory Access) architecture require special installation instructions in order to maximize IBM High IOPS Adapter performance. This includes the IBM x3850 server.

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

NOTE The example below shows the final implementation of custom affinity settings. This implementation required an analysis of the specific system, including the system architecture, type and number of IBM High IOPS Adapters installed, and the particular PCIe slots that were used. Your particular circumstances will require a custom analysis of your set-up. This analysis requires understanding of your system's NUMA architecture compared to your particular installation.

Your actual settings may be different than the example below, depending on your server configuration. In order to create the correct settings for your specific system, use `fio-status` to list all of the devices (fct numbers). Next, use `fio-beacon` to identify each of the devices in their respective PCIe slots. Then use the example below of setting the `numa_node_override` parameter as a template and modify it for your particular system.

Configuring your IBM High IOPS Adapters for servers with NUMA architecture requires the use of the `numa_node_override` parameter by modifying the `iomemory-vsl.conf` file.

numa_node_override Parameter

The `numa_node_override` parameter is a list of `<affinity specification>` couplets that specify the affinity settings of all devices in the system. Each item in the couplet is separated by a colon, and each couplet set is separated by a comma.

Syntax:

```
numa_node_override=<affinity specification>[,<affinity specification>...]
```

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


```
<fct-number>:<node-number>
```

Simple Example:

```
numa_node_override=fct4:1,fct5:0,fct7:2,fct9:3
```

Has the effect of creating :

Device	Node/Group	Processor Affinity
fct4	node 1	all processors in node 1
fct5	node 0	all processors in node 0
fct7	node 2	all processors in node 2
fct9	node 3	all processors in node 3

Advanced Configuration Example

This sample server has 4 NUMA nodes with 8 hyper-threaded cores per node (16 logical processors per node, a total of 64 logical processors in the system). This system also uses the expansion configuration and has 11 PCIe expansion slots. During system boot, the system's BIOS will assign PCIe slots 1-6 to NUMA node 2 and PCIe slots 7-11 to NUMA node 0. NUMA nodes 1 and 3 will have no assigned PCIe slots. This creates a load balancing problem in the system when IBM High IOPS Adapters are under heavy traffic. Specifically, during these periods of high use, half of the CPUs in the system will sit idle while the other half of the CPUs are 100% utilized, thus limiting the throughput of the IBM High IOPS Adapters.

To avoid this problem, you must manually configure the affinity of the IBM High IOPS Adapters using the `numa_node_override` parameter to distribute the work load across all NUMA nodes. This parameter will override the default behavior of the ioMemory VSL driver. For more information about the `numa_node_override` parameter, refer to the syntax explanation above.

What follows is an example of how to manually configure 10 IBM High IOPS Duo Adapters (each with two IBM High IOPS Adapters). Slot 1 is a Generation 1 PCI-e slot, so it is not compatible with an IBM High IOPS Duo Adapter. Thus we can fill slots 2-11 with IBM High IOPS Duo Adapters.

NOTE Because each IBM High IOPS Duo Adapter has two IBM High IOPS Adapters, there are two device numbers for each IBM High IOPS Duo Adapter (one for each IBM High IOPS Adapter). There will therefore be two device numbers for each slot.

When the system boots, the default BIOS NUMA node assignments are:

BIOS Assigned NUMA Node	PCI-e Slots	FCT device numbers	Processor Affinity
0	7-11	8,9,13,14,18,19,23,24,28,29	all processors in the node
1	none	none	none

2	2-6	135,136,140,141,145,146,150,151,155,156	all processors in the node
3	none	none	none

Here, the BIOS creates a load imbalance by assigning the cards to only two NUMA nodes in the system. In order to balance the work load, we want to make the following manual settings:

Assigned NUMA Node	PCI-e Slots	FCT device numbers	Processor Affinity
0	7-9	8,9,13,14,18,19	all processors in the node
1	10-11	23,24,28,29	all processors in the node
2	2-3	135,136,140,141	all processors in the node
3	4-6	145,146,150,151,155,156	all processors in the node

In order to configure the ioMemory VSL driver with these override settings, set the `numa_node_override` parameter with the following string:

```
numa_node_override=fct8:0,fct9:0,fct13:0,fct14:0,fct18:0,fct19:0,
fct23:1,fct24:1,fct28:1,fct29:1,fct135:2,fct136:2,fct140:2,fct141:2,
fct145:3,fct146:3,fct150:3,fct151:3,fct155:3,fct156:3
```

NOTE The above example contains line breaks for formatting purposes. There would be no line breaks in a real implementation of the `numa_node_override` parameter.

Appendix F- Upgrading Devices from VSL 2.x to 3.x

This version of the ioMemory VSL supports new features, including the latest generation of IBM High IOPS architecture and improved Flashback Protection™. These features require the latest version of the The page Information does not exist. 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 ioMemory VSL 2.2.3 with IBM High IOPS Duo Adapters previously installed, and you want to install new 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. 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.

NOTE The point at which the user data is destroyed, and the upgrade cannot be reversed, is when the IBM High IOPS Adapter is formatted using the 3.x VSL `fiio-format` utility. It may take an extended period of time to format each device, depending on the wear on the device.

Overformatting Not Supported

The `-o overformat` option is not supported in the 3.x VSL. All upgraded IBM High IOPS Adapters can only have a maximum capacity equal to the maximum advertised capacity, regardless of whether the device was overformatted before upgrading. Contact Customer Support, prior to making any changes to your environment, if you need assistance determining the implications of the upgrade to 3.1.x.

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. The following is the minimum upgrade path that must be followed. Upgrade the ioMemory VSL software on the system (and **upgrade the firmware** to the compatible version for each version of the software) in this order:

1.2.8 -> 2.2.3 -> 3.1.x

Visit <http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5083174> for all of the required software and firmware versions.

Upgrade Procedure

Be sure to follow the upgrade path listed above. Make sure that all previously installed IBM High IOPS Duo Adapters are updated with the appropriate 2.2.3-compatible firmware.

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

1. Prepare each existing IBM High IOPS Duo 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 IBM High IOPS Duo Adapters, for example:

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

For more information, see [fio-detach](#)

2. Unload the current ioMemory VSL driver, for example:

```
$ modprobe -r iomemory-vsl
```

3. Uninstall the 2.x ioMemory VSL software

- a. To uninstall the software, you must specify the kernel version of the package you are uninstalling. Run the appropriate command to find the installed packages:

```
$ rpm -qa | grep -i iomemory
```

Sample output:

```
iomemory-vsl-2.6.18-194.el5-2.2.2.82-1.0
```

- b. Uninstall the ioMemory VSL by running a command similar to this example (specify the kernel version of the package you wish to uninstall):

```
$ rpm -e iomemory-vsl-2.6.18-194.el5-2.2.0.82-1.0
```

- c. Uninstall the utilities:

```
$ rpm -e fio-util fio-smip-agentx fio-common fio-firmware  
iomanager-gui iomanager-jre libfio libfio-doc libfusionjni  
fio-sysvinit fio-smis fio-smip-mib libfio-deb
```

4. Install the new VSL and related packages.

- a. Download the ioMemory VSL binary package for your kernel and all supporting packages at <http://www.ibm.com/support/entry/portal/docdisplay?lnidocid=MIGR-5083174>

NOTE If you don't see a binary for your kernel, follow the instructions in the [Building the ioMemory VSL from Source](#) section. To see your current kernel version, run:

```
uname -r
```

- b. Install the ioMemory VSL and utilities using the appropriate commands:

```
rpm -Uvh iomemory-vsl-<kernel-version>-<VSL-version>.x86_64.rpm  
rpm -Uvh lib*.rpm  
rpm -Uvh fio*.rpm
```

See the section [Installing RPM Packages](#) for full instructions on installing those packages.

- c. Reboot the system.

5. 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 the [fio-update-iodrive](#) utility reference in the appendix.

Sample syntax:

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

Where `<iodrive_version.fff>` is the path to the firmware archive. This command will update all of the devices to the selected firmware. If you wish to update specific devices, consult the [utility reference](#) for more options.

6. Reboot the system
7. Load the ioMemory VSL, for example:

```
$ modprobe iomemory-vsl
```

For more information, see [Loading the ioMemory VSL Facility \(Driver\)](#)

NOTE If run, `fio-status` will warn that the upgraded devices are missing a lebmap. This is expected, and will be fixed in the next step.

Attention Destructive Step:

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>.

8. 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.

9. Attach all IBM High IOPS Duo Adapters, for example:

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

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

```
fio-status -a
```



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

Appendix G- Open Source License Acknowledgement

This product may include and/or use the following open source software:

iniparser configuration file parser library:

Copyright (c) 2000-2007 by Nicolas Devillard.

Refer to the [iniparser Copyright and License](#) section below.

jsoncpp c++ json parser library:

jsoncpp is Public Domain where such is allowed; elsewhere the following holds:

Copyright (c) 2007-2010 by Baptiste Lepilleur

Refer to the [jsoncpp copyright and license](#) section below.

LUA scripting language library:

Copyright (C) 1994-2006 Lua.org, PUC-Rio.

Refer to the [Lua License](#) section below.

OpenSSL:

OpenSSL (version 1.0.0d):

Copyright (c) 1998-2011 The OpenSSL Project.

Refer to the [OpenSSL License](#) and [Original SSLeay License](#) sections below.

TinyXML C/C++ XML processing library:

TinyXML is licensed under the zlib license.

Refer to the [zlib Copyright and License](#) below.

tree.h tree structure description header: ('tree.h')

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

Refer to the [tree.h Copyright and License](#) section below.

util-linux Libraries: ('uuid', 'blkid')

- blkid - esx only:
 - Copyright (C) 1996, 1997 Theodore Ts'o.
 - libblkid is subject to the terms of the LGPL, version 2.
- uuid - linux only:
 - Copyright (C) 1996, 1997, 1998, 1999, 2007 Theodore Ts'o.
 - libuuid is subject to the terms of the 3-clause BSD license.

Refer to the [GNU LESSER GENERAL PUBLIC LICENSE](#) and [libuuid Copyright and License \(3-clause BSD\)](#) sections below.

zlib compression library:

version 1.2.5.2, April 19th, 2010

Copyright (C) 1995-2010 Jean-loup Gailly and Mark Adler

Refer to the [zlib Copyright and License](#) section below.

GNU LESSER GENERAL PUBLIC LICENSE

GNU LESSER GENERAL PUBLIC LICENSE
Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.
51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts
as the successor of the GNU Library Public License, version 2, hence
the version number 2.1.]

Preamble

The licenses for most software are designed to take away your
freedom to share and change it. By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it. You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations
below.

When we speak of free software, we are referring to freedom of use,
not price. Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights. These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you. You must make sure that they, too, receive or can get the source
code. If you link other code with the library, you must provide

complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must

be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control

compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square

root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the

source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one

of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this

License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free

programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

iniparser Copyright and License

Copyright (c) 2000-2007 by Nicolas Devillard.
MIT License

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

jsoncpp Copyright and License

The JsonCpp library's source code, including accompanying documentation, tests and demonstration applications, are licensed under the following conditions...

The author (Baptiste Lepilleur) explicitly disclaims copyright in all jurisdictions which recognize such a disclaimer. In such jurisdictions, this software is released into the Public Domain.

In jurisdictions which do not recognize Public Domain property (e.g. Germany as of 2010), this software is Copyright (c) 2007-2010 by Baptiste Lepilleur, and is released under the terms of the MIT License (see below).

In jurisdictions which recognize Public Domain property, the user of this software may choose to accept it either as 1) Public Domain, 2) under the conditions of the MIT License (see below), or 3) under the terms of dual Public Domain/MIT License conditions described here, as they choose.

The MIT License is about as close to Public Domain as a license can get, and is described in clear, concise terms at:

http://en.wikipedia.org/wiki/MIT_License

The full text of the MIT License follows:

```
=====
Copyright (c) 2007-2010 Baptiste Lepilleur
```

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN

CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

(END LICENSE TEXT)

The MIT license is compatible with both the GPL and commercial software, affording one all of the rights of Public Domain with the minor nuisance of being required to keep the above copyright notice and license text in the source code. Note also that by accepting the Public Domain "license" you can re-license your copy using whatever license you like.

libuuid Copyright and License (3-clause BSD)

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

1. Redistributions of source code must retain the above copyright notice, and the entire permission notice in its entirety, including the disclaimer of warranties.
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.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ALL OF WHICH ARE HEREBY 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 NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Lua License

Lua is licensed under the terms of the MIT license reproduced below. This means that Lua is free software and can be used for both academic and commercial purposes at absolutely no cost.

For details and rationale, see <http://www.lua.org/license.html> .

=====

Copyright (C) 1994-2006 Lua.org, PUC-Rio.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

OpenSSL License

```

/* =====
* Copyright (c) 1998-2008 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the 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.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.
*
* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written
* permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
* EXPRESSED 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 OpenSSL PROJECT OR
* ITS 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.

* =====

*

* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).

*

*/

Original SSLeay License

```

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
 * All rights reserved.
 *
 * This package is an SSL implementation written
 * by Eric Young (eay@cryptsoft.com).
 * The implementation was written so as to conform with Netscapes SSL.
 *
 * This library is free for commercial and non-commercial use as long as
 * the following conditions are aheared to.  The following conditions
 * apply to all code found in this distribution, be it the RC4, RSA,
 * lhash, DES, etc., code; not just the SSL code.  The SSL documentation
 * included with this distribution is covered by the same copyright terms
 * except that the holder is Tim Hudson (tjh@cryptsoft.com).
 *
 * Copyright remains Eric Young's, and as such any Copyright notices in
 * the code are not to be removed.
 * If this package is used in a product, Eric Young should be given attribution
 * as the author of the parts of the library used.
 * This can be in the form of a textual message at program startup or
 * in documentation (online or textual) provided with the package.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 * 1. Redistributions of source code must retain the 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.
 * 3. All advertising materials mentioning features or use of this software
 *    must display the following acknowledgement:
 *    "This product includes cryptographic software written by
 *     Eric Young (eay@cryptsoft.com)"
 *    The word 'cryptographic' can be left out if the rouines from the library
 *    being used are not cryptographic related :-).
 * 4. If you include any Windows specific code (or a derivative thereof) from
 *    the apps directory (application code) you must include an acknowledgement:
 *    "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
 *
 * THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``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 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.
*
* The licence and distribution terms for any publically available version or
* derivative of this code cannot be changed. i.e. this code cannot simply be
* copied and put under another distribution licence
* [including the GNU Public Licence.]
*/

```

tree.h Copyright and License

```

* Copyright 2002 Niels Provos
* All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the 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.
*/

```

zlib License

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly
Mark Adler

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?lnodocid=MIGR-5083174>.

IBM part number 81Y1040