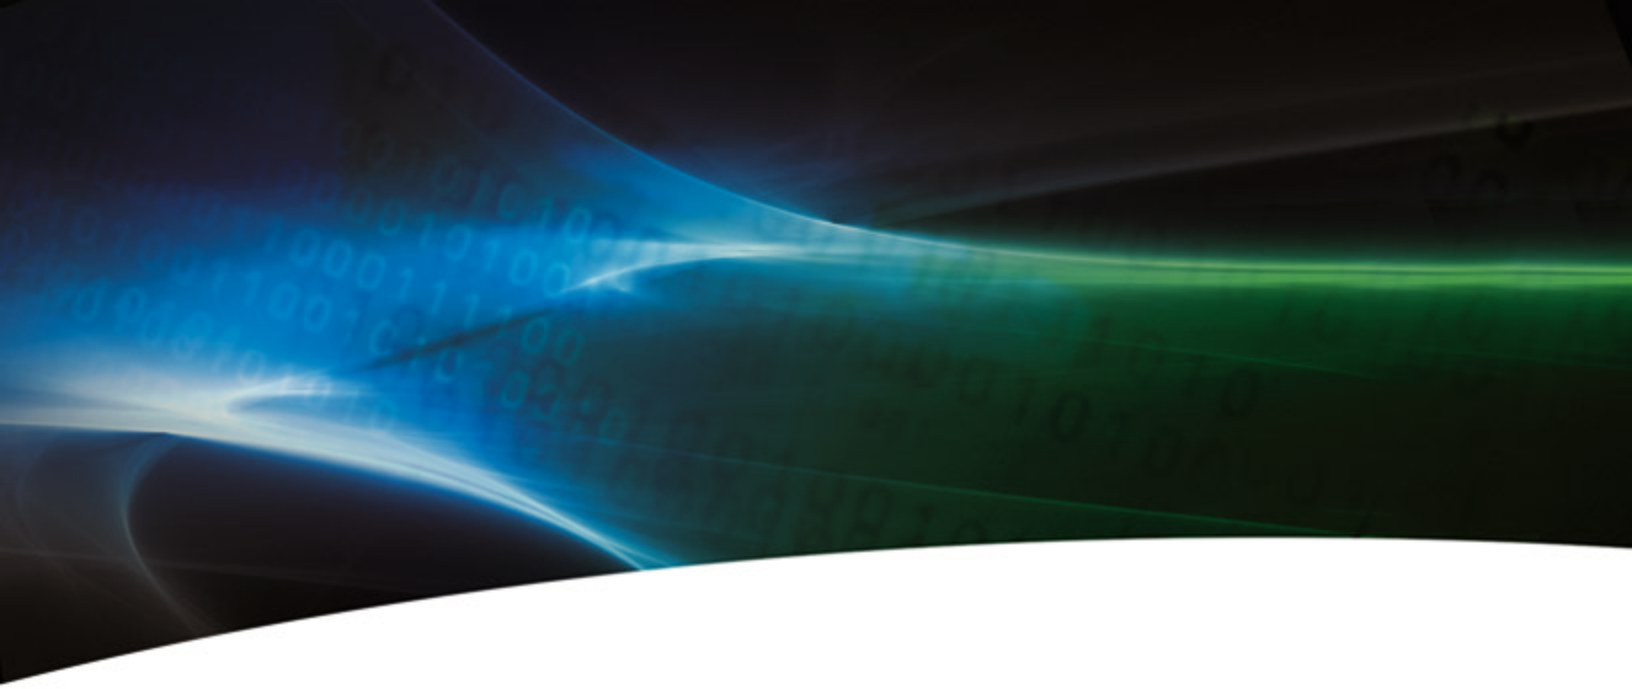

IBM ioMemory VSL 2.3.10

Release Notes

Copyright Notice

Part Number: D0001594-003_1
Published: August 20, 2012

Introduction

This document describes details about the 2.3.10 ioMemory® Virtual Storage Layer™ (VSL) release:

- System requirements, including supported operating systems and hardware requirements.
- Supported IBM High IOPS Adapters.
- Upgrade Notes, including the firmware version required for this release.
- Changes since the last generally available release.
- Issues (Errata) that may arise using this release.

System Requirements

Hardware Requirements

NOTE For complete hardware requirements and installation instructions, please refer to the *IBM High IOPS Adapter Hardware Installation Guide*.

Sufficient System Memory (RAM)

The amount of RAM the ioMemory VSL requires varies according to the average block size written to the device. Using the average block size table below, you can estimate the amount of system memory needed.

NOTE **4K Sector Sizes**: Depending on your operating system, you can reduce worst-case memory use by formatting your IBM High IOPS Adapter with a 4KB sector size and thereby force the average written block size to be 4KB or greater. However, some operating systems do not allow 4K sector sizes.

Even if you cannot force 4K sector sizes, the average write I/O size for most workloads is 4KB or larger. For this reason, 4KB average write size is typically the most accurate representation of memory utilization.

At various block sizes, the following table shows the upper limit of RAM that may be required of your system for every GB of IBM High IOPS Adapter storage space used.

Average Written Block Size (bytes)	RAM Usage (MB) per GB of storage space
8192	1.59
4096 (Most common)	2.85
2048	5.38
1024	11
512	22

For example, if your system is equipped with a device that has a total capacity of 1200GB (this example device has two IBM High IOPS Adapters) **formatted to use 4,096 byte sectors**, your system may require as much as: (1200GB capacity) x (2.85MB of RAM) = **3,420MB (or around 3.4GB) of system RAM** may be used by the ioMemory VSL in a worst-case scenario.

Attention The amount of RAM used by the VSL will depend on your use case; the table entries above are worst-case numbers. Actual RAM usage will likely be less than the amount listed.

You may run `fio-status -a` in the command line to see how much RAM the ioMemory VSL is

currently using per IBM High IOPS Adapter.

Supported Operating Systems

All operating systems must be 64-bit x86 architecture to support IBM High IOPS Adapters. Running the latest service pack / update of a release is strongly recommended.

Supported Linux Distributions

Attention The following distributions are supported. Some distribution versions may have binary packages available for download. If your version does not have a binary package available, you can build the installation package from the available source package. Check the download folders for available packages.

- Red Hat Enterprise Linux (RHEL) 5 (up to 5.8), 6 (up to 6.2)
- SUSE Linux Enterprise Server (SLES) 10 SP4, 11, 11 SP1, 11 SP2

Supported Microsoft Windows Operating Systems

- Microsoft Windows 2003 Server 64-Bit (with SP2 or higher)
- Microsoft Windows 2008 Server 64-Bit, R1 (with SP2 or higher)
- Microsoft Windows 2008 Server 64-Bit, R2

NOTE IBM High IOPS Adapters cannot be used as hibernation devices.

Hyper-V support

Hyper-V, as a Type 2 hypervisor on top of Windows 2008 R2, is supported.

Attention With Hyper-V, only a 512B sector size is supported on IBM High IOPS Adapters. For more information on sector sizes in Windows, see the following Microsoft Knowledge Base article: http://support.microsoft.com/kb/2510009.

Supported VMware Operating Systems

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

IBM High IOPS Adapters are only compatible with operating systems that are 64-bit x86 architecture. This means the following scenarios are supported:

1. Using the IBM High IOPS Adapter as VMFS datastore within the hypervisor, and then sharing that storage with guest operating systems. Guest operating systems can be 32-bit or 64-bit because they are not directly using the IBM High IOPS Adapter.

2. Using VMDirectPathIO, allow a virtual machine to directly use the IBM High IOPS Adapter. In this case, only supported operating systems can use the device.
   - VMDirectPathIO is currently supported on Windows and Linux operating systems that are supported by IBM. See either the *IBM ioMemory VSL User Guide for Linux* or the *IBM ioMemory VSL User Guide for Windows* for installation instructions.

     Attention When using VMDirectPathIO, if you upgrade the firmware on an IBM High IOPS Adapter, you must cycle the power to have the change take place. Just restarting the virtual machine won't apply the change.

Supported Devices

This section lists the legacy IBM adapters that are supported with the 2.x versions of the ioMemory VSL software. To identify the model cards installed in a server, run the management application or use the `fio-status` command line tool.

Attention Second generation adapters are not compatible with this version of the ioMemory VSL. If you have a newer IBM High IOPS Adapter that is not on this list, you may need to upgrade your system and all of your devices to ioMemory VSL 3.x software and firmware versions.

First Generation IBM High IOPS Mono Adapters (Low Profile cards)

- IBM 160GB High IOPS SS Class SSD PCIe Adapter
- IBM 320GB High IOPS MS Class SSD PCIe Adapter
- 320GB High IOPS SLC Adapter for IBM System x
- 640GB High IOPS MLC Adapter for IBM System x

First Generation IBM High IOPS Duo Adapters (Full Height cards)

- IBM 320GB High IOPS SD Class SSD PCIe Adapter
- 640GB High IOPS MLC Duo Adapter for IBM System x
- 640GB High IOPS SLC Duo Adapter for IBM System x
- 1.28TB High IOPS MLC Duo Adapter for IBM System x

Upgrade Notes

This version of the ioMemory VSL supports new features. These features require the latest version of the IBM firmware. Every IBM High IOPS Duo Adapter in a system should be upgraded to the latest version of the firmware.

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.

Firmware Version

The firmware archive that is released with this version of the ioMemory VSL contains the compatible firmware for this release.

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** -> **2.3.10**

Upgrade (both ioMemory VSL **and compatible firmware**) to the next highest version number on the path, and then continue on the path.

NOTE For example, upgrade a device configured for ioMemory VSL 1.2.8 to VSL 2.2.3 (and associated firmware) and then continue on the path.

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

Attention Note that when running multiple firmware upgrades in sequence (such as going from 1.2.8 to 2.2.3 to 2.3.10), after each subsequent firmware upgrade it is critical to reboot the system, load the ioMemory VSL (if it doesn't automatically load with your OS), and attach each device

Staged Upgrade Example

For more specific details on upgrading from one version to the next, see the *IBM ioMemory VSL Release Notes* (available at http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5083174) for each incremental version you will upgrade the device to. Then follow the upgrade instructions in that version's user guide for your operating system (including the firmware update instructions).

However, these upgrade procedures will follow this basic outline:

1. Unload the ioMemory VSL driver.
2. Uninstall the ioMemory VSL software.
3. Install the next version of the VSL software in the upgrade path.
4. Load the ioMemory VSL driver.
5. Update the firmware on the device(s) to the firmware that came with the ioMemory VSL.
6. Reboot
7. Ensure that the newly installed ioMemory VSL loads correctly and that all IBM High IOPS Adapters attach properly.
8. Repeat this procedure (if necessary) for all upgrades in the sequence.

Once you are ready to install this version of the ioMemory VSL (2.3.10), consult the User Guide for this version for further upgrade instructions.

Do Not Downgrade Device Firmware

Attention Do not downgrade the IBM High IOPS Adapter to an earlier version of the firmware. Earlier versions of the firmware may not be compatible with the device, and downgrading the firmware will result in data loss. If you need to downgrade your firmware, contact Customer Support at http://www.ibm.com/systems/support for compatibility information and to discuss your use case.

If you are installing new IBM High IOPS Adapters in a system using older devices and firmware, upgrade the older devices to the latest firmware and driver as a best practice.

Change Log

2.3.10 Change Log

The following are changes made to the ioMemory VSL from 2.2.3 to 2.3.10:

- Support for VMware operating systems.
- Expanded Linux Kernel support.
- RHEL 4 is no longer supported.
- Support for IO requests up to 1MB (up from 128KB).
- New utility, `fio-sure-erase`, see the user guide for Linux or Windows for more details.
- Support for PCIe gen3 with the `fio-pci-check` utility.
- Fixed various bugs, including:
  - When `fio-snmp-agentx` was run, the `blkid` could use up to 100% of the CPU.
  - TRIM (discard) was not working on RHEL 6.
  - `BLKGETSIZE` ioctal calls failed when called from 32-bit userspace.
  - A NUMA resource issue with Windows operating systems.

Errata

Overview

This section describes issues you may encounter when using this ioMemory VSL release.

General

Incorrect reserves indicated in fio-status

When an IBM High IOPS Adapter is in write-reduced mode (due to diminished reserves), the `fio-status` utility may erroneously report 100% health reserves.

Firmware update may fail with a TDO mismatch error

An IBM High IOPS Duo Adapter firmware upgrade may fail with the following error:

```
ERROR: TDO mismatch
```

This generally occurs when upgrading multiple devices at once.

To resolve this issue, update the devices again, but perform the update one device at a time. This includes IBM High IOPS Duo Adapters, upgrade each IBM High IOPS Adapter (on the duo product) individually.

Attention Do not reboot the system until the devices have been successfully updated.

Periodic latency bump

A latency bump every 30 seconds may be observed under some workloads, or a higher than expected average latency may be observed. This is due to normal maintenance.

Under Linux, this bump can be safely changed to happen every minute instead of every 30 seconds. To set a one minute interval, issue the following command every time after loading the driver:

```
$ echo 60000 > /proc/fusion/fio/fiox/data/groomer/groom-oldest-lebinterval
```

(where x is the letter of the device you are modifying). This is resolved in the ioMemory VSL 3.x releases.

Do not run fio-status during driver load

Run `fio-status` after the ioMemory VSL driver has loaded and not during driver load. Running fio-status while the driver is loading may yield the message:

```
Missing MIDS.  Coming up in minimal mode.
```

If this message is received when running `fio-status` while driver is loading, unload and then reload the driver and run `fio-status` after the driver has loaded.

fio-status may not display failed devices

On rare occasions, when an IBM High IOPS Adapter fails, the device may no longer appear in `fio-status`. If your device has failed, contact Customer Support.

Windows-Specific

Trim service fails to start

If the Fusion-io Trim service does not start, the following errors will appear in the Windows system log:

```
A timeout was reached (30000 milliseconds) while waiting for the Fusion-io
Trim Service service to connect.
```

```
The Fusion-io Trim Service service failed to start due to the following
error: The service did not respond to the start or control request in a
timely fashion.
```

To solve this issue, install the Microsoft Visual C++ 2010 SP1 Redistributable Package (x64). It can be found at: http://www.microsoft.com/download/en/details.aspx?id=13523.

Conversion to GPT or Dynamic disk terminates Logical Disk Manager Admin Service

This issue appears in any of these cases:

- Converting a Basic partition to GPT
- Converting a Basic partition to Dynamic Volume
- Switching between GPT and Dynamic Volume, either way

The message is: "The Logical Disk Manager Administrative Service terminated unexpectedly. Restart the 'Virtual Disk'

service."
This problem occurs only with Windows Server 2003; it is not a problem on Windows Server 2008 or Server 2008 R2.

If the IBM High IOPS Adapters are to be used in GPT or Dynamic mode, the following process should be done during the initial setup. This process will also recover devices that have had a failed conversion attempt.

Attention This will destroy any existing data on your devices. If you already have data on a device, be sure to back it up before proceeding.

1. For each IBM High IOPS Adapter in the system that is to be converted, use IBM High IOPS Management Application to a) detach the drive and b) do a low-level format (See the IBM High IOPS Management Application User Guide for details).
2. Restart the computer.
3. Go to Disk Management and select Initialize Disk.
4. Right-click and select Convert to GPT (or Convert to Dynamic Disk).

Windows Installer (DLL) Errors

If, during the installation of the Windows driver, you receive an error related to a dll, follow these steps:

1. Abort the installation process.
2. Run the uninstaller for any previous installations.
3. Stop any running Windows services related to ioMemory VSL products (`fio-agent`, for example) by opening Device Manager, navigating to the Services menu and choosing 'Stop' on these running services.
4. Reboot the machine.
5. Once the machine has rebooted, navigate to the installation directory of the driver and remove any files that were left over from the previous install (at a minimum, any leftover dll's). If you receive an 'access denied' error when attempting to remove these files, make sure you have Administrator privileges.
6. Remove fio.dll from the `Windows\System32` directory

   NOTE `fio.dll` may not exist in this directory, in which case you can continue.
7. You should now be able to run the new installer without error. If you still encounter errors, please contact Customer Support.

Linux-Specific

`kdump` crashkernel requires additional memory for the ioMemory VSL

With the ioMemory VSL installed, `kdump` may fail to generate a complete dump due to a lack of reserved memory. To ensure that there is enough memory, increase the memory allocated to the crashkernel.

The amount of memory required will very depending on your kernel version and the formatted sector sizes (more RAM is required for 512B sector sizes than for 4K sectors). The device sector size can be formatted using `fio-format`. A good starting point would be to increase the crashkernel reserve by approximately 50-250MB of RAM (depending on your configuration).

Attention Formatting the device will destroy the user data. Be sure to back up the data before running `fio-format`.

SMI-s provider registration issues when upgrading

When upgrading an existing installation of the SMI-s provider in Linux, the changes to the class schema may not be properly detected by the provider registration scripts when using the `sblim-cmpi` libraries. The provider will function normally and expose the correct data, but the schema not updating correctly may make it appear that the data is incorrect.

Software not installed post kernel update

When the ioMemory VSL is installed for a specific kernel version and the kernel version is changed for any reason, the VSL must be reinstalled to work with the new kernel version. RHEL5 has some processes that minimize the need to reinstall the software after kernel upgrade.

Rare error on driver unload using kernels older than 2.6.24

A bug in Linux kernels prior to 2.6.24 can cause a general protection fault or other kernel error when the driver is unloaded. This bug also affects non-IBM drivers. The bug has been resolved in newer kernels. More information is available here:

http://www.kernel.org/pub/linux/kernel/v2.6/ChangeLog-2.6.24

(search for "commit 5a622f2d0f86b316b07b55a4866ecb5518dd1cf7")

Because this is a bug in the Linux kernel, IBM cannot resolve this issue for older kernels.

ext4 in Kernel 2.6.33 or earlier may silently corrupt data when discard (TRIM) is enabled

The ext4 filesystem in Kernel.org kernel 2.6.33 and earlier has a bug where the data in a portion of a file may be improperly discarded (set to all 0x00) under some workloads. Use 2.6.34 or newer to avoid this bug. For more info see the patch 1 and bug report 2 below.

The fix is included in RHEL6 as of pre-release kernel `kernel-2.6.32-23.el6`. The production RHEL6 kernel is not affected by this issue.

Discard support was added to the kernel.org mainline ext4 in 2.6.28 and was enabled by default. For fear of damaging some devices, discard was set to default to disabled in v2.6.33-rc1 and was back ported to 2.6.31.8 and v2.6.32.1.

1. http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=commitdiff;h=b90f687018e6d6

2. https://bugzilla.kernel.org/show_bug.cgi?id=15579
3. http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=commitdiff;h=5328e635315734d

Kernels 2.6.34/35 don't handle switching interrupt types

Linux kernels around 2.6.34/35 may have problems processing interrupts if the driver is loaded using one interrupt type, unloaded, and then loaded again using a different interrupt type. The primary symptom is that the IBM High IOPS Adapter is unusable, and the kernel logs have errors with "doIRQ". For example, the following sequence on an affected system would likely result in errors.

1. load the driver with default of disable_msi=1 which selects APIC interrupts

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

2. load the driver, enabling MSI interrupts

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

To work around this issue, reboot if you see the error and always load with the same interrupt type selected. To change between interrupt types, reboot first.

RHEL6 udevd warning

When using an IBM High IOPS Adapter under RHEL6 (or any Linux distro with udev version 147 or greater), udevd may emit the following innocuous messages:

```
udevd[154]: worker [19174] unexpectedly returned with status 0x0100
udevd[154]: worker [19174] failed while handling
'/devices/virtual/block/fioa'
```

You can ignore this warning.

RHEL6 warn_slowpath during device attach

When attaching an IBM High IOPS Adapter under RHEL6, you may find log messages similar to the following:

```
kernel: ------------[ cut here ]------------
kernel: WARNING: at fs/fs-writeback.c:967 __mark_inode_dirty+0x108/0x160()
(Tainted: P           ---------------- )
.
.
.
[<ffffffff8106b857>] warn_slowpath_common+0x87/0xc0
[<ffffffff8106b8aa>] warn_slowpath_null+0x1a/0x20
.
.
.
```

This is due to a bug in the 2.6.32 kernel, and the warning can safely be ignored.

Switching interrupt types with newer kernels can cause errors

With newer Linux kernels, switching interrupt types after initial driver load can cause `doIRQ` errors to be reported by the kernel. As a work around, reboot your system before loading the driver with the new interrupt type specified.

Stop fio-snmp-agentx when the ioMemory VSL driver is unloaded

If running `snmpd` and `fio-snmp-agentx` with the `fio-snmp-agentx` configured to write to a log file the `fio-snmp-agentx` should be stopped if the ioMemory VSL driver is unloaded. Otherwise the `fio-snmp-agentx` will continually log error messages to the log file stating that it is unable to access IBM High IOPS Adapter information. It does not create any issues with the system other than unnecessary file space consumption.

Compiler Cache (ccache) causes ioMemory VSL src.rpm rebuild failures on some distributions

If the ccache package is installed, rebuilding the ioMemory VSL src.rpm may fail with an error similar to the following:

```
CC [M]
/root/fio/iomemory-vls-<version>/root/usr/src/iomemory-vsl/driver_init.o
/root/fio/iomemory-vls-<version>/root/usr/src/iomemory-vsl/driver_init.c:116:
error: initializer element is not constant
[...]
```

To allow the VSL to rebuild, remove the ccache package or disable ccache.

VMware-Specific

Using VMDirectPathIO with multiple-device products

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

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

ESXi 5.0 injected installer allows installation on an IBM High IOPS Adapter

IBM High IOPS Adapters are not designed to be bootable, therefore you should not install the host OS on an IBM High IOPS Adapter. The ESXi injected installer will permit you to install the OS on an IBM High IOPS Adapter, but the installation will fail on reboot.

IBM High IOPS Adapters are only labeled as a "VMware Block Device" in the installer screen. Make sure to select a bootable device during the installation.

The ioMemory VSL cannot be installed as part of an ESXi 4.1 installation

The ioMemory VSL must be installed on an existing ESXi host system.

Download Location

Software, utilities, and related documentation for this version can be found at:

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

IBM part number 81Y1045