

Release Notes

Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.0.2

Legal and notice information

© Copyright 2008 Hewlett-Packard Development Company, L.P.

Overview

This release notes discusses the recent product information about HP Device Mapper Multipath (HPDM Multipath) Enablement Kit for HP StorageWorks Disk Arrays v4.0.2. This incremental release is to provide enablement for HP StorageWorks MSA 2000 product family with Serial Attached SCSI (SAS) connectivity (MSA2012sa) and to support RedHat Enterprise Linux (RHEL) 4 Update 7 operating system.

For RedHat Enterprise Linux 4 Update 7, RedHat Enterprise Linux 5 Update 2, and SuSE Linux Enterprise Server (SLES) 10 Service Pack (SP) 2 operating systems, use Device Mapper Multipath provided with the distributions. For recommended configuration settings, see [Installing Device Mapper Multipath tools](#).

HPDM Multipath offers the following features:

- **I/O failover and failback:** Provides transparent failover and failback of I/Os by rerouting I/Os automatically to an alternative path when a path failure is sensed and routing them back when the path is restored.
- **Path grouping policies:** Paths are coalesced based on the following path-grouping policies:
 1. *Priority based path-grouping*
 - Provides priority to group paths based on Asymmetric Logical Unit Access (ALUA) state
 - Provides static load balancing policy by assigning higher priority to the preferred path
 2. *Multibus*
 - All paths are grouped under a single path group
 3. *Group by serial*
 - Paths are grouped together based on controller serial number
 4. *Failover only*
 - Provides failover without load balancing by grouping the paths into individual path groups
- **I/O load balancing policies:** Provides weighted Round Robin load balancing policy within a path group.
- **Device name persistence:** Device names are persistent across reboots and Storage Area Network (SAN) reconfigurations. Device Mapper also provides configurable device name aliasing feature for easier management.
- **Persistent device settings:** All the device settings such as load balancing policies, path grouping policies are persistent across reboots and SAN reconfigurations.
- **Device exclusion:** Provides device exclusion feature through blacklisting of devices.
- **Path monitoring:** Periodically monitors each path for status and enables faster failover and failback.
- **Online device addition and deletion:** Devices can be added or deleted to or from Device Mapper (DM) multipath without rebooting the server or disrupting other devices or applications.
- **Management Utility:** Provides Command Line Interface (CLI) to manage multipath devices.
- **Boot from SAN:** Provides multipathing for operating system installation partitions on SAN device.



NOTE:

Multipathing for SAN Boot environment is supported on the following operating systems: RHEL 5 U1/RHEL 5 U2/SLES 10 SP1/SLES 10 SP2. For more details on Multipathing support for SAN Boot environment, see the *Booting Linux x86 and x86_64 systems from a Storage Area Network with Device Mapper Multipath* document available at <http://www.hp.com>.

What's New

HPDM Multipath provides the following additional features:

- Provides support for RHEL 4 Update 7 operating system distribution

- Provides device parameter values to enable multipath support for HP StorageWorks MSA2000 product family with SAS connectivity (MSA2012sa)



NOTE:

For operating systems supported with HP StorageWorks MSA2000 product family with SAS connectivity (MSA2012sa), refer to the Support Matrix of MSA2012sa.

HPDM Multipath support matrix

Table 1 lists the hardware and software prerequisites for installing HPDM Multipath.

Table 1 Hardware and software prerequisites

System feature	Supported hardware and software
Operating system versions	RHEL 4 Update 5 RHEL 4 Update 6 RHEL 4 Update 7 RHEL 5 errata 2.6.18–8.1.6.el5 or later RHEL 5 Update 1 RHEL 5 Update 2 SLES 9 - SP 3 errata 2.6.5–7.283 or later SLES 9 - SP 4 SLES 10 SLES 10 - SP 1 SLES 10 - SP 2
Host Bus Adapters (HBA) SAN Switches	See http://h18006.www1.hp.com/storage/networking/index.html and http://h18004.www1.hp.com/products/servers/proliantstorage/adapters/index.html
Servers	HP BladeSystem c-Class Server Blades, ProLiant x86, ProLiant AMD64, ProLiant EM64T Servers, Integrity Servers
Supported arrays	EVA 3000 (HSV101) VCS 4.007 or later EVA 5000 (HSV111) VCS 4.007 or later EVA 4000 (HSV200) XCS 5.110/6.100 or later EVA 6000 (HSV200) XCS 5.110/6.100 or later EVA 8000 (HSV210) XCS 5.110/6.100 or later EVA 4100 (HSV200) XCS 6.110 or later EVA 6100 (HSV200) XCS 6.110 or later EVA 8100 (HSV210) XCS 6.110 or later EVA 4400 (HSV300) XCS 0900 or later EVA iSCSI Connectivity Option XP10000 fw rev 50-07-30-00/00 or later XP12000 fw rev 50-09-34-00/00 or later XP 20000 fw rev 60-02-04-00/00 or later XP 24000 fw rev 60-02-04-00/00 or later MSA 1000/MSA 1500 fw rev 7.0.0 or later MSA 2000 Storage product family (MSA2012fc/MSA2212fc) fw rev J200P19 or later MSA2012i fw rev J210R10 or later MSA2012sa fw rev J300R24–01 or later

HBA drivers	<p>HP SC08Ge Host Bus Adapter: 4.00.13.04-2 or later (for RHEL 5 U1/SLES 10 SP1), 3.12.14.00-2 or later (for RHEL 4 U6/ SLES 9 SP4) available at: http://h20000.www2.hp.com/bizsupport/TechSupport/DriverDownload.jsp?lang=en&cc=us&prodNameId=3759720&taskId=135&prodTypeId=329290&prodSeriesId=3759718&lang=en&cc=us</p> <p>Emulex: 8.0.16.40 or later (for RHEL 4 U7), 8.2.0.22 or later (for SLES 10 SP2/RHEL 5 U2), 8.0.16.32 or later (for SLES 9/RHEL 4 U6), 8.1.10.11 or later (for RHEL 5 U1/SLES 10 SP1) available at: http://h18006.www1.hp.com/products/storageworks/4gbpciehba/index.html</p> <p>Qlogic: 8.02.11 or later (for RHEL 4 U7), 8.02.11 or later (for SLES 10 SP2/RHEL 5 U2), 8.01.07.25 or later (for SLES 9/ SLES 10 SP1/ RHEL 4 U6/ RHEL 5 U1) available at: http://h18006.www1.hp.com/products/storageworks/fca2214/index.html</p>
-------------	--



NOTE:

HPDM Multipath provides support for iSCSI devices on the following operating systems: RHEL 5/RHEL 5 U1/RHEL 5 U2/SLES 9 SP3/SLES 9 SP4/SLES 10/SLES 10 SP1/SLES 10 SP2.

Installing Device Mapper Multipath tools

Ensure the following RPMs bundled with the operating system distributions are installed on the system:

- For RHEL 4 Update 7:
device-mapper-1.02.25-2.el4 or later, device-mapper-multipath-0.4.5-31.el4 or later
- For RHEL 5 Update 2:
device-mapper-1.02.24-1.el5 or later, device-mapper-multipath-0.4.7-17.el5 or later
- For SLES 10 SP2:
device-mapper-devel-1.02.13-6.14 or later, multipath-tools-0.4.7-34.38 or later, device-mapper-1.02.13-6.14 or later



NOTE:

For RHEL 4 U5/RHEL 4 U6/RHEL 5/RHEL 5 U1/SLES 9 SP3/SLES 9 SP4/SLES 10/SLES 10 SP1 operating systems, install HP Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.0.0.

Configuring Device Mapper Multipath to enable HP arrays

The HP Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays 4.0.2 package provides a template file with the recommended configuration for HP supported arrays. Complete the following steps to enable the HP StorageWorks Disk Arrays with Device Mapper:

1. Download the HP Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.0.2 available at <http://www.hp.com/go/devicemapper>.
2. Login as root to the host system.
3. Copy the tar package to a temporary directory (for instance, /tmp/HPDMmultipath).
4. Unbundle the package using the following commands:
#cd /tmp/HPDMmultipath
#tar -xvzf HPDMmultipath-4.0.2.tar.gz
#cd HPDMmultipath-4.0.2
5. Verify that the directory contains the README.txt, COPYING, conf and docs directories.

6. If you have `/etc/multipath.conf` file on your system, edit the file with recommended device parameter values, see [Recommended device parameter values](#), else select the correct configuration template based on your operating system from the `conf` and copy it to `/etc/multipath.conf`:

```
#cd conf
#cp multipath.conf.HPTemplate.<OS version> /etc/multipath.conf
```
7. Restart multipath services

```
#/etc/init.d/multipathd restart
```
8. Create multipath devices by executing the following command:

```
#multipath -v0
```

Recommended device parameter values

To enable the HP arrays, edit `/etc/multipath.conf` file by adding the following under devices section:

For EVA3000 / EVA5000 / EVA4x00/ EVA6x00 / EVA8x00

```
device
{
    vendor                "(COMPAQ|HP) "
    product                "HSV1[01]1|HSV2[01]0|HSV300"
    getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout           "/sbin/mpath_prio_alua /dev/%n"
    features               "0"
    hardware_handler       "0"
    path_selector           "round-robin 0"
    path_grouping_policy   group_by_prio
    failback               immediate
    rr_weight              uniform
    no_path_retry          12
    rr_min_io              100
    path_checker            tur
}
```

For MSA2012fc/MSA2212fc/MSA2012i

```
device
{
    vendor                "HP"
    product                "MSA2[02]12fc|MSA2012i"
    getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
    features               "0"
    hardware_handler       "0"
    path_selector           "round-robin 0"
    path_grouping_policy   multibus
    failback               immediate
    rr_weight              uniform
    no_path_retry          18
    rr_min_io              100
    path_checker            tur
}
```

For MSA1000/MSA1500

```
device
{
    vendor                "HP"
    product               "MSA VOLUME*"
    getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout           "/sbin/mpath_prio_alua /dev/%n"
    features               "0"
    hardware_handler       "0"
    path_selector          "round-robin 0"
    path_grouping_policy   group_by_prio
    failback               immediate
    rr_weight              uniform
    no_path_retry          12
    rr_min_io              100
    path_checker           tur
}
```

For MSA2012sa

```
device
{
    vendor                "HP"
    product               "MSA2012sa"
    getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout           "/sbin/mpath_prio_alua /dev/%n"
    features               "0"
    hardware_handler       "0"
    path_selector          "round-robin 0"
    path_grouping_policy   group_by_prio
    failback               immediate
    rr_weight              uniform
    no_path_retry          18
    rr_min_io              100
    path_checker           tur
}
```

For XP

```
device
{
    vendor                "HP"
    product               "OPEN-.*"
    getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
    features               "0"
    hardware_handler       "0"
    path_selector          "round-robin 0"
    path_grouping_policy   multibus
    failback               immediate
    rr_weight              uniform
    no_path_retry          12
    rr_min_io              1000
    path_checker           tur
}
```

**NOTE:**

- For SLES 10 SP2, in the device section, replace

```
prio_callout      "/sbin/mpath_prio_alua /dev/%n"
with
prio              alua
```

- For more information on editing `/etc/multipath.conf`, see the *Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays 4.0.0 Installation and Reference Guide*. You can find this document on the Manuals page of **Multi-path Device Mapper for Linux Software**, which is accessible at <http://www.hp.com/go/devicemapper>.

Setting up HPDM Multipath

Setting up HPDM Multipath involves the following:

- Configuring QLogic HBA parameters
- Configuring Emulex HBA parameters

Configuring QLogic HBA parameters

To configure the QLogic HBA parameters for QLogic 2xxx family of HBAs, complete the following steps:

1. Edit the `/etc/modprobe.conf` file in RHEL hosts and `/etc/modprobe.conf.local` file in SLES hosts with the following values:

```
options qla2xxx ql2xmaxqdepth=16 qlport_down_retry=10 ql2xloginretrycount=30
ql2xfailover=0 ql2xlbType=1 ql2xautorestore=0xa0 ConfigRequired=0
```

2. Rebuild the `initrd` by executing the following script:

```
/opt/hp/src/hp_qla2x00src/make_initrd
```

3. Reboot the host.

Configuring Emulex HBA parameters

To configure the Emulex HBA parameters, complete the following steps:

1. For Emulex `lpfc` family of HBAs:

- In RHEL 4 hosts, edit the `/etc/modprobe.conf` file with the following values:

```
options lpfc
        lpfc_nodev_tmo=14
        lpfc_lun_queue_depth=16
        lpfc_discovery_threads=32
```

- In SLES 9 hosts, edit the `/etc/modprobe.conf.local` file with the following values:

```
options lpfc
        lpfc_nodev_tmo=14
        lpfc_lun_queue_depth=16
        lpfc_discovery_threads=32
```

- In SLES 10 hosts, edit the `/etc/modprobe.conf.local` file with the following values:

```
options lpfc
        lpfc_nodev_tmo=28
        lpfc_lun_queue_depth=16
        lpfc_discovery_threads=32
```


- In RHEL 5 hosts, edit the `/etc/modprobe.conf` file with the following values:

```
options lpfc
    lpfc_nodev_tmo=28
    lpfc_lun_queue_depth=16
    lpfc_discovery_threads=32
```

2. Rebuild the `initrd` by executing the following script:

```
/opt/hp/hp-lpfc/make_initrd
```

3. Reboot the host.



NOTE:

Ensure you have a HBA driver installed in single path mode. For more information, see the installation and reference guides for the respective HBA drivers.

Setting up Device Mapper Multipath daemon

You must set the Device Mapper Multipath daemon to start at boot time.

For RHEL hosts, complete the following steps to start the `multipathd` daemon at boot time:

1. Run the following command to check if the daemon is configured to start at boot time:
`chkconfig --list multipathd`
2. Run the following commands to start the Device Mapper Multipath daemon:
`chkconfig [--level levels] multipathd on`
`chkconfig multipathd`

For SLES hosts, complete the following steps to start the `multipathd` daemon at boot time:

1. Run the following commands to check if the daemon is configured to start at boot time:
`chkconfig --list boot.device-mapper`
`chkconfig --list boot.multipath`
`chkconfig --list multipathd`
2. Run the following commands to start the Device Mapper Multipath daemons:
`chkconfig boot.device-mapper [levels]`
`chkconfig boot.multipath [levels]`
`chkconfig multipathd [levels]`

Known Issues

Following are the known issues in the HPDM Multipath 4.0.2 release:

- Path failure messages are seen in the log file when encountered with a Unit Attention condition. This behavior could be observed during online LUN addition, deletion, LUN transition across controller. These paths would be recovered after the polling interval set.
- `Multipath` commands may take longer time to execute on heavily loaded servers or under path failure conditions.
- User friendly multipath device names may change on reboot if `/var` is mounted on a separate partition other than the root file system. It is recommended to have `/(root)` and `/var` on the same partition.
- Blacklisting the multipath device in the file `/etc/multipath.conf` and restarting the multipath service may not remove the device on RHEL 4 distributions. Execute the following command to remove the blacklisted device:
`multipath -f <device>`
- On RHEL 4 Update 7 hosts, unrepresenting a LUN and rebooting the host may result in system hangs. It is recommended to execute the following command before rebooting the host:

```
# multipath -f <device>
```

- Using `fdisk` command to create partitions may fail to create Multipath device for the partition device. It is recommended to use `parted` command to create partitions for the device.
- `Multipathd` daemon crashes on systems configured with device paths more than the system open file limits (default system open file limit=1024). It is recommended to change the system open file limits using the '`max_fds`' parameter in `/etc/multipath.conf` file or by using the `ulimit -n` command.
- `Multipath -l` command may not reflect the correct path status for Logical Units presented from MSA2012sa array when paths fail or are restored under heavy load conditions. To refresh the path status, execute the `# multipath -v0` command.

Support

Telephone numbers for worldwide technical support are listed on the HP support website:

<http://www.hp.com/support/>

Collect the following information before calling the worldwide technical support:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP recommends that customers sign up online using the Subscriber's choice website:

<http://www.hp.com/go/e-updates>

- Subscribing to this service provides you with e-mail updates on the latest product enhancements, newer versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can locate your products by selecting **Business support > Storage** under Product Category.