

Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.2.1 release notes

Part number: AA-RWF9H-TE
First edition: March 2009



Legal and notice information

© Copyright 2009 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 v4.2.1. This incremental release is to support SUSE Linux Enterprise Server 11 operating system.

Device Mapper 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 user defined priority to the 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 the following load balancing policies within a path group:
 1. *Weighted Round Robin*

This round-robin algorithm routes `rr_min_io` number of I/Os on a selected path before switching to the next path.
 2. *Least pending I/O path*

This determines the number of non-serviced requests pending on a path and selects the path which has the least number of pending requests for service.
 3. *DM service time*

This is a service time oriented dynamic load balancer, which selects a path to complete the incoming I/O with the shortest time.
- **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 to or deleted 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:**

For 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://h18006.www1.hp.com/storage/networking/bootsan.html>

What's new

HPDM Multipath 4.2.1 provides the following additional feature:

- Provides support for SUSE Linux Enterprise Server 11

 **NOTE:**

For more information on operating systems supported with HP StorageWorks Disk Arrays, see the SPOCK website:

www.hp.com/storage/spock

Device Mapper 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	SLES 11
Host Bus Adapters (HBA) SAN Switches	See http://h18006.www1.hp.com/storage/networking/index.html , http://h18004.www1.hp.com/products/servers/proliantstorage/adapters/index.html , and http://h20000.www2.hp.com/bizsupport/TechSupport/DriverDownload.jsp?lang=en&cc=us&prodNameId=3628653&taskId=135&prodTypeId=332283&prodSeriesId=3628652&submit.y=2&submit.x=5&lang=en&cc=us
Servers	HP BladeSystem c-Class Server Blades, ProLiant x86, ProLiant AMD64, ProLiant EM64T Servers, Integrity Servers

Supported arrays	<p>EVA4000 (HSV200) XCS 5.110/6.200 or later EVA6000 (HSV200) XCS 5.110/6.200 or later EVA8000 (HSV210) XCS 5.110/6.200 or later EVA4100 (HSV200) XCS 6.200 or later EVA6100 (HSV200) XCS 6.200 or later EVA8100 (HSV210) XCS 6.200 or later EVA4400 (HSV300) XCS 0900 or later EVA6400 (HSV400) XCS 0950 or later EVA8400 (HSV450) XCS 0950 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 XP20000 fw rev 60-02-04-00/00 or later XP24000 fw rev 60-02-04-00/00 or later MSA1000/MSA1500 fw rev 7.0.0 or later MSA2000 Storage product family (MSA2012fc/MSA2212fc) fw rev J200P24 or later MSA2012i fw rev J210R10 or later MSA2012sa fw rev J300P13 or later (MSA2312fc/MSA2324fc) fw rev M100R18 or later</p>
HBA drivers and Smart Array Controller drivers	Inbox drivers

Installing Device Mapper Multipath tools

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

- For SLES 11:
device-mapper-1.02.27-8.6 or later, multipath-tools-0.4.8-40.1 or later

Installing HPDM Multipath Enablement kit 4.2.1

To install HPDM Multipath 4.2.1, complete the following steps:

1. Download the HPDM Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.2.1 available at <http://www.hp.com/go/devicemapper>.
2. Log in as root to the host system.
3. Copy the installation tar package to a temporary directory (for instance, /tmp/HPDMmultipath).
4. Unbundle the package by executing the following commands:

```
#cd /tmp/HPDMmultipath
#tar -xvzf HPDMmultipath-4.2.1.tar.gz
#cd HPDMmultipath-4.2.1
```
5. Verify that the directory contains README.txt, COPYING, conf, and docs directories. After the verification perform one of the following:

- If you do not find the `/etc/multipath.conf` file, copy the `conf/multipath.conf.HPTemplate.SLES11` template to `/etc/multipath.conf` file by executing the following command:

```
#cp conf/multipath.conf.HPTemplate.SLES11 /etc/multipath.conf
```

- If the `/etc/multipath.conf` file is already present, edit the file with the recommended device parameter values. For more details on recommended device parameter values, see [Recommended device parameter values](#).

6. Restart multipath services by executing the following command:

```
# /etc/init.d/multipathd restart
```

Configuring Device Mapper Multipath to enable HP arrays

This section describes the following:

- [Recommended device parameter values](#)
- [Setting up HPDM Multipath](#)
- [Setting up Device Mapper Multipath daemon](#)

Recommended device parameter values

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

For EVA3000/EVA5000/EVA4x00/EVA6x00/EVA8x00

```
device
{
    vendor                "(HP|COMPAQ)"
    product               "HSV1[01]1\C\)COMPAQ|HSV[2][01]0|HSV300|HSV4[05]0"
    path_grouping_policy  group_by_prio
    getuid_callout        "/lib/udev/scsi_id -g -u /dev/%n"
    path_checker          tur
    path_selector         "round-robin 0"
    prio                  alua
    rr_weight             uniform
    failback              immediate
    hardware_handler      "0"
    no_path_retry         12
    rr_min_io             100
}
```

For XP

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

For MSA A/A arrays

```
device
{
    vendor                "HP"
    product                "MSA VOLUME*"
    path_grouping_policy  group_by_prio
    getuid_callout        "/lib/udev/scsi_id -g -u /dev/%n"
    path_checker          tur
    path_selector         "round-robin 0"
    prio                  alua
    rr_weight              uniform
    failback              immediate
    hardware_handler      "0"
    no_path_retry         12
    rr_min_io             100
}
```

For MSA2012sa

```
device
{
    vendor                "HP"
    product                "MSA2012sa"
    path_grouping_policy  group_by_prio
    getuid_callout        "/lib/udev/scsi_id -g -u /dev/%n"
    path_checker          tur
    path_selector         "round-robin 0"
    prio                  alua
    rr_weight              uniform
    failback              immediate
    hardware_handler      "0"
    no_path_retry         18
    rr_min_io             100
}
```

For MSA2012fc/MSA2212fc/MSA2012i

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

For MSA2312fc/MSA2324fc

```
device
{
    vendor                "HP"
    product               "MSA2312fc|MSA2324fc"
    getuid_callout        "/lib/udev/scsi_id -g -u /dev/%n"
    hardware_handler      "0"
    path_selector         "round-robin 0"
    prio                 alua
    path_grouping_policy  group_by_prio
    failback              immediate
    rr_weight             uniform
    no_path_retry         18
    rr_min_io             100
    path_checker          tur
}
```

 **NOTE:**

- For more information on editing `/etc/multipath.conf` file, 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 includes configuring HBA and iSCSI initiator parameters for multipathed environment. This involves the following:

- [Configuring QLogic HBA parameters](#)

- [Configuring Emulex HBA parameters](#)
- [Configuring iSCSI 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.local` file in SLES hosts with the following value:


```
options qla2xxx ql2xmaxqdepth=16 qlport_down_retry=10 ql2xloginretry-count=30
```
2. Rebuild the `initrd` by completing the following steps:
 - a. Backup the existing `initrd` image by executing the following command:


```
#mv /boot/initrd-<version no.>.img /boot/initrd-<version no.>.img.old
```
 - b. Make a new `initrd` image by executing the following command:


```
#mkinitrd
```
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, 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
```
2. Rebuild the `initrd` by completing the following steps:
 - a. Backup the existing `initrd` image by executing the following command:


```
#mv /boot/initrd-<version no.>.img /boot/initrd-<version no.>.img.old
```
 - b. Make a new `initrd` image by executing the following command:


```
#mkinitrd
```
3. Reboot the host.

Configuring iSCSI parameters

To configure the iSCSI parameters, complete the following steps:

1. Update the iSCSI configuration file
 - Edit the `/etc/iscsi/iscsid.conf` file with the following value:


```
node.session.timeo.replacement_timeout=10
```
2. Restart the iSCSI service using the following command:


```
#/etc/init.d/open-iscsi restart
```

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 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]
```

Known issues

Following are the known issues in the HPDM Multipath 4.2.1 release:

- Path failure messages are seen in the log file when encountered with a Unit Attention condition. This behavior can be observed during online LUN addition, deletion, LUN transition across the controller. These paths can 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 partition other than the root file system. It is recommended to have `/(root)` and `/var` on the same partition or change the multipath bindings file location using the 'bindings_file' parameter in the `/etc/multipath.conf` file.
- 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 by using either the 'max_fds' parameter in `/etc/multipath.conf` file or by using the `ulimit -n` command and restart the `multipathd` demon.
- `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.
- Multipath devices may not be created for Logical Units when the system disks or internal controllers are `cciss` devices. It is recommended to blacklist these devices in the `/etc/multipath.conf` file and restart the `multipathd` daemon.

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/subscribe-gate1/>

- 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.