

Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.4.1 release notes



Legal and notice information

© Copyright 2009-2010 Hewlett-Packard Development Company, L.P.

Overview

This is the final release of HP Device Mapper Multipath Enablement Kit. A new Native Linux Device Mapper Multipath reference guide has been put together. User can follow the new guide to setup linux Device Mapper Multipath for HP arrays. The guide will be updated as new arrays are released and with other relevant up-to-date information. The Native Linux Device Mapper Multipath reference guide can be found on SPOCK both internally and externally at <http://www.hp.com/storage/spock>. On the web site, look for “Application Notes” on left hand side and click on “Solutions: Linux”.

This release notes discusses the recent product information about HP Device Mapper Multipath (HPDM Multipath) Enablement Kit for HP StorageWorks Disk Arrays v4.4.1. This incremental release is to provide enablement for RedHat Enterprise Linux 5 Update 5. With majority of technical features well integrated and distributed by Linux distributions, as mentioned above, this “kit” will simply be replaced by a document which user can follow. The document will include HP array profiles, especially new array, and other up-to-date relevant information.

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:
 - *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
 - *Multibus* — All paths are grouped under a single path group
 - *Group by serial* — Paths are grouped together based on controller serial number
 - *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:
 - *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.
 - *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.
 - *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 devices.
- **Cluster support:** Provides multipathing in HP Serviceguard and SteelEye LifeKeeper clustering environment.
- **Volume Manager support:** Provides support for multipathing devices to be configured under Logical Volume Manager.

NOTE:

The following features are available only on SLES 11 operating system:

- Least pending I/O path
- DM service time

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 4.4.1 provides the following:

- Provides support for RHEL 5 Update 5
-

NOTE:

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

www.hp.com/storage/spock

For more information on the inbox HBA driver parameters, see [Setting up HPDM Multipath](#).

Device Mapper Multipath support matrix

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

System feature	Supported hardware and software
Operating system versions	RHEL 4 Update 6 RHEL 4 Update 7 RHEL 4 Update 8 RHEL 5 Update 2 RHEL 5 Update 3 RHEL 5 Update 4 RHEL 5 Update 5 SLES 10 SP2 SLES 10 SP3 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

System feature	Supported hardware and software
Server	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 MSA2000 Storage product family (MSA2012fc/MSA2212fc) fw rev J200P19 or later MSA2012i fw rev J210R10 or later MSA2012sa fw rev J300P13 or later (MSA2312fc/MSA2324fc) fw rev M100R18 or later MSA2312sa/MSA2324sa fw rev M110R20 or later MSA2312i/MSA2324i fw rev M110R20 or later P2000 fc fw rev TS100R023 or later P2000 fc/iSCSI fw TS100R023 or later</p>
HBA drivers and Smart Array Controller drivers	<p>HP SC08Ge Host Bus Adapter: 4.00.13.04-2 or later (for RHEL 5/SLES 10), 3.12.14.00-2 or later (for RHEL 4) available at: 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 Emulex: 8.0.16.40-11 or later (for RHEL 4 U8), 8.0.16.40 or later (for RHEL4 U7), 8.2.0.22 or later for (SLES 10 SP2/RHEL 5 U2), 8.0.16.32 or later (for RHEL 4 U6) available at: http://h18006.www1.hp.com/products/storageworks/4gbpciehba/index.html Qlogic: 8.02.23-1 or later (for RHEL 4 U8), 8.02.11 or later (for SLES 10 SP2/RHEL 5 U2/RHEL 4 U7), 8.01.07.25 or later (for RHEL 4 U6) available at: http://h18006.www1.hp.com/products/storageworks/fca2214/index.html Brocade: 1.1.0.10 available at: http://h20180.www2.hp.com/apps/lookup?hpagetype=s-001&hlang=en&hclient=s-s-r2515-1&hcc=us&hquery=HP+StorageWorks+PCle+4Gb+Host+Bus+Adapter HP Smart Array P700m Controller: 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 RHEL 5 U3/RHEL 5 U4/SLES 10 SP3/SLES 11 FC HBA drivers — Inbox drivers</p>

NOTE:

- For more information on configuring iSCSI parameters, see [Configuring iSCSI parameters](#).
 - Device Mapper Multipath does not support coexistence with other multipath products.
 - Device Mapper Multipath does not support Active-Passive Storage Arrays.
 - HPDM Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.4.1 does not support SAS and CCISS devices on SLES 10 SP3.
 - On SLES 11, Device Mapper Multipath for iSCSI devices is supported with kernel version 2.6.27.37-0.1.1 or later
 - Brocade HBAs are supported on RHEL5U3, RHEL5U4, SLES10 SP2 and SLES11
-

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 4 Update 8:
 - device-mapper-1.02.28-2.el4 or later
 - device-mapper-multipath-0.4.5-35.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 RHEL 5 Update 3:
 - device-mapper-1.02.28-2.el5 or later
 - device-mapper-multipath-0.4.7-23.el5 or later
- For RHEL 5 Update 4:
 - device-mapper-multipath-0.4.7-30.el5 or later
 - device-mapper-1.02.32-1.el5 or later
- For RHEL 5 Update 5:
 - device-mapper-multipath-0.4.7-34.el5 or later
 - device-mapper-1.02.39-1.el5 or later
- For SLES 10 SP2:
 - device-mapper-1.02.13-6.14 or later
 - device-mapper-devel-1.02.13-6.14 or later
 - multipath-tools-0.4.7-34.43 or later
- For SLES 10 SP3:
 - device-mapper-1.02.13-6.14 or later
 - device-mapper-devel-1.02.13-6.14 or later
 - multipath-tools-0.4.7-34.50.10 or later
- For SLES 11:
 - device-mapper-1.02.27-8.6 or later
 - multipath-tools-0.4.8-40.4.1 or later

Installing HPDM Multipath Enablement kit 4.4.1

To install HPDM Multipath 4.4.1, complete the following steps:

1. Download the HPDM Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.4.1 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 example, /tmp/HPDMmultipath).
4. Unbundle the package by executing the following commands:

```
#cd /tmp/HPDMmultipath
#tar -xvzf HPDMmultipath-4.4.1.tar.gz
#cd HPDMmultipath-4.4.1
```

5. Verify that the directory contains the following files and folders:
 - INSTALL
 - README.txt
 - COPYING
 - bin
 - SRPMS
 - conf
 - docs
6. To install HPDM Multipath 4.4.1, execute the following command:

```
#./INSTALL
```

Configuring Device Mapper Multipath to enable HP arrays

This section describes the following:

- [Recommended device parameter values](#)
- [Setting up HPDM Multipath](#)
- [SAN configuration supported by DM Multipath](#)

Recommended device parameter values

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

For MSA2012fc/MSA2212fc/MSA2012i

```
device {
    vendor                "HP"
    product                "MSA2[02]12fc|MSA2012i"
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    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 EVA4x00/EVA6x00/EVA8x00

```
device {
    vendor                "HP"
    product                "HSV2[01]0|HSV300|HSV4[05]0"
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout          "/sbin/mpath_prio_alua /dev/%n"
    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 HP P2000 FC / P2000 FC/iSCSI

```
device {
    vendor                "HP"
    product                " P2000 G3 FC|P2000G3 FC/iSCSI "
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout          "/sbin/mpath_prio_alua /dev/%n"
    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 MSA2012sa/MSA2312sa/MSA2324sa

```
device{  
  
    vendor                "HP"  
    product               "MSA2012sa|MSA2312sa|MSA2324sa"  
    getuid_callout        "/sbin/hp_scsi_id -g -u -n -s /block/%n"  
    prio_callout          " /sbin/mpath_prio_alua /dev/%n"  
    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"  
    hardware_handler      "0"  
    path_selector         "round-robin 0"  
    path_grouping_policy  multibus  
    failback              immediate  
    rr_weight             uniform  
    no_path_retry         18  
    rr_min_io             1000  
    path_checker          tur  
  
}
```

For MSA2312fc/MSA2324fc/MSA2312i/MSA2324i

```
device{  
  
    vendor                "HP"  
    product               "MSA2312fc|MSA2324fc|MSA2312i|MSA2324i"  
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"  
    hardware_handler      "0"  
    path_selector         "round-robin 0"  
    prio_callout          " /sbin/mpath_prio_alua /dev/%n"  
    path_grouping_policy  group_by_prio  
    failback              immediate  
    rr_weight             uniform  
    no_path_retry         18  
    rr_min_io             100  
    path_checker          tur  
  
}
```

NOTE:

- For SLES 11, in the device section, replace

```
getuid_callout          "/sbin/scsi_id -g -u -s /block/%n"
```

with

```
getuid_callout          "/lib/udev/scsi_id -g -u /dev/%n"
```

- For SLES /RHEL 5 U4 or later, in the device section for MSA201 2sa, MSA231 2sa, and MSA2324sa, replace

```
getuid_callout          "/sbin/hp_scsi_id -g -u -n -s /block/%n"
```

with

```
getuid_callout          "/sbin/scsi_id -g -u -n -s /block/%n"
```

- For SLES 10 SP2/SLES 10 SP3/SLES 11, in the device section, replace

```
prio_callout            "/sbin/mpath_prio_alua %n"
```

with

```
prio                    alua
```

- In XP arrays, there are different LUNs, such as OPEN-<x>, 3390-3A, 3390-3B, OPC:3390-3C, 3380KA, 3380-KB, and OP-C:3380-KC where x = {3,8,9,K,T,E,V}.
The product strings for XP LUNs are based on these emulation types. A new device section must be added for each emulation type, because each product string requires a new device subsection.
OPEN-* is sufficient for the product string for all the XP LUNs with different OPEN emulations because, regular expressions are supported in the `/etc/multipath.conf` file.
- For more information on editing `/etc/multipath.conf` file, see the *Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays 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 mptsas parameters](#)
- [Configuring Brocade 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:
 - For operating systems using the native Qlogic drivers,

```
options qla2xxx ql2xmaxqdepth=16 qlport_down_retry=10 ql2xloginretrycount=30
```
 - For other operating systems using the HP Qlogic drivers,

```
options qla2xxx ql2xmaxqdepth=16 qlport_down_retry=10 ql2xloginretrycount=30 ql2xfailover=0 ql2xlbType=0 ql2xautorestore=0x00 ConfigRequired=0
```
2. Rebuild the `initrd` by executing the following commands:
 - For operating systems using the native Qlogic drivers, complete 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:
 - For SLES 10/SLES 11 operating systems: `mkinitrd -k <kernal> -i <initrd>`
 - For other operating systems: `#mkinitrd /boot/initrd-<version no.>.img `uname -r``
 - c. Edit the value for default parameter in `/boot/grub/menu.lst` file to boot with the new `initrd` image.
 - For other operating systems using the HP Qlogic drivers:

```
/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_deptg=16 lpfc_discovery_threads=32
```
 - In SLES 10/SLES 11 hosts, edit the `/etc/modprobe.conf.local` file with the following values:

```
options lpfc lpfc_nodev_tmo=114 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=14 lpfc_lun_queue_depth=16 lpfc_discovery_threads=32
```
2. Rebuild the `initrd` by executing the following commands:
 - For operating systems using the native Emulex drivers, complete 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:

- For SLES 10/SLES 11 operating systems: `#mkinitrd -k <kernal> -i <initrd>`
- For other operating systems: `#mkinitrd /boot/initrd-<version no.>.img `uname -r``

c. Edit the value for default parameter in `/boot/grub/menu.lst` file to boot with the new initrd image.

- For other operating systems using the HP Emulex drivers, execute the following command:

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

3. Reboot the host.

Configuring iSCSI parameters

To configure the iSCSI parameters, complete the following steps:

1. Update the iSCSI configuration file

- In RHEL 5, SLES 10, and SLES 11 hosts, edit the `/etc/iscsi/iscsid.conf` file with the following value:

```
node.session.timeo.replacement_timeout=15
node.startup=automatic
```

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

```
ConnFailTimeout=15
```

2. Restart the iSCSI service by executing the following command:

- In RHEL 4/RHEL 5 hosts,

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

- In SLES 10/SLES 11 hosts,

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

Configuring mptsas parameters

To configure the mptsas parameters for RHEL 5 and SLES hosts, complete the following steps:

1. Edit the `/etc/modprobe.conf` file in RHEL 5 hosts and `/etc/modprobe.conf.local` file in SLES hosts with the following values:

```
options mptsas mpt_cmd_retry_count=10 mpt_disable_hotplug_remove=1
```

2. Rebuild the initrd by executing the following commands:

- 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:

- For SLES 10 operating systems:
`#mkinitrd -k <kernal> -i <initrd>`

- For RHEL operating systems:

```
#mkinitrd /boot/<initrd-version no.>.img `uname -r`
```

- c. Edit the value for default parameter in `/boot/grub/menu.lst` file to boot with the new `initrd` image.

Configuring Brocade HBA parameters

To configure the Brocade HBA parameters for RHEL 5 and SLES hosts, set the time out value by executing the following command:

```
# bcu fcpiim --mpiomode <port_ID> off 14
```

SAN configuration supported by DM Multipath

Table 2 lists the maximum SAN configuration supported by DM Multipath.

Maximum number of LUNs supported	512
Maximum number of paths per LUN	32
Maximum number of HBAs	8
Total number of SAN devices	2048

NOTE:

If the total number of LUNs is 512, each LUN can have 4 paths which leads to 2048 (512*4) devices. Above maximum SAN configuration supported by DM Multipath is only on SLES 11 operating systems.

Known issues

Following are the known issues in the HPDM Multipath 4.4.1 release:

- `multipath` commands may take longer time to execute on heavily loaded servers or under path failure conditions.
- Blacklisting the multipath device in the `/etc/multipath.conf` file 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>
```
- 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.
- `multipath -l` command may not reflect the correct path status for Logical Units presented from MSA2xxxsa array when paths fail or are restored under heavy load conditions. To refresh the path status, execute the `# multipath -v0` command.
- `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 the `/etc/multipath.conf` file or by using the `ulimit -n` command and restart the `multipathd` demon.
- 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.

- If an existing LUN is deleted or unrepresented from RHEL host, a DM multipath device with the invalid WWN may be created which cannot be used and will be removed after the system reboots.
- For LUNs greater than 2TB in RHEL 4 operating systems, DM multipath devices may not be created with appropriate size.
- On RHEL 4 operating systems with large number of iSCSI devices, not all multipath devices may get created after a reboot. It is recommended to increase the ESTABLISHTIMEOUT value in the `/etc/sysconfig/iscsi` file depending on the number of LUNs, or run the `multipath -v0` command after the reboot.
- On SLES 11 operating systems:
 - Multipath may not always activate all partitions on reboot.
 - The `multipathd` daemon may fail to stop immediately after it is started in large SAN configurations.
 - The `multipathd` daemon may consume more memory in large SAN configurations.