

Release Notes

Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.2.0

Part number: AA-RWF9G-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 v4.2.0. This incremental release is to provide enablement for HP StorageWorks Disk Arrays EVA6400/EVA8400.

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 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 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:**

Multipathing for SAN Boot environment is supported on the following operating systems: RHEL 5 U2/RHEL 5 U3/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://h18006.www1.hp.com/storage/networking/bootsan.html>.

What's new

HPDM Multipath 4.2.0 provides the following additional features:

- Provides support for HP StorageWorks Disk Arrays EVA6400 and EVA8400.
- Provides changes for installation on RedHat based distributions.



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 HBA driver parameters on RHEL 5 U3, see [Setting up HPDM Multipath](#).

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	RHEL 4 Update 6 RHEL 4 Update 7 RHEL 5 Update 2 RHEL 5 Update 3 SLES 9 - SP4 SLES 10 - SP2
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 J200P19 or later MSA2012i fw rev J210R10 or later MSA2012sa fw rev J300P13-01 or later (MSA2312fc/MSA2324fc) fw rev M100R18 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/ SLES 9) 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) 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/RHEL 4 U6) available at: http://h18006.www1.hp.com/products/storageworks/fca2214/index.html</p> <p>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</p> <p>RHEL 5 U3 HBA drivers Qlogic: 8.02.00.06.05.03 or later Emulex: 8.2.0.33.3p or later</p>



NOTE:

- HPDM Multipath provides support for iSCSI devices on the following operating systems:
RHEL 5 U2/RHEL 5 U3/SLES 9 SP4/SLES 10 SP2.
- For more information on configuring iSCSI parameters, see [Configuring iSCSI parameters](#).

Installing Device Mapper Multipath tools

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

- For RHEL 5 Update 3:
device-mapper-1.02.28-2.el5 or later, device-mapper-multipath-0.4.7-23.el5 or later

NOTE:

For RHEL 4 U6 operating systems, install HP Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.0.0 and then install HPDM Multipath v4.2.0. For more information on installing HPDM Multipath v4.0.0, see the *Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays 4.0.0 Installation and Reference Guide*.

Prerequisites for SLES 10 SP2

Ensure the following RPMs are installed on SLES 10 SP2:

```
readline-devel sysfsutils aaa_base acl attr audit-libs autoconf automake
bash bind-libs bind-utils binutils bison brp-check-internal bzip2 coreutils
cpio cpp cracklib cvs cyrus-sasl db device-mapper device-mapper-devel
diffutils e2fsprogs file filesystem fillup findutils flex gawk gcc gdbm
gdbm-devel gettext gettext-devel glibc glibc-32bit glibc-devel glibc-locale
gpm grep groff gzip info insserv klogd less libacl libattr libcom_err
libgcc libgssapi libmudflap libnscd libstdc++ libtool libxcrypt libzio
m4 make man mktemp module-init-tools ncurses ncurses-devel net-tools
netcfg openldap2-client openssl pam pam-modules patch perl permissions
popt post-build-checks procinfo procps psmisc pwdutils rcs readline
readline-devel rpm sed strace sysfsutils sysvinit tar tcpd texinfo timezone
unzip util-linux vim zlib zlib-devel
```

Prerequisites for SLES 9 SP4

Ensure the following RPMs are installed on SLES 9 SP4:

```
aaa_base acl attr bash bind-utils bison bzip2 coreutils cpio cpp cracklib
cvs cyrus-sasl db devs diffutils e2fsprogs file filesystem fillup findutils
flex gawk gdbm-devel glibc glibc-devel glibc-locale gpm grep groff gzip
info insserv kbd less libacl libattr libgcc libselinux libstdc++ libxcrypt
m4 make man mktemp module-init-tools ncurses ncurses-devel net-tools
netcfg openldap2-client openssl pam pam-modules patch permissions popt
procinfo procps psmisc pwdutils rcs readline sed strace syslogd sysvinit
tar tcpd texinfo timezone unzip util-linux vim zlib zlib-devel autoconf
automake binutils device-mapper gcc gdbm gettext libtool perl
readline-devel rpm termcap udev
```

Installing HPDM Multipath Enablement kit 4.2.0

To install HPDM Multipath 4.2.0, complete the following steps:

1. Download the HPDM Multipath Enablement Kit for HP StorageWorks Disk Arrays v4.2.0 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-4.2.0.tar.gz
#cd HPDMmultipath-4.2.0
```
5. Verify that the directory contains README.txt, COPYING, INSTALL, bin, conf, SRPMS, and docs directories.
6. To install HPDM Multipath 4.2.0, 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](#)
- [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 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 EVA3000/EVA5000/EVA4x00/EVA6x00/EVA8x00

```
device
{
    vendor                "(COMPAQ|HP)"
    product               "HSV1[01]1|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        12
    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"
    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/hp_scsi_id -g -u -n -s /block/%n"
    prio_callout           "/sbin/mpath_prio_alua %d"
    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          12
    rr_min_io              1000
    path_checker           tur
}
```

For MSA2312fc/MSA2324fc

```
device
{
    vendor                "HP"
    product                "MSA2312fc|MSA2324fc"
    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 10 SP2, in the device section, replace

```
prio_callout      "/sbin/mpath_prio_alua %d"
```

with

```
prio              alua
```

- For SLES 10 SP2, in the device section for MSA 2012sa, 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 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` file in RHEL hosts and `/etc/modprobe.conf.local` file in SLES hosts with the following values:
 - For RHEL 5 U3 operating system using the native Qlogic drivers,

```
options qla2xxx ql2xmaxqdepth=16 qlport_down_retry=10 ql2xloginretry-count=30
```
 - For other operating systems using the HP Qlogic drivers,

```
options qla2xxx ql2xmaxqdepth=16 qlport_down_retry=10 ql2xloginretry-count=30 ql2xfailover=0 ql2xlbType=1 ql2xautorestore=0xa0 ConfigRequired=0
```
2. Rebuild the `initrd` by executing the following commands:
 - For RHEL 5 U3 operating system 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:


```
#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_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 commands:
 - For RHEL 5 U3 operating system 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:


```
#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, run the following command:

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

3. Reboot the host.

 **NOTE:**

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

Configuring iSCSI parameters

To configure the iSCSI parameters, complete the following steps:

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

```
node.session.timeo.replacement_timeout=15
```
 - In SLES 9 hosts, edit the `/etc/iscsi.conf` file with the following value:

```
ConnFailTimeout=15
```
2. Restart the iSCSI service using the following command:
 - In RHEL 5 and SLES 9 hosts,

```
#!/etc/init.d/iscsi restart
```
 - In SLES 10 hosts,

```
#!/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 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]
```

Known issues

Following are the known issues in the HPDM Multipath 4.2.0 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.
- 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.
- `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.
- 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 RHEL4 and SLES9 operating systems, DM multipath devices may not be created with appropriate size.

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.