

Linux DM Installation Details

A Dell Technical White Paper

PowerVault MD3200, MD3200i, MD3600i and MD3600f
Series Storage Arrays



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2010-2011 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, the *DELL* badge, and *PowerVault* are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

June 2011

Contents

- Device Manager Multipath Overview 2
- Device Manager Multipath Stack Overview..... 2
- Required Patches 2
 - SLES11 Installation Steps: 3
 - SLES10 SP3 Installation Steps: 5
 - RHEL 5.4 Installation Steps:..... 7
 - RHEL 5.5 Installation Steps:..... 9
 - RHEL 6.0 Installation Steps:..... 11
- Oracle Cluster Parameters 15

Device Manager Multipath Overview

This whitepaper describes the steps involved in configuring Device Mapper Multi-Path (DMMP) for the MD3200, MD3200i, MD3600i and MD3600f series (MD32/36 series) of Dell PowerVault Storage Arrays. Device Mapper (DMMP) is a generic framework for block devices provided by the Linux Operating System.

Device Mapper Multipath:

- Provides a single block device node for a multipathed Logical Unit
- Ensures that I/O is re-routed to available paths during path failure event
- Re-establishes previously failed paths (once they return)
- Provides DMMP features for newly added Logical Units

Device Manager Multipath Stack Overview

The Linux DM stack that works with Dell PowerVault MD32/36 series of arrays consists of three main components:

- DM multipath core: The generic multipath component that works with all storage devices.
- RDAC device handler: The hardware interface between the storage array and DM multipath core stack that performs hardware specific tasks.
- Multipath User Space Utility: Contains the multipath daemon *kpartx* utility and the Dell-provided *rescan_dm_devs* script.

Required Patches

The Dell PowerVault MD32/36 series DM Driver Packages includes patches to the DM multipath component and RDAC device handler required for full functionality. These packages are included in the resource DVD included with the product.

Operating System	Component	Minimum Version
SLES11 SP1	Scsi_dh_rdac DKMS	scsi_dh_rdac-1.5.2.2-3dkms.noarch.rpm
SLES11	Kernel Version	kernel-default-2.6.27.39-0.3.1
	Scsi_dh_rdac DKMS	scsi_dh_rdac-1.5.1.2-3dkms.noarch.rpm
	multipath-tools	multipath-tools-0.4.8-40.6.1.rpm
SLES10.3	Scsi_dh_rdac DKMS	scsi_dh_rdac-1.3.1.3-3dkms.noarch.rpm
RHEL 6.0 on MD36xxf	Scsi_dh_rdac DKMS	scsi_dh_rdac-1.6.1.1-1dkms.noarch.rpm
RHEL 5.5	Scsi_dh_rdac DKMS	scsi_dh_rdac-1.4.2.2-3dkms.noarch.rpm
RHEL 5.4	Scsi_dh_rdac DKMS	scsi_dh_rdac-1.4.1.3-3dkms.noarch.rpm

Details for scsi_dh_rdac package installation

The resource DVD included with the storage array automatically installs the RDAC device driver as part of management software installation. Here is the installation procedure for each OS:

SLES11 Installation:

1. Install SLES11 with the installation media provided.
2. Install the errata kernel 2.6.27.39-0.1 (available via Novell website) for the architecture. The following steps are for x86_64 architecture. If necessary, replace with the package names specific to your architecture.

```
#rpm -Uvh module-init-tools-3.4-70.7.1.x86_64.rpm
```

```
#rpm -ivh kernel-default-base-2.6.27.39-0.3.1.x86_64.rpm
```

```
#rpm -ivh kernel-default-2.6.27.39-0.3.1.x86_64.rpm
```

3. Reboot the system to boot up with 2.6.27.39-0.3.1 kernel.
4. Verify the version of the multipath tools package installed on the system using the `rpm -qa |grep multipath-tools` command. The package version should be multipath-tools-0.4.8-40.6.1 or above.

Use the following commands to install the package.

```
#rpm -Uvh multipath-tools-0.4.8-40.6.1.x86_64.rpm
```

```
#rpm -Uvh udev-128-13.3.1.x86_64.rpm
```

```
#rpm -Uvh kpartx-0.4.8-40.6.1.x86_64.rpm
```

5. Add the following to the `/etc/multipath.conf` configuration file:

```
defaults
{
    ...
    max_fds 8192
    user_friendly_names yes
    ...
}
blacklist {
    ...
    device {
        vendor "*"
        product "Universal Xport"
    }
}
devices {
    device {
        vendor "DELL"
        product "MD32xxi"
        path_grouping_policy group_by_prio
        prio rdac
        polling_interval 5
    }
}
```

Linux DM Installation Details

```
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features               "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
  }
device {
  vendor                  "DELL"
  product                 "MD32xx"
  path_grouping_policy    group_by_prio
  prio                    rdac
  polling_interval        5
  path_checker            rdac
  path_selector           "round-robin 0"
  hardware_handler        "1 rdac"
  failback                immediate
  features                 "2 pg_init_retries 50"
  no_path_retry           30
  rr_min_io               100
  prio_callout            "/sbin/mpath_prio_rdac /dev/%n"
}
devices {
  device {
    vendor                  "DELL"
    product                 "MD36xxi"
    path_grouping_policy    group_by_prio
    prio                    rdac
    polling_interval        5
    path_checker            rdac
    path_selector           "round-robin 0"
    hardware_handler        "1 rdac"
    failback                immediate
    features                 "2 pg_init_retries 50"
    no_path_retry           30
    rr_min_io               100
    prio_callout            "/sbin/mpath_prio_rdac /dev/%n"
  }
}
device {
  vendor                  "DELL"
  product                 "MD36xxf"
  path_grouping_policy    group_by_prio
  prio                    rdac
  polling_interval        5
  path_checker            rdac
  path_selector           "round-robin 0"
  hardware_handler        "1 rdac"
  failback                immediate
  features                 "2 pg_init_retries 50"
  no_path_retry           30
  rr_min_io               100
  prio_callout            "/sbin/mpath_prio_rdac /dev/%n"
}
}
```

6. Add DM-RDAC driver module parameter `rdac_blacklist` in `/etc/modprobe.conf` to support RDAC/MPP coexistence.
7. Copy the file named `99-storage-policy-fixed-drives.fdi` into `/usr/share/hal/fdi/policy/10osvendor/`:

Linux DM Installation Details

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- *- SGML *- -->
  <deviceinfo version="0.2">
    <device>
      <match key="@block.storage_device:storage.hotpluggable" bool="false">
        <match key="@block.storage_device:storage.removable" bool="false">
          <merge key="volume.ignore" type="bool">true</merge>
        </match>
      </match>
    </device>
  </deviceinfo>
```

8. Install DKMS package from the deployment DVD that was included with the PV array

```
#rpm -ivh dkms-2.1.2-1.noarch.rpm
```

9. Lay down the `scsi_dh_rdac` driver, version 1.5.1.2-3.

10. Add the following section to the end of `/etc/init.d/boot.local`

```
modprobe scsi_dh;
modprobe scsi_dh_rdac;
modprobe dm-multipath;
/etc/init.d/multipathd start;
```

11. Enable multipathd daemon using the command

```
#chkconfig multipathd on
```

12. Rebuild RAMdisk. Reboot the system, then refer to the Owner's Manual for more information.

SLES10 SP3 Installation Steps:

1. Install DKMS package from the deployment DVD that was included with the PV array.

```
#rpm -ivh dkms-2.1.2-1.noarch.rpm
```

2. Lay down the `scsi_dh_rdc` driver, version 1.3.1.2-3.

3. Add the following contents into `/etc/multipath.conf`:

```
defaults
{
  ...
  max_fds 8192
  user_friendly_names yes
  ...
}
blacklist {
  ...
  device {
    vendor "*"
    product "Universal Xport"
  }
  ...
}
devices {
```

Linux DM Installation Details

```
device {
    vendor                "DELL"
    product               "MD32xxi"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features               "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                "DELL"
    product               "MD32xx"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features               "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                "DELL"
    product               "MD36xxi"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features               "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                "DELL"
    product               "MD36xxf"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features               "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
}
}
```


Linux DM Installation Details

4. Add DM-RDAC driver module parameter `rdac_blacklist` in `/etc/modprobe.conf` to support RDAC/MPP coexistence
5. Add the following section to the end of `/etc/init.d/boot.local`

```
modprobe scsi_dh;  
modprobe scsi_dh_rdac;  
modprobe dm-multipath;  
/etc/init.d/multipathd start;
```

6. Enable multipathd daemon using the command

```
#chkconfig multipathd on
```

This command will enable multipathd during the boot up.

7. Rebuild RAMdisk. Reboot the system and then refer to the owner's Manual for more information.

RHEL 5.4 Installation Steps:

1. Install DKMS package from the deployment DVD that was included with the PV array:

```
#rpm -ivh dkms-2.1.2-1.noarch.rpm
```

2. Lay down DKMS driver package, version 1.4.1.2-3.

3. Edit `/etc/multipath.conf`:

If the default `blacklist` section is not commented out, comment out this section:

```
#blacklist {  
#   Devnode "*"   
#}
```

4. Add the following contents to `/etc/multipath.conf`:

```
defaults  
{  
    ...  
    max_fds 8192  
    user_friendly_names yes  
    ...  
}  
blacklist {  
    ...  
device {  
    vendor "*"   
    product "Universal Xport"  
}  
    ...  
}  
devices {  
    device {  
        vendor "DELL"  
        product "MD32xxi"  
        path_grouping_policy group_by_prio  
        prio rdac
```

Linux DM Installation Details

```
        polling_interval          5
        path_checker              rdac
        path_selector             "round-robin 0"
        hardware_handler         "1 rdac"
        failback                 immediate
        features                  "2 pg_init_retries 50"
        no_path_retry            30
        rr_min_io                100
        prio_callout             "/sbin/mpath_prio_rdac /dev/%n"
    }
device {
    vendor                        "DELL"
    product                      "MD32xx"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval            5
    path_checker                rdac
    path_selector               "round-robin 0"
    hardware_handler           "1 rdac"
    failback                   immediate
    features                   "2 pg_init_retries 50"
    no_path_retry              30
    rr_min_io                  100
    prio_callout               "/sbin/mpath_prio_rdac /dev/%n"
}
Device {
    vendor                        "DELL"
    product                      "MD36xxi"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval            5
    path_checker                rdac
    path_selector               "round-robin 0"
    hardware_handler           "1 rdac"
    failback                   immediate
    features                   "2 pg_init_retries 50"
    no_path_retry              30
    rr_min_io                  100
    prio_callout               "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                        "DELL"
    product                      "MD36xxf"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval            5
    path_checker                rdac
    path_selector               "round-robin 0"
    hardware_handler           "1 rdac"
    failback                   immediate
    features                   "2 pg_init_retries 50"
    no_path_retry              30
    rr_min_io                  100
    prio_callout               "/sbin/mpath_prio_rdac /dev/%n"
}
}
```

4. Add DM-RDAC driver module parameter `rdac_blacklist` in `/etc/modprobe.conf.local` to support RDAC/MPP coexistence.

- Comment out all lines in `/etc/udev/rules.d/40-multipath.rules`. Add the following lines at the end of the file:

```
# multipath wants the devmaps presented as meaningful device names
# so name them after their devmap name
SUBSYSTEM!="block", GOTO="end_mpath"
# KERNEL!="dm-[0-9]*", ACTION=="add", PROGRAM==" /bin/bash -c '/sbin/lsmmod |
/bin/grep ^dm_multipath'", RUN+="/sbin/multipath -v0 %M:%m"
KERNEL!="dm-[0-9]*", GOTO="end_mpath"
PROGRAM!="/sbin/mpath_wait %M %m", GOTO="end_mpath"
PROGRAM!="/sbin/dmsetup info -c --noheadings -j %M -m %m", GOTO="end_mpath"
RESULT!="*:~*:~*:~*:~*:~*:~*:mpath-*", GOTO="kpartx_check"
PROGRAM!="/sbin/dmsetup info -c --noheadings -o name -j %M -m %m", RESULT=="?*",
NAME="%k", SYMLINK="mpath/%c", RUN+="/sbin/kpartx -a -p p /dev/mapper/%c"
OPTIONS="last_rule"
LABEL="kpartx_check"
RESULT!="*:~*:~*:~*:~*:~*:~*:part*-mpath-*", GOTO="end_mpath"
PROGRAM!="/sbin/dmsetup info -c --noheadings -o name -j %M -m %m", RESULT=="?*",
NAME="%k", SYMLINK="mpath/%c"
OPTIONS="last_rule"
LABEL="end_mpath"
```

- Enable multipathd daemon using the command

```
#chkconfig multipathd on
```

This command will enable multipathd during the boot up.

- Rebuild RAMdisk. Reboot the system, then refer to the Owner's Manual for more information.

RHEL 5.5 Installation Steps:

- Lay down the `scsi_dh_rdac` driver, version 1.4.2.1-3.
- Make the following editions to `/etc/multipath.conf`:

If the default `blacklist` section is not commented out, comment out this section:

```
#blacklist {
# Devnode "*"
#}
```

- Add the following contents into `/etc/multipath.conf`:

```
defaults
{
...
    max_fds 8192
    user_friendly_names yes
...
}
blacklist {
...
device {
    vendor "*"
    product "Universal Xport"
}
}
```

Linux DM Installation Details

```
...
}
devices {
  device {
    vendor                "DELL"
    product                "MD32xxi"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features              "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
  }
  device {
    vendor                "DELL"
    product                "MD32xx"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features              "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
  }
  device {
    vendor                "DELL"
    product                "MD36xxi"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features              "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
  }
  device {
    vendor                "DELL"
    product                "MD36xxf"
    path_grouping_policy  group_by_prio
    prio                  rdac
    polling_interval      5
    path_checker          rdac
    path_selector         "round-robin 0"
    hardware_handler      "1 rdac"
    failback              immediate
    features              "2 pg_init_retries 50"
    no_path_retry         30
    rr_min_io             100
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
  }
}
```

```
}
```

3. Add DM-RDAC driver module parameter `rdac_blacklist` in `/etc/modprobe.conf.local` to support RDAC/MPP coexistence.
4. Rebuild RAMdisk. Enable `multipathd` daemon using the command:

```
#chkconfig multipathd on
```

This command will enable `multipathd` during the boot up.

5. Reboot the system and then refer to Owner's Manual for more information.

RHEL 6.0 Installation Steps:

1. Lay down `scsi_dh_rdac` driver, version 1.6.1.1.
2. Make the following editions to `/etc/multipath.conf`:

If the default `blacklist` section is not commented out, comment out this section:

```
#blacklist {  
# Devnode "*"   
#}
```

3. Add the following contents into `/etc/multipath.conf`:

```
defaults  
{  
  ...  
  max_fds 8192  
  user_friendly_names yes  
  ...  
}  
blacklist {  
  ...  
  device {  
    vendor "*"   
    product "Universal Xport"  
  }  
  ...  
}  
devices {  
  device {  
    vendor "DELL"  
    product "MD32xxi"  
    path_grouping_policy group_by_prio  
    prio rdac  
    polling_interval 5  
    path_checker rdac  
    path_selector "round-robin 0"  
    hardware_handler "1 rdac"  }  
}
```

Linux DM Installation Details

```
        failback                immediate
        features                 "2 pg_init_retries 50"
        no_path_retry            30
        rr_min_io                100
        prio_callout             "/sbin/mpath_prio_rdac /dev/%n"
    }
device {
    vendor                       "DELL"
    product                     "MD32xx"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval             5
    path_checker                 rdac
    path_selector                "round-robin 0"
    hardware_handler             "1 rdac"
    failback                    immediate
    features                    "2 pg_init_retries 50"
    no_path_retry                30
    rr_min_io                    100
    prio_callout                 "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                       "DELL"
    product                     "MD36xxi"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval             5
    path_checker                 rdac
    path_selector                "round-robin 0"
    hardware_handler             "1 rdac"
    failback                    immediate
    features                    "2 pg_init_retries 50"
    no_path_retry                30
    rr_min_io                    100
    prio_callout                 "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                       "DELL"
    product                     "MD36xxf"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval             5
    path_checker                 rdac
    path_selector                "round-robin 0"
    hardware_handler             "1 rdac"
    failback                    immediate
    features                    "2 pg_init_retries 50"
    no_path_retry                30
    rr_min_io                    100
    prio_callout                 "/sbin/mpath_prio_rdac /dev/%n"
}
}
```

4. Rebuild RAMinitrd. Enable multipathd daemon using the command:

```
#dracut --dorce
```

This command will enable multipathd during the boot up.

5. Reboot the system and then refer to Owner's Manual for more information.

SLES11 SP1 Installation Steps:

1. Lay down scsi_dh_rdac driver, version 1.5.2.1-3.
2. Make the following editions to `/etc/multipath.conf`:
3. If the default `blacklist` section is not commented out, comment out this section:

```
#blacklist {  
# Devnode "*"   
#}
```

4. Add the following contents into `/etc/multipath.conf`:

```
defaults  
{  
  ...  
  max_fds 8192  
  user_friendly_names yes  
  ...  
}  
blacklist {  
  ...  
device {  
  vendor "*"   
  product "Universal Xport"  
}  
  ...  
}  
devices {  
  device {  
    vendor "DELL"  
    product "MD32xxi"  
    path_grouping_policy group_by_prio  
    prio rdac  
    polling_interval 5  
    path_checker rdac  
    path_selector "round-robin 0"  
    hardware_handler "1 rdac"  
    failback immediate  
    features "2 pg_init_retries 50"  
    no_path_retry 30  
    rr_min_io 100  
    prio_callout "/sbin/mpath_prio_rdac /dev/%n"  
  }  
  device {  
    vendor "DELL"  
    product "MD32xx"  
    path_grouping_policy group_by_prio  
    prio rdac  
    polling_interval 5  
    path_checker rdac  
    path_selector "round-robin 0"  
    hardware_handler "1 rdac"
```

Linux DM Installation Details

```
        failback                immediate
        features                 "2 pg_init_retries 50"
        no_path_retry            30
        rr_min_io                100
        prio_callout             "/sbin/mpath_prio_rdac /dev/%n"
    }
device {
    vendor                       "DELL"
    product                     "MD36xxi"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval             5
    path_checker                 rdac
    path_selector                "round-robin 0"
    hardware_handler             "1 rdac"
    failback                    immediate
    features                    "2 pg_init_retries 50"
    no_path_retry                30
    rr_min_io                    100
    prio_callout                 "/sbin/mpath_prio_rdac /dev/%n"
}
device {
    vendor                       "DELL"
    product                     "MD36xxf"
    path_grouping_policy         group_by_prio
    prio                         rdac
    polling_interval             5
    path_checker                 rdac
    path_selector                "round-robin 0"
    hardware_handler             "1 rdac"
    failback                    immediate
    features                    "2 pg_init_retries 50"
    no_path_retry                30
    rr_min_io                    100
    prio_callout                 "/sbin/mpath_prio_rdac /dev/%n"
}
}
```

4. Add DM-RDAC driver module parameter `rdac_blacklist` in `/etc/modprobe.conf` to support RDAC/MPP coexistence
5. Enable multipathd daemon using the command

```
#chkconfig multipathd on
```

This command will enable multipathd during the boot up.

6. Copy file named `99-storage-policy-fixed-drives.fdi` into `/usr/share/hal/fdi/policy/10osvendor/`:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!-- -*- SGML -*- -->
<deviceinfo version="0.2">
<device>
    <match key="@block.storage_device:storage.hotpluggable" bool="false">
        <match key="@block.storage_device:storage.removable" bool="false">
            <merge key="volume.ignore" type="bool">true</merge>
        </match>
    </match>
</device>
```



```
</deviceinfo>
```

7. Add the following section to the end of `/etc/init.d/boot.local`:

```
modprobe scsi_dh;  
modprobe scsi_dh_rdac;  
modprobe dm-multipath;  
/etc/init.d/multipathd start;
```

8. Rebuild RAMdisk. Reboot the system and then refer to the Owner's Manual for more information

Oracle Cluster Parameters

In Oracle cluster configurations, avoid LUN thrashing between multiple initiators by setting the **failback** parameter in the MD32/36 device section to *manual* in `/etc/multipath.conf`. This change will disable LUN failback.