

Application Note



Enabling Linux Device Mapper Multipath on EonStor[®]

For Linux RHEL 5 Update 2, SLES 10
SP2, & Linux Enterprise SUSE 11

Abstract

This document describes how to enable multi-pathing configuration using Linux's Device Mapper service. With additional configuration files, customers using RedHat Enterprise Linux (RHEL) 5 Update 2 and SuSE Linux Enterprise Server (SLES) 10 Service Pack (SP) 2 operating systems can utilize multi-pathing benefits.

Infotrend[®] Technology, Inc.

Networked Storage Solution Provider

Revision 1.1a

Feb, 2010

Infotrend Technology makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Infotrend reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify of such changes.

Infotrend, Infotrend logo, and EonStor are registered trademarks of Infotrend Technology, Inc.

Benefits

Multi-pathing with Device Mapper offers the following features:

- I/O failover and failback: Transparent failover and failback of I/Os by rerouting I/Os automatically to an alternative path in the event of a path failure and redirecting them back when the failed path is restored.
- Choices with grouping policies:
 1. Path priority:
 - Provides priority to grouped paths based on Target Port Group Service (TPGS).
 - Provides static load balancing policy by assigning a high-priority path as the preferred route
 2. Group by serial no.
 - Paths are grouped together based on the firmware-generated serial number
 3. Failover-only:
 - Failover without load balancing by grouping paths into individual path groups
- I/O load balancing policies: Weighted Round Robin load balancing policy among grouped paths.
- Persistent device name: Device names stick to individual arrays across reboots and Storage Area Network.
- Ease of reconfigurations. Device Mapper 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: Access control via device exclusion by blacklisting unwanted devices.
- Path monitoring: Periodically monitoring path status and thus enables fast failover and failback.
- Attaching / Removing device online: Devices can be added or deleted to or from Device Mapper (DM) Multipath without rebooting the server, or disrupting other devices or applications.
- Boot from SAN: Provides multipathing for operating systems installed on SAN devices.

Device Mapper Multipath support matrix

Hardware and Software prerequisites:

Operating system versions

RHEL 5 Update 2, SLES 10 - SP 2, SUSE 11

Host Bus Adapters (HBA) and SAN switches

See <http://viproom.infortrend.com.tw/index.aspx> (VIP room access required)

Supported arrays: all EonStor ASIC400 models

Device Mapper Availability

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

- For SLES 10 SP2:

device-mapper-devel-1.02.13-6.14, multipath-tools-0.4.7-34.38, device-mapper-1.02.13-6.14

- For RHEL 5 Update 2:

device-mapper-1.02.24-1.el5, device-mapper-multipath-0.4.7-17.el5

- For SUSE 11:

device-mapper-1.02.27-8.6.i586, multipath-tools-0.4.8-40.1.i586.rpm

Complying Device Mapper Multipath with EonStor

Arrays

Step 1. Uninstall Infortrend's EonPath driver if it has previously been installed on your application server.

Step 2. Ensure the related RPMs are installed on the system.

Step 3. Configure your RAID systems and complete the mapping process, parameter tuning, etc.

Step 4. Enable the OS-embedded iSCSI initiator. For example:

```
rpm -ivh iscsi-initiator-utils-6.2.0.868-0.7.el5.i386.rpm
```

(Skip steps 4 through 8 if you are using the Fibre Channel SAN models)

Step 5. Check iSCSI initiator version.

```
iscsiadm -v
```

Step 6. Let initiators find your targets (logical drives mapped to host LUNs) by specifying iSCSI host port addresses.

```
iscsiadm -m discovery -t sendtargets -p 192.168.10.100:3260
```

```
iscsiadm -m discovery -t sendtargets -p 192.168.10.101:3260
```

```
iscsiadm -m discovery -t sendtargets -p 192.168.10.102:3260
```

```
iscsiadm -m discovery -t sendtargets -p 192.168.10.103:3260
```

Step 7. Log on these iSCSI targets.

```
iscsiadm -m node --loginall=all (can be abbreviated as -L all)
```

Step 8. Check if all targets are present to the OS file system.

```
cat /proc/partitions
```

Step 9. Open and edit the multipath.conf. Change the last section of the configuration file. The file can be found in **/etc/multipath.conf**. An example is given below. Note that for SUSE 11, the conf file is located elsewhere. Please check the Notice below for how to direct the callout.

For RedHat

```
devices {
```

Enabling Device Mapper Multipathing for EonStor Series

```
# If you are already using a valid configuration file and do not have a
# device subsection for EonStor arrays, please add the appropriate device
# subsection for the respective arrays from the entries below.
# If you already have a device subsection for EonStor arrays which has
# different parameters from the entries below, modify it appropriately.
# For EonStor arrays

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

}
```

For SLES

```
devices {

# If you are already using a valid configuration file and do not have a
# device subsection for EonStor arrays, please add the appropriate device
# subsection for the respective arrays from the entries below.
# If you already have a device subsection for EonStor arrays which has
# different parameters from the entries below, modify it appropriately.

# For EonStor arrays

device {
    vendor          "IFT*"
#   product         "*"
    path_grouping_policy  group_by_prio
    getuid_callout   "/sbin/scsi_id -g -u -s /block/%n"
```

Enabling Device Mapper Multipathing for EonStor Series

```
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 more details, please refer to Device Mapper documentation.

NOTE:

On SUSE 11, locate the multipath config file, and the callout directory to the following. The changes take place at the highlighted lines:

```
# For EonStor arrays

device {
    vendor          "IFT*"
#   product         "*"
    path_grouping_policy  group_by_prio
    getuid_callout  "/lib/udev/scsi_id --page=0x83 --whitelisted --device=/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
}
}
```

Step 10. Locate and open an iSCSI configuration file at `/etc/iscsi/iscsid.conf` (the file can also be `iscsi.conf`) and change the following parameters.

`node.conn[0].timeo.noop_out_interval = 30` (default is 5)

`node.comm[0].timeo.noop_out_timeout = 90` (default is 5)

(This step is not necessary for FC SAN models)

Step 11. Reboot your system.

Step 12. Restart the Multipath service.

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

```
Stopping multipathd daemon [ OK ]
```

```
Stopping multipathd daemon [ OK ]
```

Step 13. Check the availability of your multipath devices.

Check the device list:

```
multipath -ll
```

Devices should appear as follows:

```
mpath1 (3600d023100037dfb000000007e2ae926) dm-2 IFT,S16E-R1130
```

```
[size=20G][features=0][hwandler=0]
```

```
\_round-robin 0 [prio=1][active]
```

```
\_4:0:0:0 sdb 8:16 [active][ready]
```

```
\_round-robin 0 [prio=1][enabled]
```

```
\_5:0:0:0 sdc 8:32 [active][ready]
```

Check disk partitions, you should be able to see fault-tolerant links dm-2 and dm-3:

```
more /proc/partitions
```

```
major minor #blocks name
```

```
8 0 244198584 sda
```

```
8 1 104391 sda1
```

```
8 2 244091610 sda2
```

```
253 0 242057216 dm-0
```

```
253 1 2031616 dm-1
```

```
8 16 20971520 sdb
```

```
8 17 20964793 sdb1
```

```
8 32 59180032 sdc
```

```
8 33 59180016 sdc1
```

```
253 2 20971520 dm-2
```

```
253 3 20964793 dm-3
```

If you cannot see multi-path devices, see the note at the end of document.

Step 14. You may then start formatting and using these devices.

```
mkfs /dev/dm-0
```

```
mkfs /dev/dm-1
```

```
mkfs /dev/dm-2
```

```
mkfs /dev/dm-3
```

```
....
```

NOTE:

If you can not find multi-path devices, check the following:

Enabling Device Mapper Multipathing for EonStor Series

1. Check blacklist devices to see if the EonStor arrays have been indicated in the blacklist. And then restart the multipath devices. `/etc/multipath restart`. Device mapper's default is to blacklist all devices. Add # marks in front of blacklist lines to disable the blacklist setting.
2. Enter 5 spaces after the vendor name to cover the range of Device Mapper version variables, e.g., "IFT ". Restart multipath service.