

# **How to Guide: DMRAID Eventing - What To Expect**

---

Author: Brian Wood

Revision 0.7  
02/25/08

# Disclaimers

## Intel Open Source License

Copyright ©Intel Corporation All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Intel Corporation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE INTEL OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright © Intel Corporation 2008.

## Table of Contents

Overview.....	4
Files Referenced In This Document.....	4
Expected Output.....	4
1.What You Should Expect To See With "dmraid".....	4
2.What You Should Expect To See With "dm_dso_reg_tool".....	5
3.What You Should Expect To See With "dmsetup".....	5
4.What You Should Expect To See In Syslog.....	6
5.What You Should Expect To See In Your Email.....	7
Appendix.....	8
Troubleshooting.....	8
No Events Generated For Raid 1 Devices.....	8
User Cannot Register Volume Group With Dmeventd.....	8
Notes About Example Used In This Document.....	9
Other Reading/Learning Material.....	9

# Overview

This document will give some examples of what you should expect to see if you have dmraid eventing setup properly. We will be discussing in this document:

- dmraid
- dm\_dso\_reg\_tool
- dmsetup
- Syslog
- Logwatch/Email

*(\*Note: It is assumed that the user of this document has followed the instructions described in the reference documents listed below. If you have not done so do not expect to get help on your questions from the author of this document.)*

## Files Referenced In This Document

- HOWTO - LogWatch Setup For Dmeventd\_0.71.pdf
- HOWTO - How to Setup Device Mapper Raid Event Monitoring\_0.74.pdf

# Expected Output

## 1. What You Should Expect To See With "dmraid"

- I. Using a Raid 4/5 volume as an example, the "dmraid -s" command should display the following when the Raid 4/5 is functioning properly:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmraid -s
*** Group superset isw_ghgifjaah
--> Active Subset
name   : isw_ghgifjaah_raid45
size   : 625154560
stride : 128
type   : raid5_la
status : ok
subsets: 0
devs   : 3
spares : 0
[mpsd@dmraid-devhost linux-2.6.25-rc2]$
```

- II. Using this same command "dmraid -s" when a Raid 4/5 volume is broken (i.e. one of the drives in the array is disconnected) should

display the following:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmraid -s
*** Group superset isw_ghgifjaah
--> Active Subset
name   : isw_ghgifjaah_raid45
size   : 625154560
stride : 128
type   : raid5_la
status : broken
subsets: 0
devs   : 3
spares : 0
[mpsd@dmraid-devhost linux-2.6.25-rc2]$
```

## 2. What You Should Expect To See With "dm\_dso\_reg\_tool"

- I. Using a Raid 4/5 volume as an example, the "dm\_dso\_reg\_tool -m" command should display the following when the Raid 4/5 is functioning properly:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dm_dso_reg_tool -m
Device Name: isw_ghgifjaah_raid45
Registered DSO:          libdmraid-events.so
UUID:                   DMRAID-isw_ghgifjaah_raid45
Status:                 Active
Major Device #:         253
Minor Device #:         0
Read-only Device:       No
Error Events Recorded by Kernel: 0
SATA drives in this volume group: /dev/sdb /dev/sdc /dev/sdd
```

- II. Using this same command "dm\_dso\_reg\_tool -m" when a Raid 4/5 volume is broken (e.g., one of the drives in the array is disconnected) should display the following:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dm_dso_reg_tool -m
Device Name: isw_ghgifjaah_raid45
Registered DSO:          libdmraid-events.so
UUID:                   DMRAID-isw_ghgifjaah_raid45
Status:                 Active
Major Device #:         253
Minor Device #:         0
Read-only Device:       No
Error Events Recorded by Kernel: 1
SATA drives in this volume group: /dev/sdb /dev/sdc /dev/sdd
```

Notice the "Error Events Recorded by Kernel" has a value greater than "1" to indicate that there has been an error recorded.

## 3. What You Should Expect To See With "dmsetup"

- I. Using a Raid 4/5 volume as an example, the "dmsetup status" command should display the following when the Raid 4/5 is functioning properly:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmsetup status
isw_ghgifjaah_raid45: 0 625154560 raid45 3 8:16 8:32 8:48 1 AAA
```

II. Using this same command “dmsetup status” when a Raid 4/5 volume is broken (e.g., one of the drives in the array is disconnected) should display the following:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmsetup status
isw_ghgifjaah_raid45: 0 625154560 raid45 3 8:16 8:32 8:48 1 AAD
```

Notice the last string of letters has changed from “AAA” to “AAD” with the D signifying the “disabled” drive in the listing (in this example that would correspond to the “8:48” *major:minor* number).

III. Again, using a Raid 4/5 volume as an example, the “dmsetup info” command should display the following when the Raid 4/5 is functioning properly:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmsetup info
Name:                isw_ghgifjaah_raid45
State:               ACTIVE
Tables present:     LIVE
Open count:         1
Event number:       0
Major, minor:       253, 0
Number of targets:  1
UUID: DMRAID-isw_ghgifjaah_raid45
```

IV. And using the command “dmsetup info” when a Raid 4/5 volume is broken (e.g., one of the drives in the array is disconnected) should display the following:

```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmsetup info
Name:                isw_ghgifjaah_raid45
State:               ACTIVE
Tables present:     LIVE
Open count:         1
Event number:       1
Major, minor:       253, 0
Number of targets:  1
UUID: DMRAID-isw_ghgifjaah_raid45
```

Notice that the “`Event number`” has incremented from zero to indicate the number of kernel errors recorded against this raid volume.

## 4. What You Should Expect To See In Syslog

I. This is an example of the output you should see in the syslog file (/var/log/messages on a Fedora/RHEL Operating System):

```

Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: Processing Raid Volume "isw_eecebgiaa_raid0" for Events
Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: Stripe device, 8:16 (/dev/sdb) has reported an I/O
error.
Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: The kernel has recorded 7 event(s) against this
device.
Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: Associated Userspace
Names: /dev/sdb=Disabled /dev/sdc=Active
Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: Associated SATA Port Mapping: /dev/sdb=2 /dev/sdc=3
Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: Associated UUID: DMRAID-isw_eecebgiaa_raid0
Jan  7 06:40:55 dmraid-devhost dmeventd[3254]: End of event processing for Raid Volume
"isw_eecebgiaa_raid0"

```

This output will look very similar for Raid 0, Raid 1, and Raid 4/5 volumes that have been registered to be monitored by the dmeventd daemon for events. (The main difference will be the number of "Associated Userspace Names" displayed.)

## 5. What You Should Expect To See In Your Email

- I. After a kernel event for a raid device is passed to the dmeventd daemon a syslog message will be created that the Logwatch daemon will process on its next scheduled run. As part of this Logwatch processing an email will be sent to a recipient listed inside the dmeventd logwatch configuration scripts (see "*BKM - LogWatch Setup For Dmeventd\_0.71.pdf*" for more detailed information). Here's a sample of the text that's sent in the email message body to the registered recipient:

```

##### Logwatch 7.3.4 (02/17/07) #####
Processing Initiated: Mon Jan  7 06:41:01 2008
Date Range Processed: today
                    ( 2008-Jan-07 )
                    Period is day.
Detail Level of Output: 5
Type of Output: unformatted
Logfiles for Host: dmraid-devhost.localdomain
#####

----- Dmeventd Errors Begin -----

There were a total of 28 new log entries

Date                Message
-----
Jan  7 06:40:24:    Processing Raid Volume "isw_eecebgiaa_raid0" for Events
Jan  7 06:40:24:    Stripe device, 8:16 (/dev/sdb) has reported an I/O error.
Jan  7 06:40:24:    The kernel has recorded 1 event(s) against this device.
Jan  7 06:40:24:    Associated Userspace Names: /dev/sdb=Disabled /dev/sdc=Active
Jan  7 06:40:24:    Associated SATA Port Mapping: /dev/sdb=2 /dev/sdc=3
Jan  7 06:40:24:    Associated UUID: DMRAID-isw_eecebgiaa_raid0
Jan  7 06:40:24:    End of event processing for Raid Volume "isw_eecebgiaa_raid0"

----- Dmeventd Errors End -----

##### Logwatch End #####

```

# Appendix

## Troubleshooting

### ■ **No Events Generated For Raid 1 Devices**

If you're setup for event monitoring and a device within that volume group encounters an error there should be an event generated. If this does not occur as expected run the following command to make sure the device mapper table has the proper parameters for processing:

```
$ /sbin/dmsetup table  
  
isw_cijcgddfhh_vol0: 0 488390920 mirror core 2 131072 nosync 2 8:16 0  
8:32 0 1 handle_errors
```

In the output displayed to the screen you should see a "1 handle\_errors" at the end of the line describing the stored data for the Raid 1 volume group that is to be monitored. If this is not present you need to update your version of the utility "dmraid" to the latest version dmraid-1.0.0.rc\*\*.

*(\*Note: At the present time the latest release of dmraid "dmraid-1.0.0.rc14-6" is posted on the maintainers website; link is available at the end of this document. This and all subsequent versions of dmraid contain the necessary code for Raid 1 activation)*

### ■ **User Cannot Register Volume Group With Dmeventd**

There could be a couple of reasons why this might be happening:

- If you have a bootable volume group that the operating system has been installed on, *at this time*, there will not be a UUID associated when the system comes up. This will prevent the registration of this volume group with dmeventd.

*This is due to the initrd kernel image not having the updated dmraid activation patch. (In the case of the Anaconda installer for Fedora/RHEL this is where the changes would need to be integrated by the installers maintainer).*

- If you're using an old version of the dmraid utility that does not contain the code patch for generating a UUID for a volume group upon activation.

Using the utility "dmsetup" run "dmsetup info" at the command line and see if you have a UUID value in the output, if you don't you either have an old version of dmraid or the UUID patch has not been applied:



```
[mpsd@dmraid-devhost linux-2.6.25-rc2]$ sudo /sbin/dmsetup info
Name:                isw_ghgifjaah_raid45
State:               ACTIVE
Tables present:     LIVE
Open count:         1
Event number:       1
Major, minor:       253, 0
Number of targets:  1
UUID: DMRAID-isw_ghgifjaah_raid45
```

## Notes About Example Used In This Document

- All examples in this document assume the user has basic Linux knowledge.
- The Author takes no responsibility for what happens to your system by using this guide. All instructions and testing should be performed on non-production systems before deploying to general service.
- As of the time of this writing both Fedora 8 and RHEL5.x were able to build all of the items described in this document; your experiences may vary in future operating system releases.

## Other Reading/Learning Material

- DMRAID maintainers homepage: <http://people.redhat.com/heinzm/>
- Device Mapper homepage: <http://sources.redhat.com/dm/>
- Logwatch Utility homepage: <http://www.logwatch.org/>