

RAID1 Support Notes

The following are support notes for RAID1, included as a standard item in the XE2000 and XE3000 Xorcom IP-PBX lines.

Document version: 2.1

Modification date: November 4, 2012

1. Disk Partitioning

XE2000/3000 servers are supplied with a RAID1 mirror that is constructed from two SATA hard disks. Partitioning of the disks is as follows:

Drive	Partition	Mount point	Size
SATA 0	/dev/sda1	/boot	100 MB
	/dev/sda2	/	72 GB
	/dev/sda3	swap	380 MB
SATA 1	/dev/sdb1	/boot	100 MB
	/dev/sdb2	/	72 GB
	/dev/sdb3	swap	380 MB

These partitions are used to create the RAID1 devices as described in the table below:

RAID device	Members	Mount point
/dev/md0	/dev/sda1	/boot
	/dev/sdb1	
/dev/md1	/dev/sda2	/
	/dev/sdb2	
/dev/md2	/dev/sda3	/swap
	/dev/sdb3	

2. RAID Device Status Information

2.1. Method for obtaining RAID device status information manually

You can manually inspect RAID status using commands such as:

```
cat /proc/mdstat mdadm --query --detail /dev/md*
```

Following is an example of the commands output:

```
# cat /proc/mdstat

Personalities : [raid1]

md0 : active raid1 sdb1[1] sda1[0]
      104320 blocks [2/2] [UU]

md2 : active raid1 sdb3[1] sda3[0]
      1044160 blocks [2/2] [UU]

md1 : active raid1 sdb2[1] sda2[0]
      76999424 blocks [2/2] [UU]
```

where [UU] denotes the normal status of both partitions. If, for example, partition of the first hard disk is not available, then you will see [_U] instead.

```
# mdadm --query --detail /dev/md*

/dev/md0:
  Version : 00.90.03

  Creation Time : Thu Jul 3 06:57:24 2008

  Raid Level : raid1

  Array Size : 104320 (101.89 MiB 106.82 MB)

  Device Size : 104320 (101.89 MiB 106.82 MB)

  Raid Devices : 2

  Total Devices : 2

  Preferred Minor : 0

  Persistence : Superblock is persistent

  Update Time : Sun Jul 6 08:02:39 2008

  State : clean

  Active Devices : 2

  Working Devices : 2
```

Failed Devices : 0

Spare Devices : 0

UUID : 13e490d5:38bccb19:7190d57d:10210724

Events : 0.20

Number	Major	Minor	RaidDevice	State	
0	8	1	0	active sync	/dev/sda1
1	8	17	1	active sync	/dev/sdb1

/dev/md1:

Version : 00.90.03

Creation Time : Thu Jul 3 06:55:47 2008

Raid Level : raid1

Array Size : 76999424 (73.43 GiB 78.85 GB)

Device Size : 76999424 (73.43 GiB 78.85 GB)

Raid Devices : 2

Total Devices : 2

Preferred Minor : 1

Persistence : Superblock is persistent

Update Time : Sun Jul 6 09:02:52 2008

State : clean

Active Devices : 2

Working Devices : 2

Failed Devices : 0

Spare Devices : 0

UUID : 3a8fb1ef:7bae1705:37f57d6c:eb2abc34

Events : 0.2724

Number	Major	Minor	RaidDevice	State	
0	8	2	0	active sync	/dev/sda2
1	8	18	1	active sync	/dev/sdb2

Xorcom USA
 145 S. Jefferson Ave., Suite G
 Cookeville, TN 38501 USA
 Tel: 1-866-XORCOM1/1-866-967-2661
info.usa@xorcom.com


www.xorcom.com

Xorcom Ltd.
 POB 60
 D.N. Misgav 20174, Israel
 Tel: +972-4-9951999
info@xorcom.com

```
/dev/md2:

    Version : 00.90.03

Creation Time : Thu Jul  3 06:55:47 2008

    Raid Level : raid1

    Array Size : 1044160 (1019.86 MiB 1069.22 MB)

    Device Size : 1044160 (1019.86 MiB 1069.22 MB)

Raid Devices : 2

Total Devices : 2

Preferred Minor : 2

    Persistence : Superblock is persistent

Update Time : Thu Jul  3 07:11:27 2008

    State : clean

Active Devices : 2

Working Devices : 2

Failed Devices : 0

Spare Devices : 0

    UUID : 7d6d33b3:dba23ecf:7b8e0b1c:8aafc2cd

    Events : 0.4

Number   Major   Minor   RaidDevice State
  0         8       3         0     active sync   /dev/sda3
  1         8      19         1     active sync   /dev/sdb3
```

2.2. Method for obtaining RAID device status information by e-mail

The `mdadm` daemon is responsible for monitoring RAID devices status.

- Define the destination e-mail address for `mdadm` notification messages.

The e-mail address must be defined in the parameter `MAILADDR` in the `/etc/mdadm.conf` file. For example:

Xorcom USA
145 S. Jefferson Ave., Suite G
Cookeville, TN 38501 USA
Tel: 1-866-XORCOM1/1-866-967-2661
info.usa@xorcom.com



Xorcom Ltd.
POB 60
D.N. Misgav 20174, Israel
Tel: +972-4-9951999
info@xorcom.com

MAILADDR mymail@mycompany.com

- Restart the mdadm daemon:

```
/etc/init.d/mdmonitor restart
```

3. How to replace a faulty hard disk

In order to replace a failed hard disk you must perform the following actions:

1. Discover which hard disk has failed
2. Replace the faulty hard disk with the new one
3. Create a partition table on the new hard disk
4. Make the new disk bootable
5. Synchronize the disks

The detailed instructions are listed below.

3.1. Find the faulty disk

Please note that the hard disk connected to the SATA0 connector is defined as the first disk and the hard disk connected to the SATA1 connector is defined as the second disk. When both disks are connected then you can see the first disk as device `/dev/sda` and the second disk as device `/dev/sdb`. In `/proc/mdstat` you will see:

```
Personalities : [raid1]

md0 : active raid1 sdb1[1] sda1[0]
      104320 blocks [2/2] [UU]

md2 : active raid1 sdb3[1] sda3[0]
      1044160 blocks [2/2] [UU]

md1 : active raid1 sdb2[1] sda2[0]
      76999424 blocks [2/2] [UU]
```

But if, for example, the first disk is disconnected, then the second disk will be defined in Linux as device `/dev/sda` but in `/proc/mdstat` you will see:

```
Personalities : [raid1]

md0 : active raid1 sda1[1]
      104320 blocks [2/1] [_U]

md2 : active raid1 sda3[1]
```

```
1044160 blocks [2/1] [_U]
md1 : active raid1 sda2[1]
      76999424 blocks [2/1] [_U]
```

As you can see, instead of [UU] now there is [_U]. In other words, the first disk is disconnected. You can get more detailed information by using mdadm. For example:

```
# mdadm --query --detail /dev/md0
/dev/md0:
    Version : 00.90.03
  Creation Time : Thu Jul  3 06:57:24 2008
    Raid Level : raid1
    Array Size : 104320 (101.89 MiB 106.82 MB)
    Device Size : 104320 (101.89 MiB 106.82 MB)
    Raid Devices : 2
    Total Devices : 1
Preferred Minor : 0
    Persistence : Superblock is persistent

    Update Time : Thu Jul 10 10:35:27 2008
    State : clean, degraded

    Active Devices : 1
    Working Devices : 1
    Failed Devices : 0
    Spare Devices : 0

    UUID : 13e490d5:38bccb19:7190d57d:10210724
    Events : 0.42

   Number   Major   Minor   RaidDevice State
     0         0         0         0        removed
     1         8         1         1        active sync  /dev/sda1
```

Xorcom USA
145 S. Jefferson Ave., Suite G
Cookeville, TN 38501 USA
Tel: 1-866-XORCOM1/1-866-967-2661
info.usa@xorcom.com



Xorcom Ltd.
POB 60
D.N. Misgav 20174, Israel
Tel: +972-4-9951999
info@xorcom.com

3.2. *Replace the faulty disk*

If the failed drive was not removed automatically from RAID mirroring, then you must remove it manually. To remove the first disk partitions from the RAID devices use:

```
mdadm -remove /dev/md0 /dev/sda1
```

```
mdadm -remove /dev/md1 /dev/sda2
```

```
mdadm -remove /dev/md2 /dev/sda3
```

Now you can replace the faulty disk:

- a. switch off the XR2000/3000 device
- b. remove the cover
- c. replace the disk with 2.5" hard disk of equal or larger volume
- d. switch on the device

3.3. *Create partition table*

You can use `sfdisk` utility to copy the partition table from the good disk to the new disk. The example below shows how you can copy the partition table from the second disk to the first disk:

```
sfdisk -d /dev/sdb | sfdisk /dev/sda.
```

3.4. *Add the disk partitions to the RAID devices*

The example below shows how to add the first disk partitions to the RAID devices.

```
mdadm -add /dev/md0 /dev/sda1
```

```
mdadm -add /dev/md1 /dev/sda2
```

```
mdadm -add /dev/md2 /dev/sda3
```

3.5. *Make the new disk bootable*

Use command `grub` in order to make the new disk bootable. It will allow you to boot the device if the old disk should fail in the future. The example below shows how to make the first disk bootable:

```
grub --batch <<EOF
device (hd0) /dev/sda
root (hd0,0)
setup (hd0)
quit
EOF
```

If you will need to make the second disk bootable then run the following:

```
grub --batch <<EOF
device (hd1) /dev/sdb
root (hd1,0)
setup (hd1)
quit
EOF
```

3.6. Synchronize the disks

Now you will see in `/proc/mdstat`:

```
# cat /proc/mdstat

Personalities : [raid1]

md0 : active raid1 sdb1[1]
      104320 blocks [2/1] [_U]

md2 : active raid1 sdb3[1] sda3[0]
      1044160 blocks [2/2] [UU]

md1 : active raid1 sdb2[1]
      76999424 blocks [2/1] [_U]
```

As you can see there is a problem with `/dev/md0` and `/dev/md1`.

```
mdadm /dev/md0 -a /dev/sda1
mdadm /dev/md1 -a /dev/sda2
```

If everything is OK you will see that `mdadm` has started to synchronize the partitions:

```
# cat /proc/mdstat

Personalities : [raid1]

md0 : active raid1 sda1[0] sdb1[1]
      104320 blocks [2/2] [UU]

md2 : active raid1 sdb3[1] sda3[0]
      1044160 blocks [2/2] [UU]

md1 : active raid1 sda2[2] sdb2[1]
      76999424 blocks [2/1] [_U]

      [>.....] recovery = 0.9% (736960/76999424)
finish=32.7min speed=38787K/sec
```