

RocketRAID 18xx Controller Fedora Linux Installation Guide

Version 1.15

Copyright © 2006 HighPoint Technologies, Inc.

All rights reserved.

Last updated on Apr 20, 2006

Table of Contents

1 Overview	1
2 Installing Fedora Linux on RR18xx controller	1
Step 1 Prepare Your Hardware for Installation.....	1
Step 2 Check System BIOS Settings.....	1
Step 3 Prepare the Driver Diskette.....	1
Step 4 Install Fedora Linux	2
3 Installing RR18xx driver on an Existing System.....	3
Remove sata_mv module on Fedora Core 5	3
Step 1 Obtain the Driver Module	3
Step 2 Test the Driver Module	3
Step 3 Configure System to Automatically Load the Driver.....	4
Step 4 Configure System to Mount Volumes when Startup	5
4 Monitoring the Driver	5
Checking Devices Status.....	5
Rebuilding a Critical Array	5
Verifying RAID 1/RAID 5.....	6
5 Updating the Driver	6
6 Installing RAID Management Software.....	6
7 Uninstalling.....	6

1 Overview

The purpose of this document is to provide clear instructions on how to install and use RocketRAID 18xx Controller on Fedora Linux system.

2 Installing Fedora Linux on RR18xx controller

If you would like to install Fedora Linux onto drives attached to RR18xx controller, please perform the following operations:

Step 1 Prepare Your Hardware for Installation

After you attach your hard disks to RR18xx controller, you can use RR18xx BIOS Setting Utility to configure your hard disks as RAID arrays, or just use them as single disks.

Before installation, you must remove all the disk drives, which are not physically attached to RR18xx controller, from your system.

Note

If you have other SCSI adapters installed, you must make sure the RR18xx controller BIOS will be loaded firstly. If not, try to move it to another PCI slot. Otherwise you may be unable to boot up your system.

Step 2 Check System BIOS Settings

In your system BIOS SETUP menu, change **Boot Sequence** in such a way that the system will first boot from floppy or CDROM, and then from SCSI. Refer to your BIOS manual to see how to set boot sequence.

If your BIOS settings do not support such a boot sequence, you can first set it to boot from floppy or CDROM. After you finish installation, set SCSI as the first boot device to boot up the system.

Step 3 Prepare the Driver Diskette

Driver is contained in a floppy diskette image file, you can get it from our web site for free.

On a DOS or Windows system, you can make the Fedora driver diskette using rawrite.exe. It can be found on the Fedora Linux CD (under /dosutils). Just run it under a command window and follow its prompt.

On a Linux system, you can use the “dd” command to make the boot diskette. Insert a floppy disk into the floppy drive and type the command (example):

```
# dd if=fc2-i386.img of=/dev/fd0
```

Step 4 Install Fedora Linux

Installation steps for Fedora Core 5

- 1) Start installing Fedora Linux by booting from the installation CD.
- 2) On "**Welcome to Fedora Linux**" installation screen, a prompted label "**boot:**" will appear at the bottom of the screen. Type in "**linux dd updates**" (without quotation mark) and then press **Enter**.
- 3) When prompted "**Do you have a driver disk?**". Select "**Yes**". When prompted "**Insert your driver disk and press OK to continue**", insert the driver diskette in the floppy drive and then select "**OK**". The system will load RR18xx driver automatically.
- 4) When prompted for an update disk source, select **fd0**. The installer will read updates from the driver floppy.
- 5) Continue the installation as usual. **Do not remove the floppy diskette** until installation completed.

Installation steps for Fedora Core 1 & Core 3 & Core 4

- 6) Start installing Fedora Linux by booting from the installation CD.
- 7) On "**Welcome to Fedora Linux**" installation screen, a prompted label "**boot:**" will appear at the bottom of the screen. Type in "**linux dd**" (without quotation mark) and then press **Enter**.
- 8) When prompted "**Do you have a driver disk?**". Select "**Yes**". When prompted "**Insert your driver disk and press OK to continue**", insert the driver diskette in the floppy drive and then select "**OK**".
- 9) Now the system will load RR18xx driver automatically.
- 10) Continue the installation as usual.

Installation steps for Fedora Core 2

- 1) Start installing Fedora Linux by booting from the installation CD.
- 2) On "**Welcome to Fedora Linux**" installation screen, a prompted label "**boot:**" will appear at the bottom of the screen. Just press **Enter**.
- 3) When prompted the warning "**No hard drives have been found**", press "**ALT+F2**" to switch the shell on console 2. Insert the driver diskette into floppy drive and type the following commands:

```
# mkdir /dd
# mount /dev/fd0 /dd
# sh /dd/fedora-install-step1.sh
# umount /dev/fd0
```

Then press “**ALT+F1**” to switch back to installation screen and select “**No**” to continue.

- 4) Continue the installation as usual. You can refer to Fedora Linux installation guide. Make sure you have finished the next step before system reboot.
- 5) When installation finishes and prompts you to reboot the system, press “**CRL+ALT+F2**” to the shell and type the following commands:

```
# chroot /mnt/sysimage
# mount /dev/fd0 /mnt
# sh /mnt/fedora-install-step2.sh
# umount /dev/fd0
# exit
```

Then switch back to console 1 and finish the installation.

3 Installing RR18xx driver on an Existing System

Remove sata_mv module on Fedora Core 5

If you are using Fedora Core 5, you must remove the sata_mv module from system otherwise it will conflict with RR18xx driver:

```
# rm -f /lib/modules/`uname -r`/kernel/drivers/scsi/sata_mv.ko
# depmod
```

Step 1 Obtain the Driver Module

You can extract the module file from the file modules.cgz on the driver disk. Using the following commands:

```
# mount /dev/fd0
# cd /tmp
# gzip -dc /mnt/floppy/modules.cgz | cpio -idumv
```

Driver modules for all supported kernel versions will be extracted. You can find the driver module for your running kernel under the directory that matches your kernel version.

Step 2 Test the Driver Module

You can test out the module to ensure that it works for your system by changing working directory to the location where hptmv.o resides and typing in the command “**insmod hptmv.o**”. If you are using a distribution with 2.6 kernel it should be “**insmod hptmv.ko**”.

Sometimes insmod will report “**unresolved symbols**” when you attempt to load the

module. This can be caused by two ways:

- 1) You haven't loaded the SCSI module before loading `hptmv.o`. Try to load SCSI modules first.

E.g. **# modprobe scsi_mod**
 # modprobe sd_mod
 # insmod ./hptmv.o

- 2) You are using a kernel that is build off a different configuration with the driver. In this case the precompiled drivers cannot be used. You can build a driver for your kernel using Open Source package for RocketRAID 18xx controller.

To ensure the module has been loaded successfully, you can check the driver status by typing in the command "**cat /proc/scsi/hptmv/x**", where **x** is the filename you found under `/proc/scsi/hptmv/`. You should see the driver banner and a list of attached drives. You can now access the drives as a SCSI device (the first device is `/dev/sda`, then `/dev/sdb`, etc.).

Example

You have configured a RAID 0 array using 2 disks. It will be registered to system as device `/dev/sda`. You can use "**fdisk /dev/sda**" to create a partition on it, which will be `/dev/sda1`, and use "**mkfs /dev/sda1**" to setup a file system on the partition. Then you can mount `/dev/sda1` to somewhere to access it.

Step 3 Configure System to Automatically Load the Driver

Most likely, you will not want to type in "**insmod hptmv.o**" each time you boot up the system. Therefore you must install the module and tell the system about it. To install the module, type in the following commands (first change directory to where the proper `hptmv.ko` locates):

```
# cp hptmv.ko /lib/modules/^uname -r^/kernel/drivers/scsi  
# depmod
```

Then you should inform the system load the module when system boots up with the following command:

```
# echo "modprobe hptmv" > /etc/init.d/hptdriver  
# chmod 755 /etc/init.d/hptdriver  
# ln -sf /etc/init.d/hptdriver /etc/rc.d/rc3.d/S01hptdriver  
# ln -sf /etc/init.d/hptdriver /etc/rc.d/rc4.d/S01hptdriver  
# ln -sf /etc/init.d/hptdriver /etc/rc.d/rc5.d/S01hptdriver
```

Step 4 Configure System to Mount Volumes when Startup

Now you can inform the system to automatically mount the array by modifying the file `/etc/fstab`. E.g. you can add the following line to tell the system to mount `/dev/sda1` to location `/mnt/raid` after startup:

```
    /dev/sda1    /mnt/raid    ext3    defaults    0 0
```

4 Monitoring the Driver

Once the driver is running, you can monitor it through the Linux `proc` file system support. There is a special file under `/proc/scsi/hptmv/`. Through this file you can view driver status and send control commands to the driver.

Note

The file name is the SCSI host number allocated by OS. If you have no other SCSI cards installed, it will be 0. In the following sections, we will use `x` to represent this number.

Checking Devices Status

Using the following command to show driver status:

```
# cat /proc/scsi/hptmv/x
```

This command will show the driver version number, physical device list and logical device list.

Rebuilding a Critical Array

A RAID 1 array may become critical after a disk member fails. When an array is in critical status, it will loss the ability of fault tolerance until you finish rebuilding.

Generally rebuilding will automatically start if you have a spare disk or you have put back the failed disk. In these cases, the array only needs to be synchronized to ensure data consistency. If the array is broken, you must first add a disk to the array. To add a disk to an array and start rebuilding, you can use the following command:

```
# echo "hpt rebuild a,b" > /proc/scsi/hptmv/x
```

In the command, "a" is array number shown in the logical device list. "b" is channel number. E.g.

```
# echo "hpt rebuild 1,2" > /proc/scsi/hptmv/x
```

will rebuild the array with logical device number 1 using the disk on secondary channel.

If rebuilding cannot be automatically started, you can use command

```
# echo "hpt rebuild start" > /proc/scsi/hptmv/x
```

to start rebuilding. To stop the rebuilding process, use command

```
# echo "hpt rebuild stop" > /proc/scsi/hptmv/x
```

Verifying RAID 1/RAID 5

To RAID 1/RAID 5, verifying will ensure data consistency.

You can use the following command to start verifying:

```
# echo "hpt verify start a" > /proc/scsi/hptmv/x
```

To stop the verifying process, use command:

```
# echo "hpt verify stop a" > /proc/scsi/hptmv/x
```

In the command, "a" is array number shown in the logical device list.

5 Updating the Driver

- 1) If the original driver is installed in the system initrd file (e.g. when OS is installed to RR18xx controller), you can update the driver as below:

```
# cp hptmv.ko /lib/modules/`uname -r`/updates/  
# mkinitrd `uname -r`  
For FC2, the commands should be  
# cp hptmv.ko /lib/modules/`uname -r`/kernel/drivers/scsi/  
# mkinitrd --preload scsi_mod --preload sd_mod --with=hptmv  
/boot/initrd-`uname -r`.img `uname -r`
```

It's recommended to backup the original initrd file before you update the driver.

- 2) If the original driver is installed in the /lib/modules/`uname -r`/kernel/drivers/scsi/ directory, and load it by the script file (e.g. /etc/init.d/hptdriver) during the init process, please replace it with the new driver (hptmv.o or hptmv.ko)

6 Installing RAID Management Software

HighPoint RAID Management Software is used to configure and keep track of your hard disks and RAID arrays attached to RR18xx controller. Installation of the management software is optional but recommended.

Please refer to HighPoint RAID Management Software documents about more information.

7 Uninstalling

You can only uninstall the driver when your system is not booting from devices attached to RR18xx controller. Just remove the lines you added to /etc/fstab, and remove the file files you created in the /etc/init.d.