

The NetWinder Rescue-HOWTO

Ralph Siemsen, ralphs@netwinder.org

\$Revision: 1.5 \$, \$Date: 2000/11/19 17:15:26 \$

This document explains how to restore a NetWinder using the 'Rescue Filesystem' that is now being shipped on all units. This is expected to replace, for most cases, the other methods of rescue (which are described in the [<Disk-Update-HOWTO>](#)). This document also contains a section on troubleshooting common hardware and software problems on the NetWinder.

Contents

1	Introduction	2
1.1	Other source of information	2
2	Using the rescue partition	2
2.1	Overview	2
2.2	Booting the rescue partition	3
2.3	Configuring networking	3
2.4	Now what?	4
2.5	Using mount smb	4
2.6	Using mount nfs	5
2.7	Using ftprescue	5
2.8	Using nfsserver	5
2.9	Using smbserver	5
2.10	Actual installation	5
3	Rescue partition installation	6
3.1	Do I already have it?	6
3.2	Installing the image	7
3.3	If you don't have /dev/hda4	7
4	Troubleshooting	9
4.1	Normal boot procedure and terminology	9
4.1.1	NeTTrom / BIOS	9
4.1.2	Minikernel	9
4.1.3	Second stage NeTTrom	9
4.1.4	Main kernel	10

5 Misc	10
5.1 Author	10
5.2 To-do	10
5.3 History	10
5.4 Contributors	10
5.5 Legal stuff	10

1 Introduction

Recovering a NetWinder to "factory" conditions has traditionally been a difficult task, since it required a second computer, configured to act as a boot server for the NetWinder. Although this is a flexible solution for technical users, it can be quite difficult for novices to get it to work. Thus, the OfficeServer includes a 'Rescue' partition which eliminates the need for a boot server. This document explains how to use the rescue partition to recover a NetWinder to factory condition.

This document therefore is aimed primarily at users of the OfficeServer product, though owners of the Developer model can also benefit. The rescue partition software began shipping on all models beginning in October 1999. For machines shipped prior to that date, the rescue software may be retro-fitted (see chapter 3).

1.1 Other source of information

A lot of NetWinder-specific information can be found in my home page at

<<http://www.netwinder.org/~ralphs/>> including a number of other HOWTO's on disk images, kernel and firmware installation and usage.

There is also a wealth of information in the general Linux HOWTO's, most of which apply directly to the NetWinder as well. They can be found in many places on the net, including <<http://www.linux.org/help/howto.html>>. I'd particularly recommend the Ethernet-HOWTO, the NET-3-HOWTO, and (if networking is all new to you), the Networking-HOWTO. Actually all of them :)

2 Using the rescue partition

This chapter explains the various ways that the rescue partition can be used to recover a NetWinder back to its factory state. Keep in mind that this is normally a 'last resort' measure for fixing your system; often you can more easily repair the damage in other ways. The rescue partition can be used as an emergency boot device, which allows you to go and fix stuff on the main partition.

2.1 Overview

All recent NetWinder machines include a small (10 MB) rescue partition, that contains enough software to reformat the NetWinder's internal hard disk, and then to reinstall the normal software load. Naturally,

an image of all the files on the hard disk is also necessary; for OfficeServer this image is included on the CDROM in the `rescue` folder.

In the most common scenario, the OfficeServer CDROM is placed into a PC that has network access, so that the NetWinder can retrieve the data from the CDROM via the network. It is necessary for the PC to have enabled file sharing, and a ‘shared folder’ for the CDROM has to be created.

The NetWinder rescue partition is then booted, and networking is configured so that the PC can be reached. There are a series of scripts to guide you through the process of formatting and then mounting the NetWinder hard disk. Then, the drive image is retrieved from the CDROM on the PC and installed on the NetWinder hard disk. (There is also the option to fetch the drive image via FTP).

The following sections describe the process in greater detail.

2.2 Booting the rescue partition

When you need to use the NetWinder’s rescue partition, here are the steps to access it. You’ll need to connect a keyboard and monitor to the NetWinder to carry out this rescue process.

1. Turn on your NetWinder (or reboot it, if it was running)
2. Stop the autoboot sequence by as the NetWinder boots. (e.g. when it says ‘Press any key to abort autoboot’)
3. Type the following commands at the firmware prompt

```
setenv kerndev /dev/hda4
setenv rootdev /dev/hda4
boot
```

That will do it, the NetWinder will now boot from the rescue partition. In short time, a shell prompt will appear, along with a message telling you to run `netconfig` to configure the network.

2.3 Configuring networking

The `netconfig` script will allow you to set up a network interface. It will ask a number of questions about your network, such as the IP address and netmask to be used. Some options, like DNS servers and gateways, are not required if your rescue computer is on the same subnet.

The `netconfig` script will ask you which interface to use. Normally, the OfficeServer uses `eth1` (the 10/100-base-T port) for its internal gateway. So that is generally the one you would select. Then give an IP address and a netmask. The script will try to compute the broadcast address for you.

If you normally operate using DHCP, you’ll have to ‘guess’ a free IP address to be used during this rescue boot. Go to some other computer on your network, check out what its IP address is, and then add one or two to the number. You can use `ping` or other tools to verify that the address is free for use. Then enter the free address into the NetWinder’s script.

It is a good idea to test the network connection once it’s been configured. From the NetWinder you can try to `ping` another machine on your network. DNS name resolution might not work, but numeric IP’s should. Note that the rescue partition shell does not support job control, which means you cannot abort a `ping` with CTRL-C. Instead, you have to use `ping -c 5 aa.bb.cc.dd` which tells ping to only try 5 times.

2.4 Now what?

At this point, there are five possible options for re-imaging the NetWinder's hard disk. Three of them are quite common:

1. `mountsmb` is used if the rescue image is going to be loaded from a Windows 95/98/NT computer on your network,
2. `mountnfs` is used if the rescue image is going to be loaded via NFS from a unix system on your network, and
3. `ftprescue` is used if the rescue image will be downloaded by FTP from an FTP server.

In some cases, instead of connecting from the NetWinder to a rescue server, you'll want to turn the NetWinder into a server so that other computers can connect to it. If this seems like the same thing to you, then don't worry about it, and ignore the following options:

1. `nfsserver` turns your NetWinder into an NFS server, with the root filesystem exported to the whole network,
2. `smbserver` similarly turns the NetWinder into a Samba server, so that other (Windows) clients can connect to it.

These options are described further in the following sections. There are a few more helpful scripts that are used, `wipefs` which erases the hard disk, and `mountfs` which mounts the partitions in preparation for the untarring of the disk image.

2.5 Using mountsmb

This is the option that most people will use. It requires that you have a computer running Windows on your network. You place the OfficeServer CD-ROM into this machine and allow the CD to be shared across the network. Click on 'My Computer', then right-click on the CD-ROM icon. A menu will appear, select 'Properties' and then click on the 'Sharing' tab. Turn on sharing and give it a name, for example, 'CDROM'.

On the NetWinder, you should now run the `mountsmb` script. It will ask for the name of the Windows computer (if you don't know what it is, then go to the Windows machine, right-click on 'Network Neighborhood' and then click the 'Identification' tab). Next, you'll be prompted for the name of the share ('CDROM' in the example above). Finally, you should enter the username (which matches the name you used to log into Windows). The NetWinder will then try to establish the connection to the Windows machine.

If the connection fails, you'll have to check your settings carefully and try again. Make sure the network cables are plugged in and that you can `ping` the Windows computer from your NetWinder, and vice-versa. Try entering the computer name and share name in uppercase, as some Windows systems seem to want it that way. If your DNS server is dodgy or nonexistent, then you'll need to use the IP address of the Windows machine in place of its name.

Once the mount is successful, then the contents of the CDROM should be visible on the NetWinder. To verify, type `ls -l /mnt/rescue`. You should see a directory called 'recovery' (or 'Recovery') and inside that directory, the OfficeServer disk image. You can now skip down to the [2.10](#) () section to complete the process.

2.6 Using mountnfs

If you have other computers on your network that run a Linux or some other UNIX-like operating system, then this option is the one to use. Place the CDROM into the drive and then do whatever is necessary to mount and share the CD to the network. For Linux, this would mean mounting the disk (`mount /dev/cdrom /mnt/cdrom`) and then editing the `/etc/exports` file to allow the `/mnt/cdrom` directory to be shared. And then the NFS service would need to be restarted.

On the NetWinder, the `mountnfs` script will prompt you for the IP address (or name) of the rescue server, and the name of the share (e.g. `/mnt/cdrom`). It will then try to mount the volume so that it can be accessed on the NetWinder as `/mnt/rescue`.

If the mount fails, check the network cables, IP addresses, and the settings on your server. Try mounting the server from elsewhere on your network, to see if it is correctly configured. Often you have to restart both NFS and portmap services on the server. Try ping tests to verify that the NetWinder can talk to the server.

Once the mount is successful, then the contents of the CDROM should be visible on the NetWinder. To verify, type `ls -l /mnt/rescue`. You should see a directory called 'recovery' (or 'Recovery') and inside that directory, the OfficeServer disk image. You can now skip down to the 2.10 () section to complete the process.

2.7 Using ftprescue

To be written.

2.8 Using nfsserver

To be written.

2.9 Using smbserver

To be written.

2.10 Actual installation

At this point, the new disk image you want to install should be mounted under `/mnt/rescue` somewhere, and you should know the exact path and filename. Since the CDROM's have the old DOS limitations on filenames, you may find that the image is called something strange, like `os-1.0_2~.gz` when really it should be something more meaningful like `os-1.0-2.tar.gz`. In the following examples, just substitute the actual filename for the examples listed.

You can now proceed to erase the `hda1` and `hda3` partitions and then to transfer, via the network, the new disk image on to the empty partitions. Two scripts are provided to facilitate this process: `wipefs` is used to clear the two disk partitions, and `mountfs` sets the partitions up so they can be accessed from `/mnt/hdroot`.

Note: there is a bit of a bug in the early versions of the rescue system. If you type `cat /proc/version` and it reports linux version 2.2.9-3, then you will likely have trouble with formatting the two partitions. The format command (`mke2fs`) will fail randomly with a 'memory violation' error. If this happens to you, your

options are to replace the kernel with a newer version (2.2.12), or to repeat the command until it succeeds, or to use `rm -rf` to delete all the files instead of `mke2fs`.

After you've used `wipefs` and `mountfs`, the new disk image can be installed directly. Just to keep you on your toes, we did not include a script for doing this. You have to type the commands yourself:

```
cd /mnt/hdroot
tar zxvpf /mnt/rescue/recovery/os-1.0-2.tar.gz
```

Adjust the pathname on the `tar` command as necessary to reflect the actual path and filename where the new image is located. It is critical to use the 'p' option so that permissions will be set correctly on the files. The 'v' option can be omitted if you don't want to see the names of the files scrolling by.

It should take about 15 minutes to copy all the data across. Once it's done, you should wait a little longer (30 seconds or so) to let the data be flushed to disk. Then, type `exit`, wait until the message appears that it's safe to shutdown. Then press the reset button to reboot. At this point, the new image will be loaded and hopefully all will be well.

3 Rescue partition installation

This chapter explains how to install and use the 'rescue partition' software package. NetWinder OfficeServer and DM models shipped after October 1999 include this software package by default; older systems need to be retrofitted (or sent back for upgrade) in order to make use of the new package.

3.1 Do I already have it?

If you've received your machine after October 1999, then you should already have the rescue package installed on your system. To be sure, there are two things to check. As `root`, run the command `fdisk -l /dev/hda`. This will list the current partition table, which should look something like this:

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1		1	3895	1963048+	83	Linux native
/dev/hda2		3896	4026	66024	82	Linux swap
/dev/hda3		4027	7921	1963080	83	Linux native
/dev/hda4		7922	7944	11592	83	Linux native

The rescue partition is `/dev/hda4`, and it's just a bit over 11 Megs in size. This is a pretty sure sign that you have the image, or at least, you have the space for the rescue image. To verify that the data is actually there, you need to mount the partition (temporarily):

```
mount /dev/hda4 /mnt
cd /mnt
ls
```

If the `mount` command fails with 'You must specify the filesystem type' then `/dev/hda4` probably is not formatted and therefore does not contain the rescue image. Otherwise, you should see a fairly standard directory structure listed:

```
bin  dev  lib          mnt  sbin  usr
boot etc  lost+found  proc tmp   var
```

If you see these directories, then you're all set. Note that from time to time, the rescue package will be updated, so it's a good idea to periodically install a newer version anyways. There currently isn't a way to find out which version of the rescue package you have installed, but in the future, we'll include a **README** file in the root directory (shown above) that will tell you which version you are looking at.

3.2 Installing the image

The following steps explain how to install the rescue image onto your system (or how to upgrade to a newer rescue image; it's the same procedure). I'm assuming that you do actually have a `/dev/hda4` partition of at least 10 Megs. See below for advice if you do not have this partition.

To install or update the rescue image on `/dev/hda4`, follow these steps:

1. Download the latest rescue image by anonymous FTP from <ftp://ftp.netwinder.org/pub/netwinder/images/> . The filename is `rescue.tar.gz` or there may be a newer version.
2. Log in as `root` or use the `su -` command to become root.
3. If you had previously mounted the partition, unmount it with the command `umount /dev/hda4`.
4. Format the hda4 partition, then mount it on `/mnt`:

```
mke2fs /dev/hda4
mount /dev/hda4 /mnt
```

5. Change directory to the mount point, and untar the rescue image.

```
cd /mnt
tar zxvpf /root/rescue.tar.gz
```

You will of course need to adjust the pathname on the `tar` command to reflect the location where you downloaded the rescue image.

3.3 If you don't have `/dev/hda4`

If you have an older system where the disk is already fully allocated to partitions 1 through 3, then it's a bit difficult to install the rescue system. I would recommend using one of the other rescue methods, which are described in the [Disk-Update-HOWTO.html](#) . Instead of installing the full disk image, though, you can repartition the drive and install the rescue package only. Then the rescue package can be used to reinstall everything else.

Another option is to try and merge two partitions together. If there is enough space free, then you can copy e.g. `/dev/hda3` over to `/dev/hda1`, and then can safely split 10MB or so off from `/dev/hda3` to be used as the rescue partition. Sadly, there is no way to resize an `ext2` partition without erasing the data on it. (There is `fips`, but that only works for DOS partitions).

Supposing you want to try this, then the first thing to do would be to run `df` to check how much disk space is available. It should look roughly like so:

```
Filesystem    1k-blocks      Used Available Use% Mounted on
/dev/hda1      1477028    301819  1098880  22% /
/dev/hda3      1521792    1151033  292110  80% /usr
```

In this case, there are about 1.15 Gig on `hda3` and only 1.09 Gig of space remaining on `hda1`, so it won't fit on `hda1`. It could be copied the other way (making `hda3` the root filesystem) but in that case you'd need to carefully adjust `/etc/fstab` to reflect that fact that the root filesystem is then on `/dev/hda3`, and remember to delete `/etc/mntab` before shutting down.

To copy the data between the partitions, you would use the following series of commands. Note that in my case, `/dev/hda3` was mounted as `/usr` (as indicated in the output from `df` above). On the older systems, it was mounted on `/home` instead. If that is the case for you, then substitute `home` for `usr` below.

```
umount /dev/hda3
mount /dev/hda3 /mnt
cp -avx /mnt /usr
umount /dev/hda3
```

Now you have to edit `/etc/fstab` and comment out (with the `#` character) the line that begins with `/dev/hda3` (You don't have to do this if you plan to move everything right back again, after having re-partitioned. Just don't reboot in the meantime).

You can then safely split `/dev/hda3` into two smaller pieces, using `fdisk /dev/hda`. First delete the entry for partition 3, then create a new primary partition 3. When prompted for the size, put in 10 MB less than you have left. You can either do the math (total cylinders divided by the total drive size times 10 MB) or just fiddle by trial and error.

Then create a 4th primary partition with the remaining 10 MB of space. Save the partition table, and format both partitions. You might also want to copy the stuff from back over from `/dev/hda1`:

```
mke2fs /dev/hda3
mke2fs /dev/hda4
# Now copy back /usr back from /dev/hda1 if desired:
mount /dev/hda3 /mnt
cp -avx /usr /mnt
umount /mnt
rm -rf /usr/.          # Careful with this !!
mount /dev/hda3 /usr
```

Don't forget to restore the `/etc/fstab` file if you changed it. Then you can install the rescue image onto `/dev/hda4` as described above.

4 Troubleshooting

4.1 Normal boot procedure and terminology

This section describes the stages that the NetWinder goes through when booting up, from the moment the power is applied until the login prompt appears. It also covers the common things that can go wrong.

4.1.1 NeTTrom / BIOS

When power is first applied, the first block of flash memory (64k) gets mapped in and executed. The first visible action is a quick probe of video ram, to determine how much memory there is. The screen is then cleared and the firmware version number and build date are displayed. Any logos that might be found are also rendered, along with the NetWinder logo animation if it is enabled. Meanwhile, the remainder of flash memory is read into RAM and the code therein is decompressed. There is a red progress meter shown at the bottom of the screen during this time. When the decompression is completed successfully, the screen fades to black, then the decompressed code is executed.

If the progress meter stops, then flash memory has been corrupted (or bad data was written to it). The only way to boot the NetWinder in this case is to hook up a serial terminal and to download a kernel via the serial port. For more details, see section 3.7 of the [<Firmware-HOWTO.html>](#) .

4.1.2 Minikernel

The system now boots into a small linux kernel. The screen clears and reverts back to text mode. In older versions, the full boot-up messages were displayed as the minikernel boots. In recent versions, only selected messages are shown to describe the hardware found. This kernel has the ability to mount a root filesystem in a variety of ways, as well as to fetch the main kernel in a variety of ways. There is a ‘firmware control menu’ available here.

Normally, the minikernel loads a real kernel from the hard disk. The parameters `kerndev` and `kernfile` specify the actual file in this case (default values are `/dev/hda1` and `/boot/vmlinux` respectively).

If an invalid kernel filename is given, the firmware will stop with an error message. The root filesystem however is a different matter: since it is not mounted until the kernel boots, the firmware cannot report if an invalid value is specified. So you won’t find out until later, when the kernel says `VFS: Unable to mount root fs` and proceeds to try booting from the non-existent floppy disk.

4.1.3 Second stage NeTTrom

After loading the main kernel into RAM, a reset is performed. Execution once again starts in the first block of flash code. However, this time it notices that its the second boot. Quickly, the RAM refresh is turned on and we jump directly to the main kernel.

If the main kernel is not bootable, the screen will stay dark at this point. This can also be caused by having inappropriate args passed from firmware to the main kernel (in particular, the amount of RAM on the system). Using old firmware with a new kernel will generally trigger this condition. Please see <http://netwinder.org/~ralphs/compat.html> for details on this.

4.1.4 Main kernel

The main kernel, generally loaded from disk, then goes through its normal boot sequence. Hardware is probed, devices are reported, and eventually the root filesystem gets mounted. This could fail, particularly if an NFS root is being used, for a variety of reasons.

Once the root filesystem is mounted, the kernel tries to start the `init` program, which will then run through the SysV-style init process. It will source `/etc/inittab`, which in turn sources `/etc/rc.d/rc.sysinit` and then all of the `/etc/rcN.d/S*` scripts (where N is the current runlevel, as defined in `inittab`). Finally, `getty`'s are launched on the various virtual consoles.

5 Misc

5.1 Author

The author and maintainer of the NetWinder Rescue-HOWTO is Ralph Siemsen (ralphs@netwinder.org). Please send me any comments, additions, corrections so that they can be included in the next release. The latest version of this document can be obtained from <http://www.netwinder.org/~ralphs/howto/Rescue-HOWTO.html> .

5.2 To-do

The 'sgml2info' version of this document doesn't show the examples properly - for some reason the linefeeds are removed. Why is this and how do I fix it?

5.3 History

Sep 21, 1999 (version 1.0): First public release of this document.

Nov 09, 1999 (version 1.1): Reorganization, and significant rewrite.

5.4 Contributors

Phil Petruzzo (philpe@rebel.com) contributed the section on how to install and use the rescue partition.

Douglas Paul (douglasp@netwinder.org) put together the rescue partition software.

5.5 Legal stuff

This document is copyright (c) Ralph Siemsen, 1999.

Permission is granted to make and distribute copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

There is no warranty whatsoever.