RED HAWK Linux®
Version 5.1.3
Release Notes

January 2009

0898003-5.1.3

concurrent

Disclaimer

The information contained in this document is subject to change without notice. Concurrent Computer Corporation has taken efforts to remove errors from this document, however, Concurrent Computer Corporation's only liability regarding errors that may still exist is to correct said errors upon their being made known to Concurrent Computer Corporation.

License

Trademark Acknowledgments

Concurrent Computer Corporation
2881 Gateway Drive
Pompano Beach, FL 33069

Note: Information subject to change without notice.

Contents

1.0. Introduction

This document provides release information and installation instructions for Concurrent Computer Corporation's RedHawk™ Linux® operating system Version 5.1, including kernel updates.

1.1 Product Description

RedHawk Linux is a real-time version of the open source Linux operating system. Modifications are made to standard Linux version 2.6 to support the functionality and the performance required by complex real-time applications. RedHawk uses a single kernel design to support a single programming environment that directly controls all system operation. This design allows deterministic program execution and response to interrupts while simultaneously providing high I/O throughput and deterministic file, networking, and graphics I/O operations. RedHawk is the ideal Linux environment for the deterministic applications found in simulation, data acquisition, industrial control and medical imaging systems.

Included with RedHawk is the popular Red Hat® Enterprise Linux 5 distribution. The RedHawk installation media provide additional real-time kernels and libraries for accessing RedHawk specific kernel features. Optionally, the NightStar™ RT development tool set is available for developing real-time applications, and the Frequency-Based Scheduler and Performance Monitor can be used for monitoring performance.

The RedHawk kernel integrates both open source patches and Concurrent developed features derived from the real-time UNIX® implementations that Concurrent has supported in its over 42 years experience developing real-time operating systems.

RedHawk is included with each Concurrent iHawk™ system. iHawks are symmetric multi-processor (SMP) systems available in a variety of architectures and configurations. Either 32-bit or 64-bit versions of RedHawk and its supporting software products are installed, depending upon the iHawk architecture type and the support included in RedHawk for that architecture.

Support for SMPs is highly optimized. A unique concept known as *shielded CPUs* allows a subset of processors to be dedicated to tasks that require the most deterministic performance. Individual CPUs can be shielded from interrupt processing, kernel daemons, interrupt bottom halves, and other Linux tasks. Processor shielding provides a highly deterministic execution environment that guarantees processes fast and dependable response to external events.

RedHawk Linux exceeds the level of POSIX® conformance of other Linux distributions based on the 2.6 series of kernels. Concurrent has added additional POSIX compliance by adding some of the POSIX real-time extensions that are not present in standard Linux. Linux on the Intel® x86 architecture has defined a defacto binary standard of its own which allows shrink-wrapped applications that are designed to run on the Linux/Intel x86 platform to run on Concurrent's iHawk platform.

NightStar RT is Concurrent's powerful real-time tool set that provides a robust graphic interface for non-intrusive control, monitoring, analysis, and debugging of real-time multiprocessing applications. All tools can be run natively on the same system as the application or remotely. The NightStar RT tools include the following:

- NightView™ source-level debugger – allows multi-language, multi-processor, multi-program and multi-thread monitoring and debugging from a single graphical interface.
- NightTrace™ run-time analyzer – analyzes the dynamic behavior of a running application.
- NightSim™ periodic scheduler – allows the user to easily schedule applications that require periodic execution.

- NightProbe™ data monitor – used to sample, record or modify program data in multiple running programs.
- NightTune™ performance tuner – used for analyzing system and application performance.

1.2 Product Media

The tables below list the CDs or DVDs that are delivered with RedHawk Linux Version 5.1 for each Concurrent system.

WARNING

While much of the software on the RedHawk media is licensed under the GNU GPL license, some components are not. Therefore, the media cannot be freely copied.

32-Bit Product Media

Processor Type:	**Intel® Xeon™ i386 and EM64T 32-bit**	
Model:	**iHawk 860 Series**	**[HQ685, 680, 665, 660, 645, 460, 295, 290, 285, 280, 265, 260, 195, 185, 175, 160, 075, 074, 069, 067, 065, 060, 058, 053, 049, 047 045, 036, 034, T74, T54, 000]**
	iHawk 880 Series	**[HQHS2]**
	ImaGen	**[HQ0G1]**
Operating System CD/DVDs:		
Red Hat Enterprise Linux 5.1 (32-bit) (1 DVD or 7 CDs)		
Red Hat Enterprise Linux 5.1 Source (32-bit) (1 DVD -- CDs available by request)		
Red Hat Enterprise Linux 5.1 Updates (32-bit) (1 DVD or 2 CDs) [LXUPDATE WA-IA32 Version 5.1]		
RedHawk Linux OS (32-bit) (1 DVD or 1 CD) [WA-IA32 Version 5.1]		
Optional Product CD/DVDs:		
RedHawk Linux Frequency-Based Scheduler (32-bit) (1 DVD or 1 CD) [WU1210-JL Version 5.1]		
NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version *x.x* (*x.x*=latest version)]		
RedHawk Linux PCI-to-VME Bridge Software Library for i386 & EM64T Systems [WU-HS630-LIB (WCS-SBS-620)] Version 5.1		
RedHawk Linux Cluster Manager (32-bit) [WA9017-L] Version 5.1		
RedHawk Global File System for i386 & EM64T Systems [WA9018-L] Version 5.1		
RedHawk High Availability NFS for i386 & EM64T Systems [WA9019-L] Version 5.1		

64-Bit Product Media

Processor Type:	**AMD Opteron™ 64-bit and EM64T x86_64 64-bit**
Model:	**iHawk 860 Series [HQ685, 680, 285, 280, 185, 069, 067, 049, 047, T74, T54, 000]** **iHawk 870 Series [HR844, 430, 221, 210, 00W, 00T]** **iHawk 880 Series [HQHS2]** **ImaGen [HR0G1]**
Operating System CD/DVDs:	
Red Hat Enterprise Linux 5.1 (64-bit) (1 DVD or 7 CDs)	
Red Hat Enterprise Linux 5.1 Source (64-bit) (1 DVD -- CDs available by request)	
Red Hat Enterprise Linux 5.1 Updates (64-bit) (1 DVD or 2 CDs) [LXUPDATE WA-AMD64 Version 5.1]	
RedHawk Linux OS (64-bit) (1 DVD or 1 CD) [WA-EM64T-AMD64 Version 5.1]	
Optional CD/DVDs:	
RedHawk Linux Frequency-Based Scheduler for x86_64 Systems (1 DVD or 1 CD) [WU1210-JA Version 5.1]	
NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version *x.x* (*x.x*=latest version)]	
RedHawk Linux PCI-to-VME Bridge Software Library for x86_64 Systems [WU-HS630-LIB (WCS-SBS-620)] Version 5.1	
RedHawk Linux Cluster Manager (64-bit) [WA9017-L] Version 5.1	
RedHawk Global File System for x86_64 Systems [WA9018-L] Version 5.1	
RedHawk High Availability NFS for x86_64 Systems [WA9019-L] Version 5.1	

1.3 Version Numbers

All RedHawk Linux version numbers take the following form:

major.*minor*[.*update*]

where:

major	is the product's major version number
minor	is the product's minor version number
update	is the update's incremental version number

For example, 5.1 is the initial release of major version 5, and 5.2 is a subsequent release to version 5.1. Both 5.1 and 5.2 constitute complete product releases and do not require a previous release to be currently installed on the system; both are available on CD and DVD from Software Distribution.

Individual kernel updates available via the RedHawk Updates website (see the "System Updates" section below) are not a complete product and can only be installed as an update to the matching *major*.*minor* release. For example; a 5.1.1 kernel update can only be installed on a system currently running RedHawk Linux version 5.1.

1.4 System Updates

As RedHawk Linux updates are issued, they are made available for downloading from Concurrent's Software Portal.

NOTE

Concurrent does not recommend downloading Red Hat updates.

The RedHawk Linux kernel replaces the standard Red Hat kernel and is likely to work with any version of the Red Hat distribution. However, installing upgrades, especially to **gcc** and **glibc**, from sources other than Concurrent may destabilize the system and is not recommended. Security updates from outside sources may be installed if desired.

Instructions for downloading updates from the website are provided in the section "Software Updates" on page 63.

1.5 Syntax Notation

The following notation is used throughout this document:

italic	Books, reference cards, and items that the user must specify appear in *italic* type. Special terms may also appear in *italic*.
`list bold`	User input appears in **`list bold`** type and must be entered exactly as shown. Names of directories, files, commands, options and man page references also appear in **`list bold`** type.
`list`	Operating system and program output such as prompts, messages and listings of files and programs appears in `list` type.
`[]`	Brackets enclose command options and arguments that are optional. You do not type the brackets if you choose to specify these options or arguments.
hypertext links	When viewing this document online, clicking on chapter, section, figure, table and page number references will display the corresponding text. Clicking on Internet URLs provided in ***blue*** type will launch your web browser and display the web site. Clicking on publication names and numbers in ***red*** type will display the corresponding manual PDF, if the file is available and accessible.

2.0. Documentation

2.1 RedHawk Linux Documentation

Table 2-1 lists the documentation provided with RedHawk Version 5.1 as PDF files. After installing RedHawk Linux documentation, the files are available for viewing:

- by clicking on the "Documents" icon on the desktop
- by clicking on the red document entry hyperlinks within the documents, including Table 2-1 below. Optional product documentation is available only if the product has been installed.
- from Concurrent's documentation web site at **http://redhawk.ccur.com/docs**

Table 2-1 RedHawk Linux Documentation

RedHawk Linux Operating System Documentation	Document Number
RedHawk Linux Online Documentation Roadmap	0898002-531
RedHawk Linux Version 5.1.3 Release Notes	0898003-5.1.3
RedHawk Linux User's Guide	0898004-610
Real-Time Clock & Interrupt Module (RCIM) User's Guide	0898007-400
iHawk Optimization Guide	0898011-110
RedHawk Linux FAQ	N/A
Optional RedHawk Product Guides	
RedHawk Linux Frequency-Based Scheduler (FBS) User's Guide	0898005-300
RedHawk High Availability NFS (HA-NFS) Installation Guide	0898018-100
RedHawk Global File System (GFS) Installation Guide	0898020-100
Optional RedHawk Product Documentation Sets	
RedHawk Linux Cluster Manager	
PCI-to-VME Bridge	

The Optional RedHawk Product Documentation Sets include Concurrent's user guides as PDF files along with additional vendor documentation PDFs that facilitate using these RedHawk products. By clicking on the red document entries, the product's Roadmap displays, from which all available documentation can be accessed.

The latest documentation for all Concurrent software products, including updated Release Notes, User Guides and FAQ, can be accessed from Concurrent's documentation web site at **http://redhawk.ccur.com/docs**.

Printed copies can be ordered by contacting the Concurrent Software Support Center. Refer to "Direct Software Support" on page 63 for details.

2.2 Vendor Documentation

Red Hat Enterprise Linux operating system documentation is included as PDF files on the Red Hat media delivered with RedHawk Linux.

Some RedHawk optional products include vendor documentation that facilitates using these RedHawk products.

Commercial off-the-shelf (COTS) documentation applicable to your iHawk system is included with the system. Contact the vendor’s sales office to purchase additional copies.

3.0. Prerequisites

3.1 Software

No previously installed software is required in order to install RedHawk Linux. The complete RedHawk Linux installation consists of the following software:

- Red Hat Enterprise Linux 5.1 with updates
- RedHawk Linux operating system
- NightStar RT for RedHawk tools (optional)
- RedHawk Linux Frequency-Based Scheduler (FBS) (optional)
- RedHawk Linux PCI-to-VME bridge software (optional)
- RedHawk Linux Cluster Manager (optional)
- RedHawk High Availability NFS (HA-NFS) (optional)
- RedHawk Global File System (GFS) (optional)

These packages operate in either a 32-bit or 64-bit environment. The operating environment on a system depends upon the processor type and the support included in the software for that processor. Refer to “Product Media” on page 2 for details.

NOTE

Real-time extensions and features are ***not*** available to 32-bit applications running under a 64-bit operating system (i.e. x86_64). In order to use real-time features, migrate 32-bit applications to 64-bit (described in Appendix E of the *RedHawk Linux User's Guide*) or boot a 32-bit operating system instead.

3.2 Hardware

- Any Concurrent iHawk or ImaGen system
- Real-Time Clock and Interrupt Module (RCIM)

4.0. Changes in this Release

This section describes enhancements and other changes introduced in Version 5.1.

4.1 Kernel Updates

The following kernel updates are available from the RedHawk Updates web site. These updates can be installed on a system currently running RedHawk Version 5.1. See the section "Software Updates" on page 63 for instructions for downloading and installing kernel updates.

Version 5.1.3 This kernel update includes:

- fix to prevent a SIGSEGV when writing to the user stack in the kernel during PTRACE_SETJUMPADDR int3 processing
- disabled CONFIG_ATA_OVER_ETH in the RedHawk kernel configuration files
- fix to prevent RCIM-III GPS lockups
- added `pci=norom` grub option
- fix for BUG() that occurred when a file that contained pages mmapped with **`/proc/pid/mem`**, **`usermap(3)`** or **`nprobe(1)`** was truncated or removed

Version 5.1.2 This kernel update includes:

- kernel support for GFS
- adjusted header line in **`run(1)`** to accommodate more than four hex digits (16 CPUs)
- updated FBS libraries to support up to 32 CPUs on a single host
- updates to IRQ shielding and IRQ CPU affinity support on x86_64 kernels with more than 8 cpus
- preallocated graphics page support for NVIDIA cards having a device class value of PCI_CLASS_DISPLAY_3D (0x0302)
- fix to prevent user segmentation faults during page replication
- fix to prevent potential corruption to kernel data structures tracking replicated pages
- eliminated reordering zonelists when a NUMA node is memory shielded on two-node systems
- fixed occasional very large erroneous values appearing in **`/proc/sys/kernel/tsc_to_nsec`** and **`/proc/sys/kernel/tsc_to_nsec_current`** (x86_64 only)

Version 5.1.1 This kernel update includes the following:

- Support has been added for the new RCIM III board, which provides Real-time Clock and Interrupt functions with a PCI Express connector. Chaining of RCIM III boards and use of distributed interrupt channels on RCIM III boards is not yet supported.
- Kernel support was added for HA-NFS.
- Intel gigabit network connection support has been updated to support the latest drivers. The igb driver was integrated and the e1000 driver was split out into e1000 and e1000e.

Driver	Version	Connection Type	IRQ type
e1000	8.0.1	PCI/PCI-X	Pin IRQ Assertion
e1000e	0.2.9.5	PCI/X/PCI Express	MSI
igb	1.2.24	PCI Express 2.0	MSI-X

- Sound support is now built entirely as modules.
- Support has been added for the Intel HD Audio.
- The **`memexact(1)`** tool has been updated so it does 64-bit arithmetic on 32-bit platforms (required for 32-bit PAE kernels addressing over 4GB of memory).
- A fix was made for the **`shield(1)`** command **`-m`** option to wait/block in the kernel until all memory shielding processing is complete before enabling/disabling memory shielding.
- A fix was made for a hang condition with **`shield -p`** kernel execution interfering with kernel stop machine execution.
- The rcim init script now recognizes the RCIM III and reports pass or fail “start/stop” status correctly.
- The ccur-rt RPM updates **`/usr/share/hwdata/pci.ids`** with a snapshot of the latest **`pci.ids`** registered at sourceforge. The RPM also replaces the **`/sbin/update-pciids`** script with a symlink to **`/usr/share/hwdata/update-pciids`**, which downloads the sourceforge **`pci.ids`** and updates **`/usr/share/hwdata/pci.ids`** so new devices are recognized by **`lspci(8)`** and **`lsirq(8)`**.

4.2 Enhancements in Version 5.1

4.2.1 RedHawk based on RHEL 5.1

RedHawk Version 5.1 uses Red Hat Enterprise Linux 5.1 as its base with a merge of kernel.org version 2.6.23.17.

4.2.2 UIO User-level Driver Support

The Userspace I/O (UIO) patch has been incorporated into RedHawk 5.1. UIO can be used to write user-level device drivers for any number of I/O boards. UIO requires a small per-device kernel module, with the main part of the driver written in user space, utilizing the tools and libraries commonly used for userspace applications.

A complete how-to document and sample UIO drivers are provided to guide you in writing your own user-level device drivers. Refer to Chapter 14 of the *RedHawk Linux User’s Guide.*

4.2.3 Expanded NUMA Support

NUMA node(s) can be shielded so that user pages are migrated to and from the shielded nodes’ memory to reduce the amount of remote memory accesses for real-time applications, and to reduce memory contention from other processes on the system. Shared read-only pages will be replicated if they reside on a node different from the requesting CPU.

The **shield(1)** command now contains the **--mem**/**-m** option for memory-shielding a NUMA node. Details about memory-shielded NUMA nodes can be viewed using **run(1)**, **cpu(1)** and a new utility, **numapgs(1)**.

Refer to Chapter 10 of the *RedHawk Linux User's Guide* and the man pages for details.

4.2.4 New Desktop Icon

RedHawk provides desktop icons that provide immediate access to software updates, documentation and Concurrent's web site. The former "Updates" icon is now labeled "Concurrent Software Portal," which displays Concurrent's software update web site, **redhawk.ccur.com**. You can use this web site to browse software repositories or to create a customized Update CD for your system. A new icon, labeled "Updates (NUU)" launches NUU to streamline the process of downloading RedHawk updates. See "Software Updates and Support" on page 63 for more information.

4.3 Modifications in Version 5.1

4.3.1 Alternative glibc Removed

Previous versions of RedHawk included an alternative **glibc** that provided robust POSIX mutexes and priority inheritance POSIX mutexes through pthread mutex services described in Chapter 5 of the *RedHawk Linux User's Guide*. Applications using these mutexes were compiled with **ccur-gcc** or **ccur-g++**.

The additional functionality provided by the alternative **glibc** is now part of standard **glibc**. Concurrent's alternative **glibc** and **ccur-*** compilation scripts have been removed.

The standard **glibc** additions are completely binary compatible with the extensions provided through the alternative **glibc** in previous versions of RedHawk. These applications, when loaded, will automatically redirect their dynamic library loading from the now-nonexistant alternative **glibc** to the standard **glibc**.

New applications that use priority inheritance and/or robust mutexes must be compiled with the standard **cc(1)**, **gcc(1)** and **g++(1)** tools. Note that this may require that Makefiles that specify **ccur-*** be changed to use the standard tools. Alternatively, symbolic links can be created in **/usr/bin** to point the names **ccur-gcc** and **ccur-g++** to **gcc** and **g++**, respectively.

4.3.2 Accessing PCI BAR Resources Depends on Kernel Tunable

Support for accessing the PCI base address registers (BARs) for use in user-level device drivers is now dependent upon a kernel parameter. This parameter is enabled by default in all RedHawk pre-built kernels. No change is required to use these interfaces. This feature is described in the Device Drivers chapter of the *RedHawk Linux User's Guide*.

4.3.3 kgdb Not Yet Available

kgdb is not available on RedHawk Version 5.1. It will be available in an upcoming kernel update.

4.3.4 Availability of RedHawk Optional Products

The Frequency-based Scheduler (FBS) was released at the same time as RedHawk Version 5.1; however, some other optional RedHawk products that usually accompany a RedHawk release were released independently.

Cluster Manager, SBS PCI-to-VMEBridge, HA-NFS, GFS and High Performance Math Package are now available and can be downloaded from the RedHawk Updates web site. See “Software Updates” on page 63 for instructions for downloading these products.

SIMulation Workbench will become available shortly.

The userland component for SNARE, SNARE Agent for Linux, can be downloaded from the **www.intersectalliance.com** web site. Instructions for configuring SNARE in a 5.1 RedHawk kernel are included in the RedHawk-FAQ. For the latest RedHawk FAQ, go to **http://redhawk.ccur.com/docs**.

5.0. Installation Procedures

A full system generation must be performed. If you are running previous versions of RedHawk or Red Hat, you must backup your files and then proceed with a completely new system generation.

Review the following points before installing the software:

- Concurrent recommends these procedures be followed as documented here; however, advanced users may decide to make other choices based on onsite requirements.
- The installation media in use is accessed at various times and must be kept in the drive during the installation until you are instructed to install different media.
- Prior to installation, remove any floppy media that may have been left in the machine.

Before you start:

- **Main Memory Size** — You need to know the system's main memory size in order to properly set up the file system swap size. If you are not sure of the main memory size, access the system's BIOS to ascertain this information.
- **Network Addresses** — If you do not plan to configure networking DHCP, you will need to have the following information for your system before beginning the installation:
  - IP address
  - netmask address
  - gateway address
  - primary and secondary DNS addresses
- **BIOS Settings** — BIOS settings should be reviewed before installation and whenever problems involving non-working devices are encountered. Exercise caution before changing default settings. Concurrent makes the following recommendations:
  - "Linux" should be specified for the "Installed OS" setting (or its equivalent) if it is an available option. If not, use the default "other."
  - Only the default settings for memory functions and PCI timing functions are supported.
  - ACPI support should be enabled, but care should be taken before changing other ACPI options from their defaults. ACPI sleep states should ***not*** be enabled in the BIOS or in any custom kernel build.

5.1 Installing Red Hat Software

The version of Red Hat Enterprise Linux that is appropriate to your system's architecture must be installed on the system first by specifically completing the steps below. Installation of Red Hat is accomplished in two parts:

- Red Hat Installation
- Installation of Red Hat Updates

5.1.1 Red Hat Installation

This section typically takes between 40-70 minutes to complete.

1. Power on the system, if needed, to open the CD or DVD drawer.
2. Insert Red Hat Enterprise Linux Installation Disc 1 in the drive and power down the system.
3. Power on the system with the installation disc installed to boot from the disc. After a short delay, a screen containing the `boot:` prompt should appear.

4. Press Enter (or allow this screen to timeout) to start the installation program using the GUI installation interface.

NOTES

If your system contains disk drives or devices that are not supported by the Red Hat installation program, you will need to obtain the necessary driver(s) in order to install Red Hat Enterprise Linux. A diskette containing the driver may be included with your system or it may be available on the Red Hat installation CD. If not available, contact the driver manufacturer for an updated driver or call Concurrent Support (see "Direct Software Support" on page 63).

For this type of installation, type `linux dd` at the `boot:` prompt. Insert the diskette in the drive when instructed to do so.

If your hard drive is not recognized and no driver diskette is supplied, a screen appears stating that no hard drives are found, would you like to select drivers. The installation disc contains a number of device drivers. Select Yes, then Add Device to display a list of drivers. Select the desired driver and proceed through the screens, completing this action by selecting Done.

If your system has a graphics card that is not recognized by the Red Hat installation program, the installation GUI will not appear and you will have to install Red Hat Enterprise Linux in **text mode**. Aids at the bottom of the screen describe how to navigate through the text mode screens and make selections. The instructions given below can be applied to both GUI and text mode installation; however, the number and order of the configuration screens differ slightly between the two modes. During RedHawk installation, your graphics card will be configured properly.

5. If you wish to check the Red Hat media, press Enter, otherwise, navigate to the Skip selection using the Tab key and press Enter to immediately start the installation.

   NOTE: The Source DVD media cannot be checked using this procedure because Red Hat does not include the required checksum in the DVD content.

6. At the Welcome screen, no action is needed. Press Enter to proceed.

7. Select your desired language and keyboard layout on the next two screens, pressing Next at each to proceed.

8. After selecting your keyboard layout, a dialog will appear which requests an installation number. Select the Installation Number radio button, enter **`6d70-bd33-6d3c-fcbc`** in the text field, and click the OK button.

9. The dialog as shown below will appear.

   Red Hat does not recommend upgrading from a previous version of Red Hat and these release notes assume that a full system generation will be done.

   Select the Install Red Hat Enterprise Linux Client radio button and click Next.

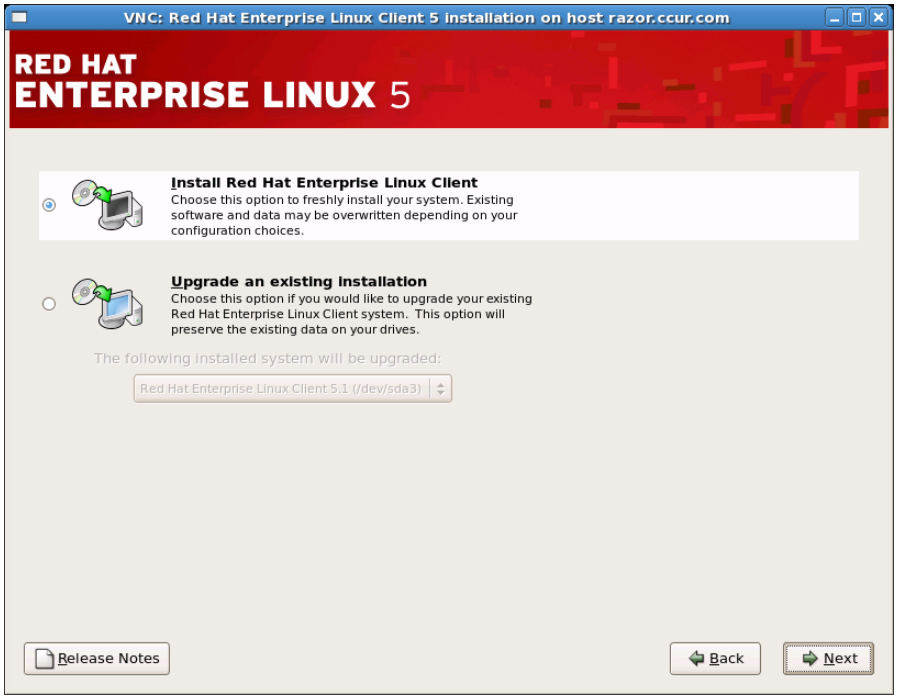

10. The next step involves partitioning the disk. Select the Create custom layout option from the option list and check any physical disks shown in the table that you may wish to partition. Then click the Next button.

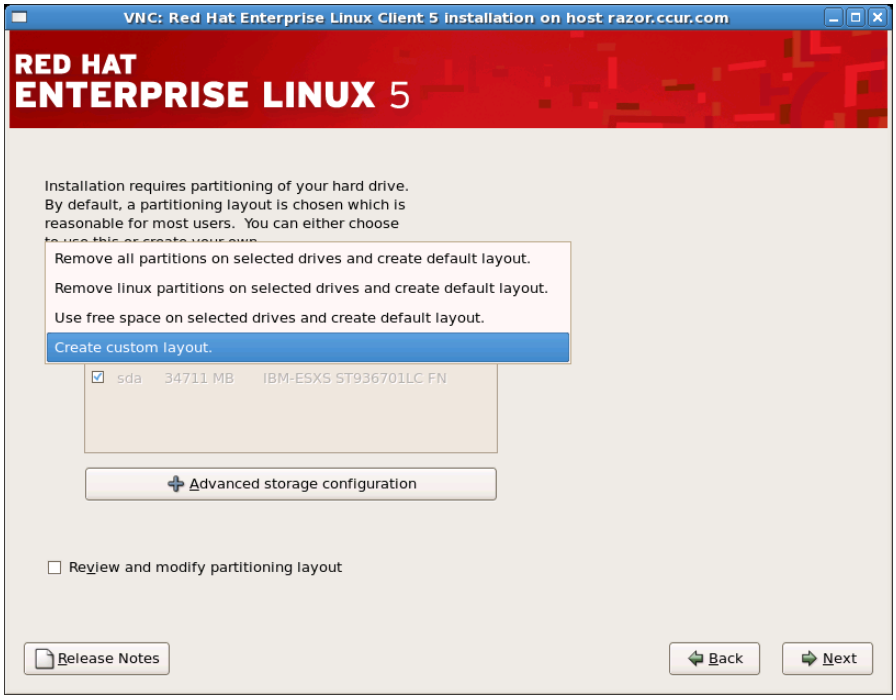

a. Select a disk of interest and delete all partitions on the disk so that all that remains is free space.

Warning! This action will erase all data on the disk when the disk is formatted in the next step.

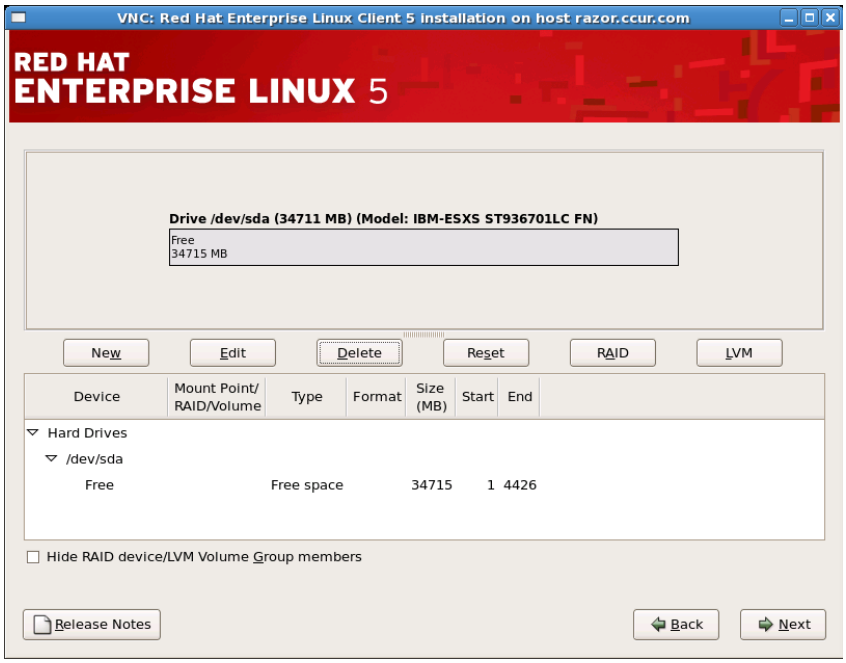

b. Now press the New button to configure the **/boot** partition. Type **/boot** into the Mount Point text field, select ext3 as the File System Type, and specify a Size of 1000 MB. Select the Fixed size radio button and check the Force to be a primary partition checkbox. The Add Partition dialog should appear as below, except perhaps for the notation of the disk in the Allowable Drives section. Click the OK button.

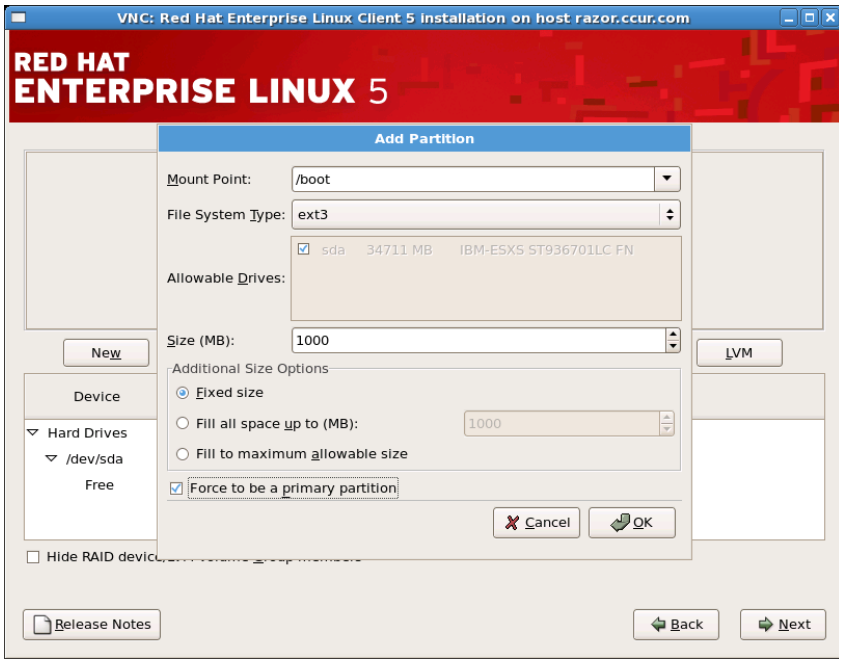

c. Click the New button to add the swap partition. Select swap as the File System Type. Select the Fixed size radio button and check the Force to be a primary partition checkbox.

Select an appropriate size for the swap partition -- you should specify a size which matches the amount of physical memory you have in your system, plus 5% (refer to "File System Swap Size Guidelines" on page 57).

The Add Partition dialog should appear as below, except perhaps for the notation of the disk in the Allowable Drives section and the specific size you have selected for the partition.

Click the OK button.

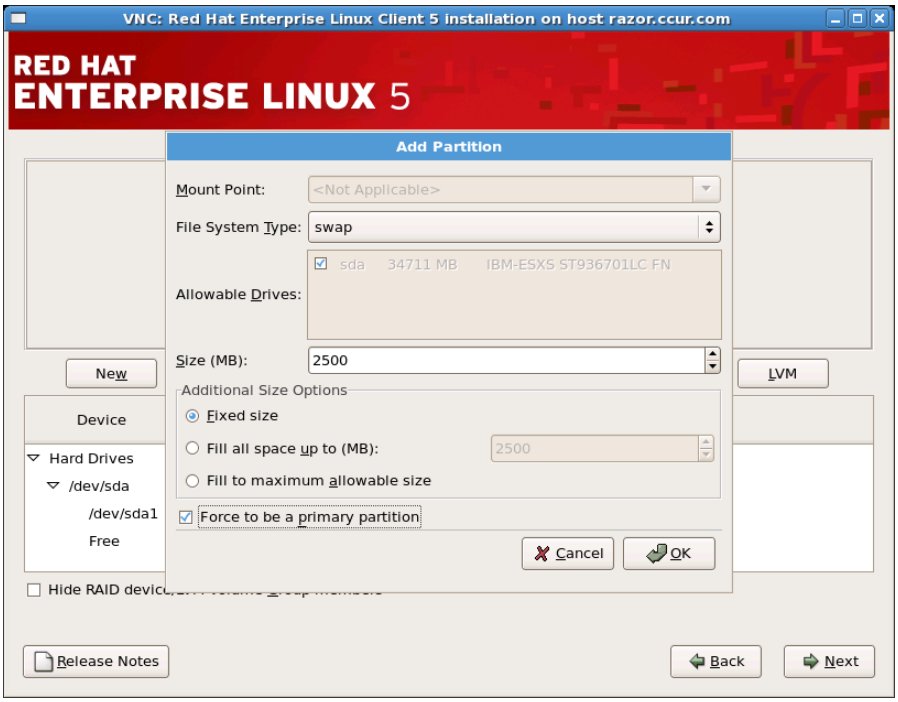

d. Click the New button to add the root partition. Type / into the Mount Point text field, select ext3 as the File System Type. Select the Fixed size radio button and check the Force to be a primary partition checkbox.

Select an appropriate size for the root partition. For large disks, a size of 60,000MB is recommended. The remaining space will be used for a partition of general use in the next step. It is recommended to separate the root partition from the remainder of the disk to make regenerating systems simpler, avoiding the need to back up the lion's share of the disk. For smaller disks, you may want to use all of the remainder of the disk as the root partition; if so, click the Fill to maximum allowable size radio button.

In all cases, the root partition should be no less that 12,000MB. The Add Partition dialog should appear as below, except perhaps for the notation of the disk in the Allowable Drives section and the specific size you have selected for the partition. Click the OK button.

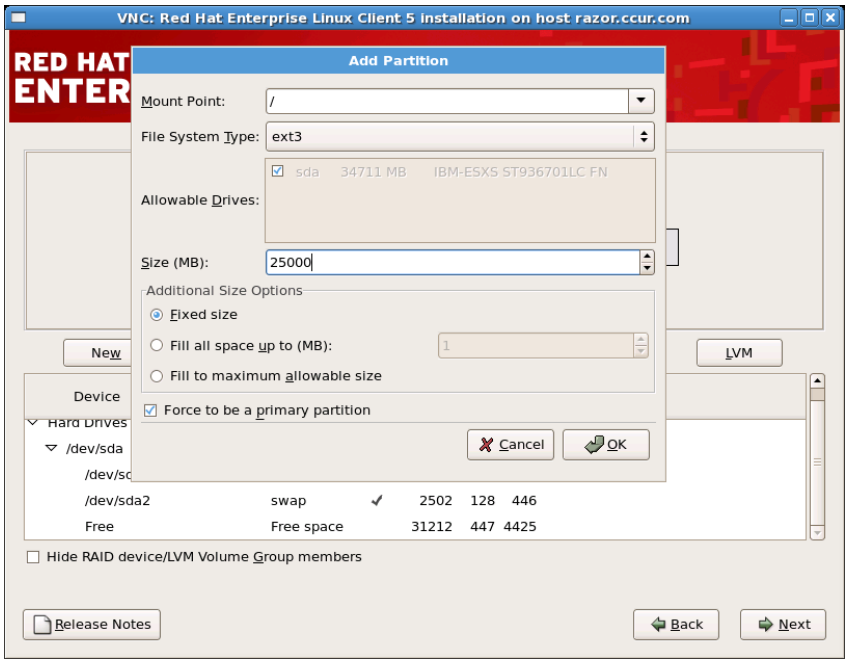

e. Unless you used the entire contents of the disk in the previous step, configure a **`/home`** partition for general use which uses the remaining free space on the disk.

f. The disk partitioning dialog should look similar to the following, except for device names and sizes as selected in the previous steps.

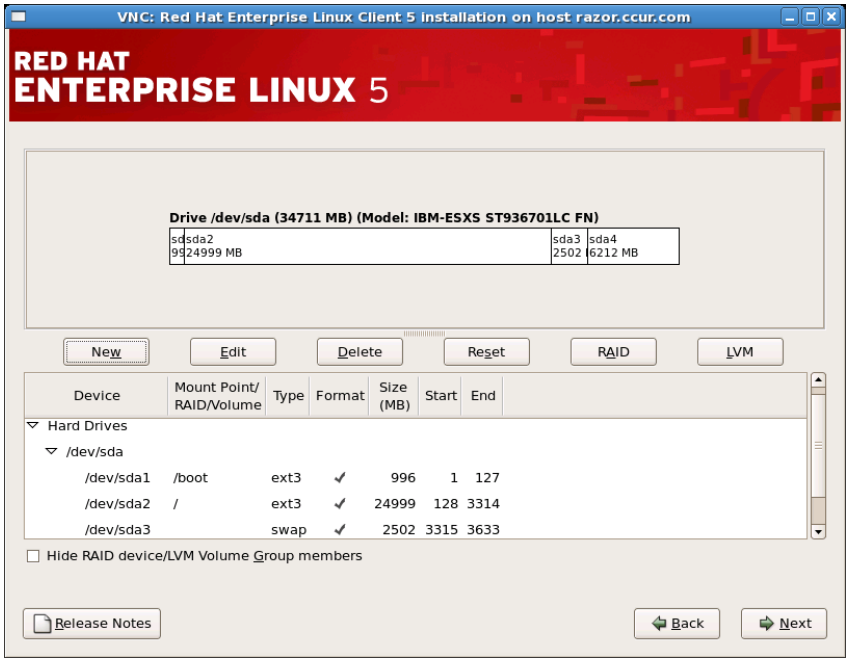

Click Next to continue.

11. The next dialog determines where the system boot information will be stored. Select "The GRUB boot loader will be installed on" radio button and click Next.

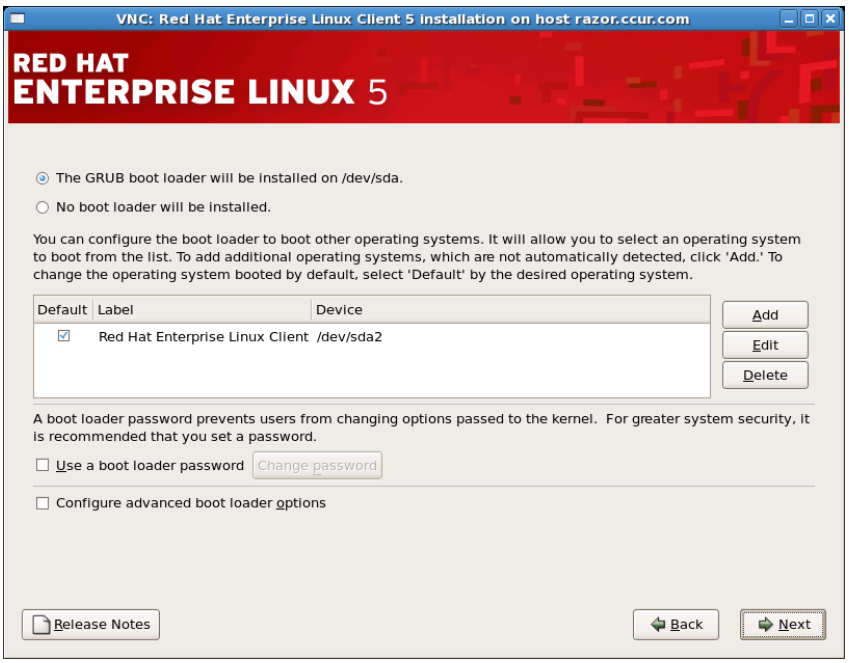

12. The next dialog defines networking configuration options. Select a network device from the list and then select either the DHCP or Manual radio button. Then click Edit.

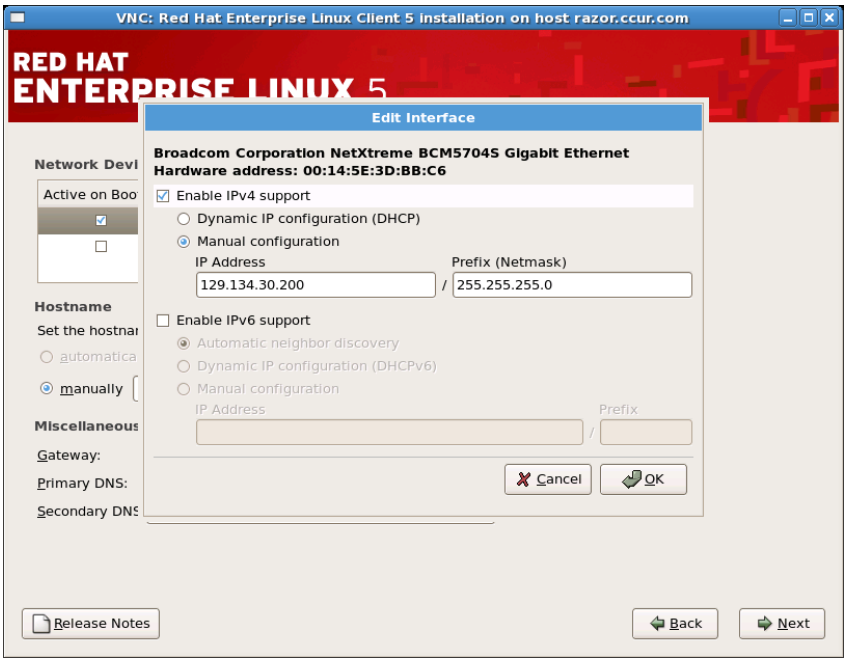

a. Ensure that the Enable IPv4 support checkbox is checked and the Enable IPv6 support checkbox is not checked.

b. For Manual configuration, enter your IP Address and Netmask in the text fields.

c. Click OK to close the Edit dialog.

d. The Networking screen now shows your configuration. If you choose Manual configuration, you also need to fill in the host name next to the manually text field, and the Gateway, Primary DNS and Secondary DNS fields in the Miscellaneous Settings section of the Networking dialog. Make any final required changes.

e. Follow the same procedure for any other network devices that are shown in the list and which you wish to be enabled.

f. Click Next to continue.

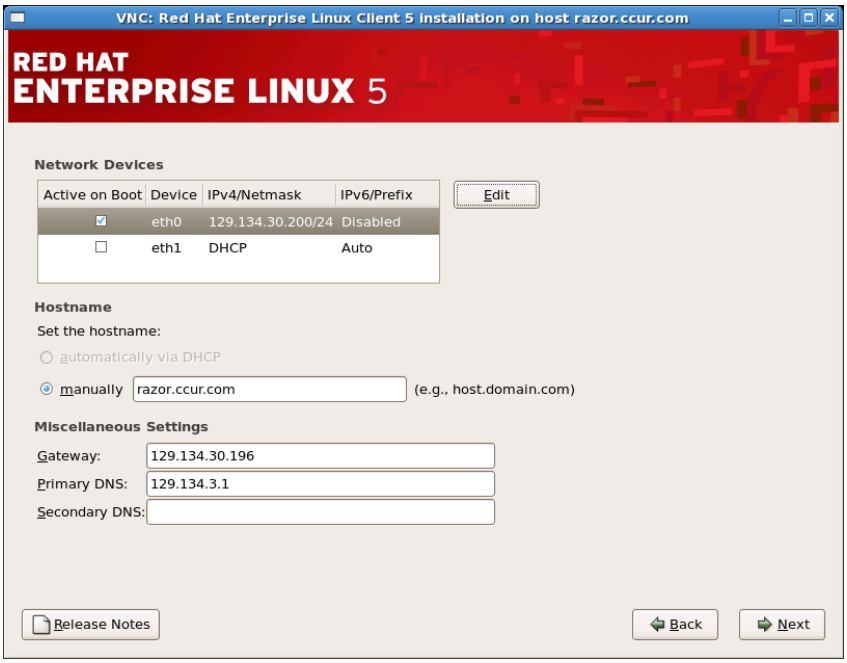

13. Select the time zone and enter root's password on the next two screens.

14. The Select Packages screen now appears. Ensure that Office, Multimedia and Software Development boxes are checked. Select the Customize Now radio button and click Next.

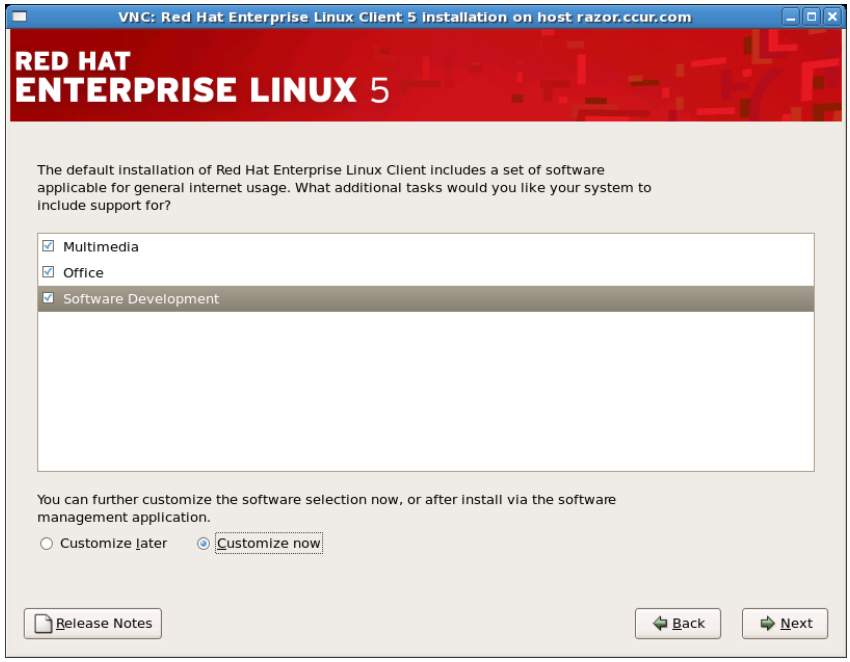

15. The following screen will appear which lists the major categories of software on the left and the major sub-sections within each category on the right. For Desktop Environments, check both GNOME and KDE selections.

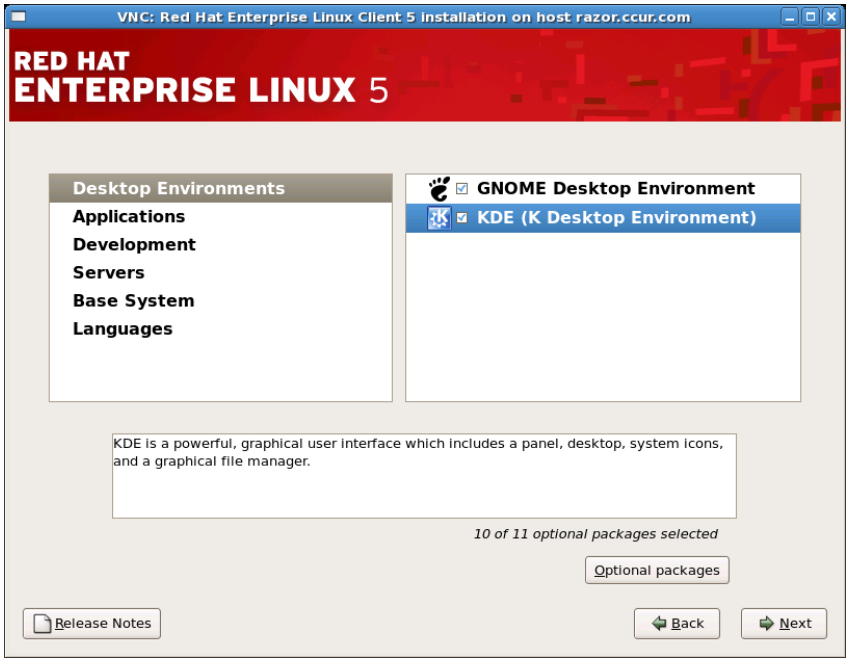

16. In addition to selecting the major sub-sections in each category, you can add additional optional packages to the installation. Select every package available in **every** sub-section.

NOTE

A quick way to select all sub-sections and all optional packages for all sub-sections within a category is to use multiple selection and the right-click menu.

Select the first sub-section in a category by clicking on it. Scroll to the bottom of the list of sub-sections if necessary, and then click on the last sub-section while holding the Shift key down. Now all sub-sections are selected. Right-click and choose Select all optional packages from the context menu.

To verify that you didn't miss any packages, look at the optional package statistic on the window which is immediately above the Optional packages button. It displays the number of packages selected out of the total number possible for the selected category.

17. Repeat this procedure for all software categories, selecting all major sub-sections, and all the optional products **UNTIL** you get to the Language category. That category is discussed in the next step.

NOTE

The Base System category has an optional package in the Base sub-section which cannot be installed. Even if you select it for installation by checking its checkbox, it will automaticaly uncheck itself when you close the Optional Packages dialog.

18. The Language category contains a wide variety of language support sub-sections.

These take up vast amounts of disk space. Only include the language sub-sections that you expect to use and only check the sub-section without adding additional optional packages.

If you do require additional optional packages, you may attempt to install them manually after all Red Hat updates have been installed, but it is likely that the Red Hat Installation media does not have the required (updated) versions. Contact Concurrent Software Support, as described in "Direct Software Support" on page 63, for aid in obtaining and installing such packages.

NOTE

Support for the Zulu language is not included on the RHEL 5.1 Updates CD. Do not select that language if you are installing from CD as the update step will subsequently fail.

19. After completing software customization, press the Next button. This causes the installer to check for package dependencies.

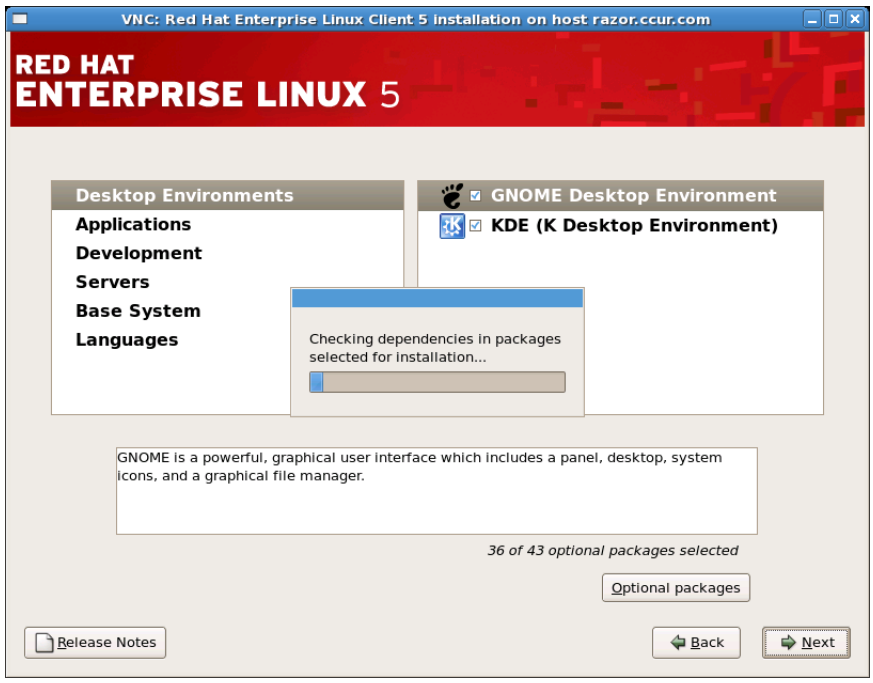

20. After dependency checking is complete, the following screen is shown. Click Next to continue.

21. After package dependencies have been completed, installation begins. No action is required from you during this stage unless you are installing from

CDs. In that case, you will occasionally be prompted to insert specific CDs from the installation set.

22. When installation is complete, the following dialog appears. Remove the media from the CD or DVD device and click Reboot to continue.

The system will now restart and boot Red Hat from your disk. Several more steps are required before we are finished with the Red Hat Installation.

23. After the initial boot screen completes, the following screen appears. Click Forward to proceed.

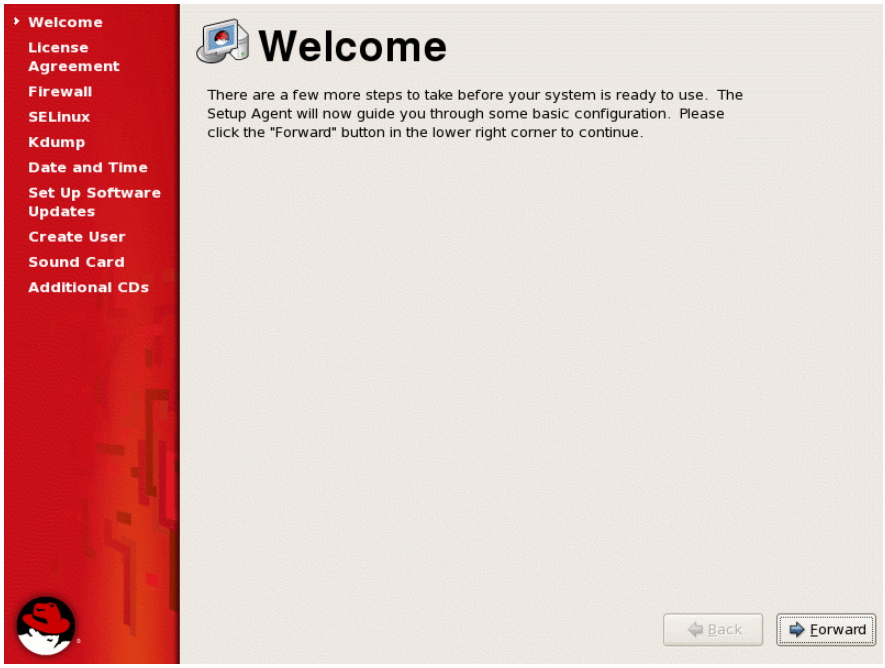

24. The License Agreement screen appears. Click Yes, I agree to the License Agreement, and then click Forward.

25. The next screen controls the firewall configuration. Select Disabled from the option list to disable the firewall. If your system will be on a public network that is not behind a firewall, you may wish to enable it. Click Forward to continue.

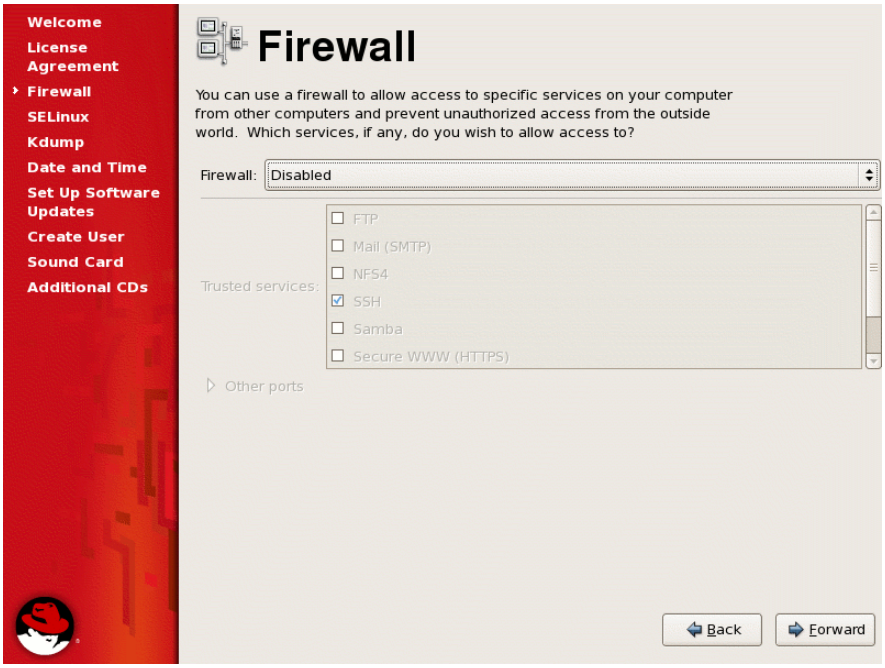

26. If you disabled the firewall in the previous step, the following warning dialog appears. Click Yes.

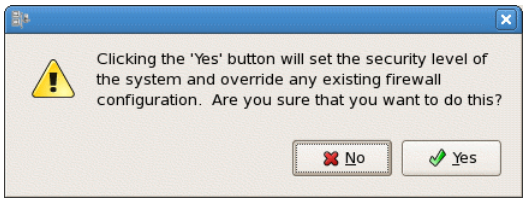

27. The next screen controls the settings of SELinux. Select Disabled from the option list and click Forward.

28. The following warning dialog will appear. Click Yes.

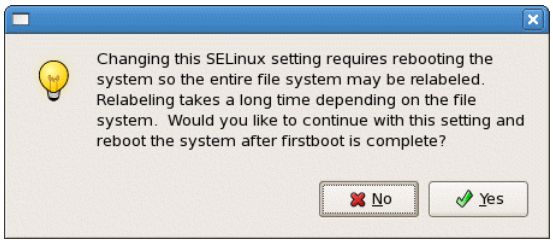

29. The next screen controls the configuration of Kdump. Ensure that the Enable Kdump checkbox is not set. The installation of RedHawk in a subsequent step installs an improved system dump utility which will be automatically configured. Click Forward.

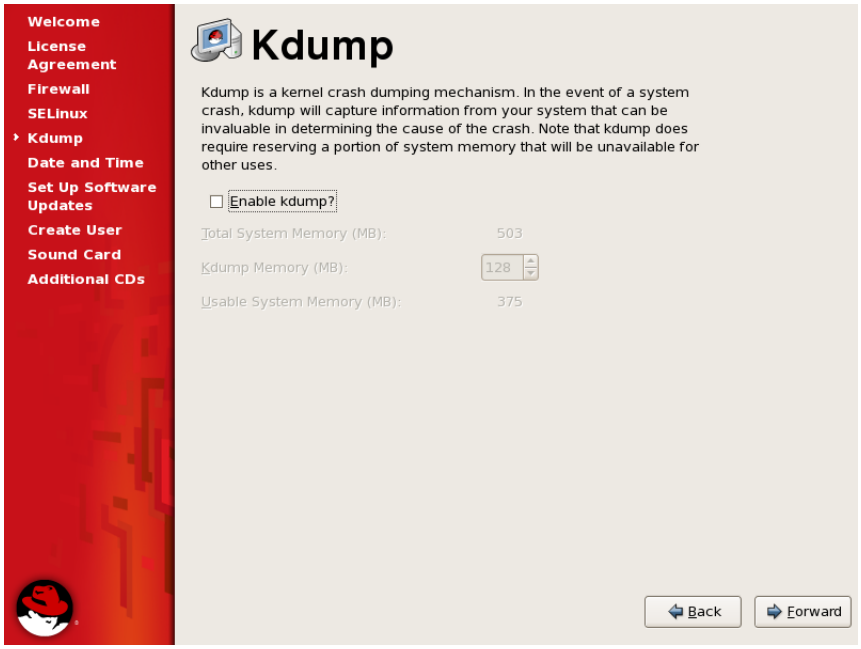

30. On the next screen, set the date and time and configure the Network Time Protocol daemon if desired.

31. The next screen asks you to register with Red Hat for software updates. Select the No, I prefer to register at a later time radio button. Software updates to the Red Hat Installation software will be installed in a subsequent step. Concurrent does not recommend installing continual updates from Red Hat because these may not have been validated for real-time operation or may include changes to **`glibc`** or **`gcc`** that are incompatible with other software installed with RedHawk.

    If you still wish to register with Red Hat you may do so after all the software installation steps are completed. Click Forward to continue.

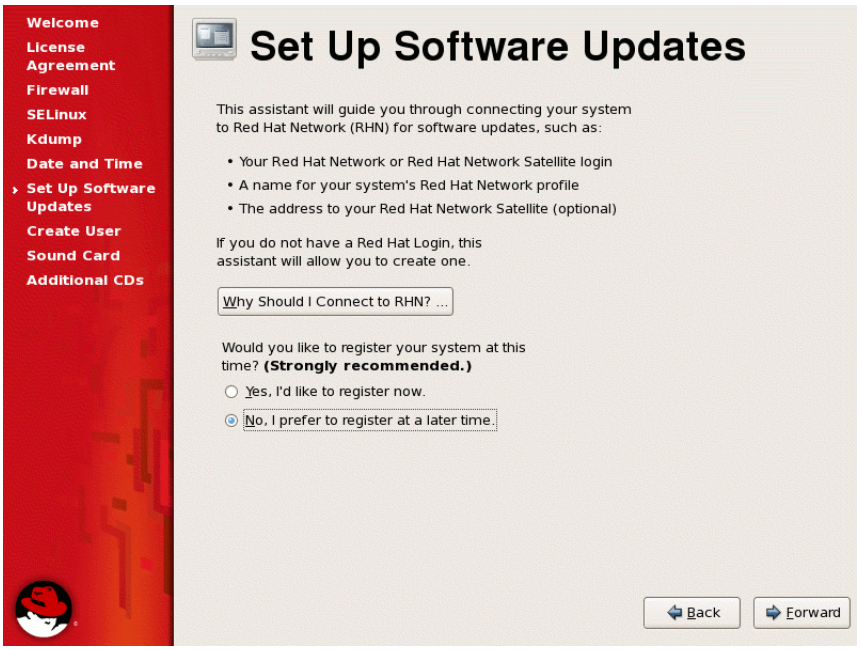

32. A warning dialog will appear. Click No thanks, I'll connect later, to continue.

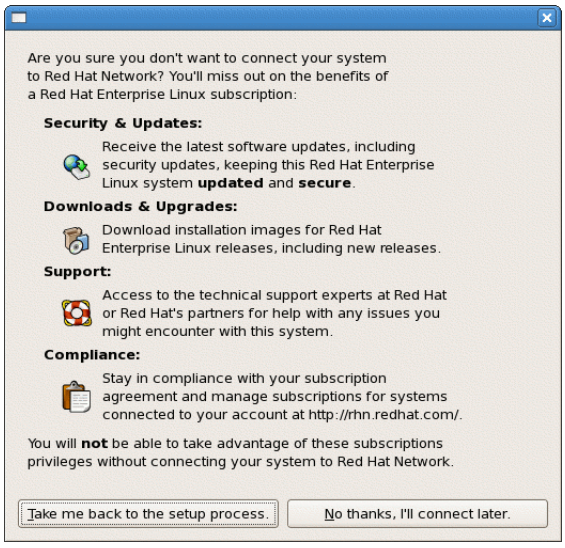

33. Another warning will appear on the Finish Updates Setup window. Click Forward to continue.

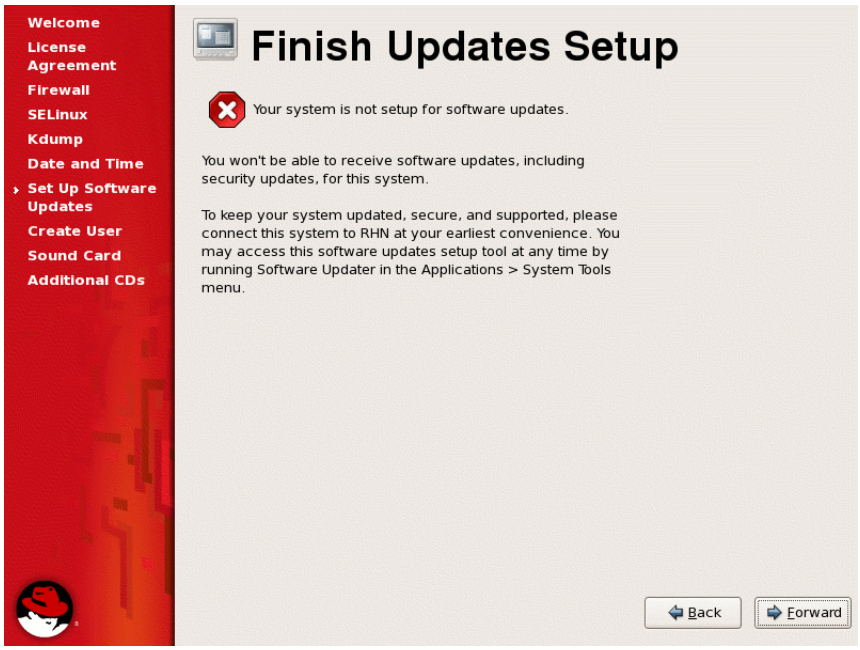

34. Optionally configure a user, NIS, and your sound card on the next two screens.

35. The final screen will prompt you for additional CDs. Do not click the Install… button -- the remaining software will be installed in a subsequent step. Click Finish, and then subsequently OK in the dialog which indicates a reboot is required.

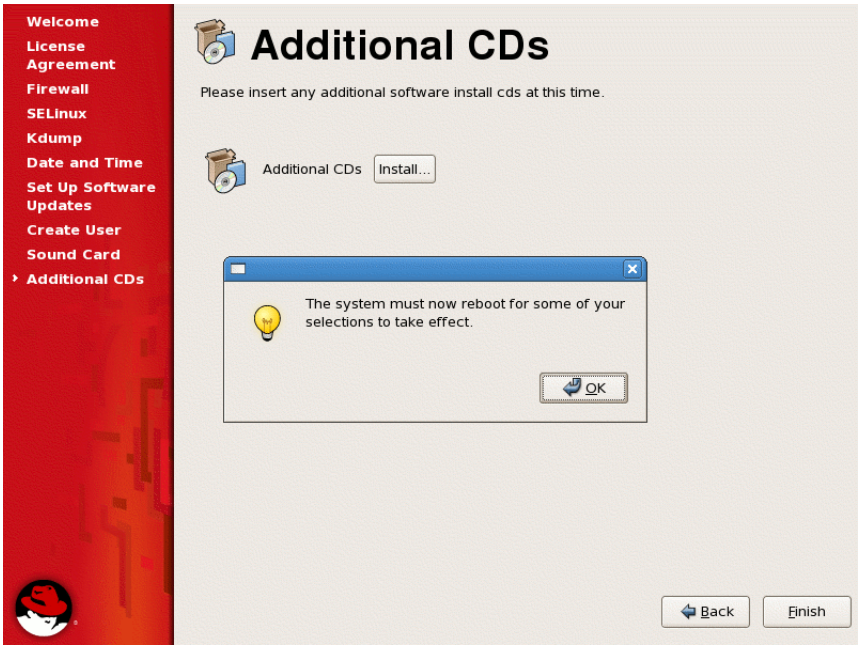

After the system reboots, preliminary installation of Red Hat is complete. Proceed to the next section to complete the Red Hat installation and update.

5.1.2 Installation of Red Hat Updates

Install the DVD labeled RHEL 5.1 Updates or the CD labeled RHEL 5.1 Updates 1 of 2. The disk contains the latest updates released by Red Hat which have been validated by Concurrent on RedHawk Linux. These updates are important to the proper operation of RedHawk Linux.

Execute the following instructions on the console or in a terminal window running as root.

1. Log in as root and take the system down to single-user mode:
   a. If in GUI mode, right click on the desktop and select Open Terminal.
   b. At the system prompt, type **`init 1`**.
2. If the CD/DVD device did not mount automatically, execute the **`mount`** command appropriate for mounting your CD or DVD device.

NOTES

NOTE 1: A specific mount point may not have been created under **`/mnt`** or **`/media`**.

If not, you can often identify the device file associated with your CD or DVD device by looking for files matching the following patterns:

```
ls /dev/cd* /dev/dvd*
```

If these files are not present, you should examine the boot messages in **/var/log/messages** and look for your CD or DVD device and find the associated **/dev** file.

Once located, you can mount the device with commands similar to the following:

[-d /media/cdrom] || mkdir -p /media/cdrom
mount *device-name* **/media/cdrom -o ro**

where *device-name* is the **/dev/** file you located.

NOTE 2: On systems with multiple optical media drives, you may need to manually create separate mount points for each drive under **/media** and include explicit mount instructions in **/etc/fstab** for each device. For example, on a system with a reader and a writer, you may need to manually create these mount points:

```
/media/cdrom
/media/cdrecorder
```

and include the following in **/etc/fstab**:

```
/dev/hda /media/cdrom auto pamconsole,exec,noauto,managed 0 0
/dev/hdc /media/cdrecorder auto pamconsole,exec,noauto,managed 0 0
```

3. Execute the following commands to install the update software, substituting the actual mount point used for your CD or DVD device for **/media/cdrom**, if necessary:

```
cd /media/cdrom
./install-updates
```

Package installation and update begins automatically. No action is required of you during this step.

NOTE

If you cannot execute the **install-updates** script (it fails with Permission Denied), or if you attempt to execute it by launching it from a file browser and it immediately returns, the CD or DVD may have been mounted with the **noexec** option.

Run the **mount** command (without arguments) and check the options. If **noexec** is present, remount the CD or DVD using:

```
mount -o remount,exec mount-point
```

where *mount-point* is the directory name in the output of the previous mount command for that CD/DVD. For example:

```
> mount | fgrep noexec
/dev/scd0 on /media/RHEL5.1-Updates type iso9660 (ro,noexec)

> mount -o remount,exec /media/RHEL5.1-Updates
```

The next stage involves reading in package header information from each package. No action is required of you during this step.

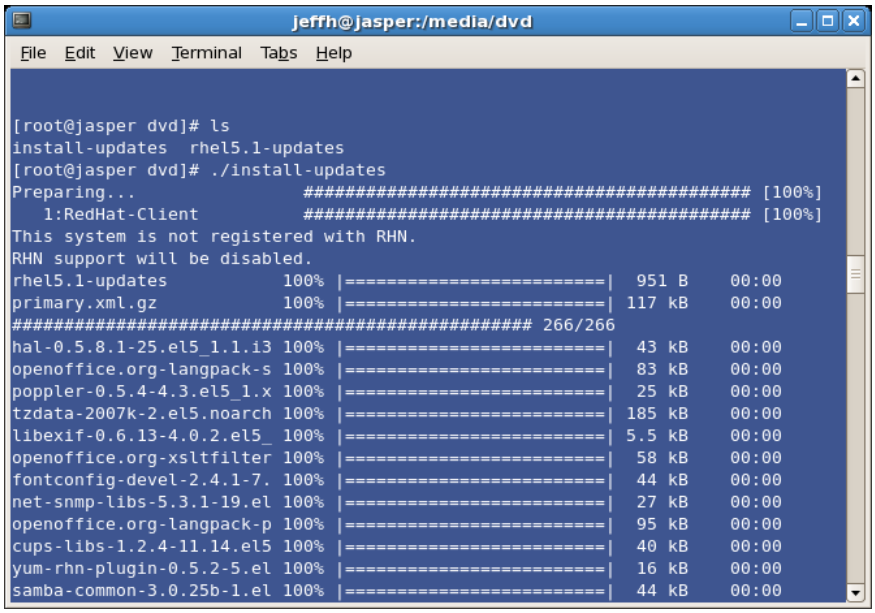

Finally, updating of packages is executed. No action is required of you during this step.

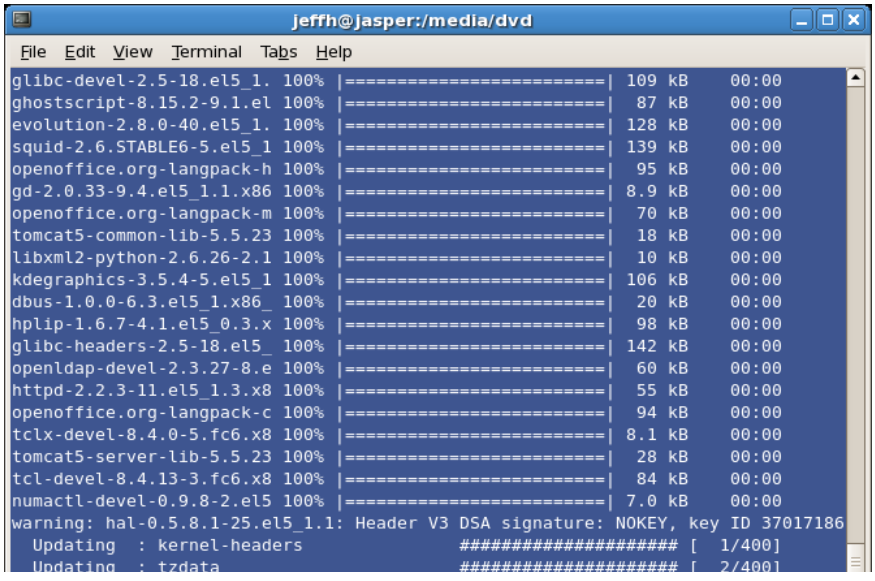

4. When the update completes, execute the following commands, substituting the actual mount point for your CD or DVD device for **/media/cdrom**:

   ```
   cd /
   umount /media/cdrom
   eject
   ```

5. Remove the disc from the CD/DVD drive.

6. If installing from CD, a second RHEL 5.1 Update disk is required. Insert the CD labeled RHEL 5.1 Updates 2 of 2 and repeat steps 3-5.

7. **Perform this step ONLY if you are installing on a single processor system.** If installing on a multiprocessor system or on a system with one physical CPU that supports hyper-threading, skip to the next step.

   On a single processor system, execute the following:

   ```
   rpm -e kernel-smp
   ```

NOTE: It is possible that some warning messages may display at this time. These can be ignored and will be resolved when the system is rebooted.

8. Type “reboot” at the system prompt and verify that the new Red Hat kernel boots.

9. Continue with the steps in the next section to install RedHawk Linux.

5.2 Installing RedHawk Linux

After completing the installation of Red Hat in the previous section, perform the following steps to install RedHawk Linux:

1. Log in as root and take the system down to single-user mode:

   a. If in GUI mode, right click on the desktop and select Open Terminal.
   b. At the system prompt, type **`init 1`**.

2. Edit the **`/etc/X11/xorg.conf`** file and ensure that Module Option lines of the following form are ***not*** part of the file:

   ```
   Section "Module"
           ...
           Option "dri"
           Load "dri"
           ...
   End Section
   ```

   The "dri" option has been known to cause serious system instability and must not be specified.

3. Locate the disc labeled "RedHawk Linux OS Version 5.1" appropriate to your system's architecture and insert it into the CD/DVD drive.

4. If the cdrom device does not mount automatically, execute the **`mount`** command appropriate for your CD or DVD device.

5. Execute the following commands to install RedHawk Linux, substituting the actual mount point used for your CD or DVD device for **`/media/cdrom`**, if necessary:

   ```
   cd /media/cdrom
   ./install-redhawk
   ```

   NOTES

   NOTE 1: Special instructions are displayed during the installation that may require appropriate action, such as video card installation and configuration. Follow the on-screen instructions until the installation script completes.

   NOTE 2: The installation script will prompt you to accept or reject the End User License Agreement. You can view the agreement by entering "v" at the prompt.

   NOTE 3: Note that the **`/etc/X11/XF86Config`** has been renamed to **`/etc/X11/xorg.conf`**. Any references to the **`XF86Config`** file in message output refers to **`xorg.conf`**.

6. When the installation completes, execute the following commands:

   ```
   cd /
   umount /media/cdrom
   eject
   ```

7. Remove the disc from the CD/DVD drive.

8. Exit single-user mode (Ctrl-D).
9. If you installed in text mode because full support for the video card was not present during Red Hat installation, you need to complete this configuration now. Type `init 5` at the system prompt. Proceed through the windows that display to configure screen resolution and color depth for your system. When you exit this session, the RedHawk Linux desktop should appear.
10. Reboot the system (see **NOTES** below) choosing the new RedHawk kernel and verify that the system boots properly.

NOTES

If this system is a model 868 8-way iHawk, disable hyper-threading through the BIOS (refer to the hardware documentation for the appropriate BIOS setting).

On an i386 system with greater than 4 GB of memory installed, support must be configured into the kernel. Refer to section **7.1.3** on page 57 for details.

11. Continue with the next section if an RCIM is to be installed. Otherwise, proceed to the following section to install Frequency-Based Scheduler (FBS).

5.3 Installing the RCIM

The following serves as an installation checklist for installing a PCI-based RCIM board in an iHawk system. If an RCIM is not installed in your iHawk system, it should be installed now. See the *Real-Time Clock and Interrupt Module (RCIM) User's Guide* for complete details. The PDF file for this manual is available for viewing by clicking on the "Documents" icon on the desktop.

5.3.1 Quick Hardware Installation Checklist

1. Before installing the RCIM, determine if you will be using the RCIM to accept or deliver external interrupts and the mode in which the RCIM will run (see "RCIM Connection Modes" on page 57).

NOTE

The initial release of the RCIM III board does not support cabling (chaining) boards together or distributing RTC or PIG operations to local Distributed Interrupt devices.

2. Verify that the **`ccur-rcim`** RPM has been installed by executing the command **`rpm -q ccur-rcim`**. The system will inform you if it is not installed. It is a standard package in the RedHawk installation.
3. Power down the system and remove all power cords.

NOTE

Concurrent Computer Corporation strongly recommends that you use an antistatic wrist strap and a conductive foam pad when installing and removing printed circuit boards.

4. Open the case of your system and identify the PCIe slot (RCIM III) or PCI slot (RCIM II or RCIM I) where you want the RCIM to reside. In general, it is best for the RCIM to be configured in a slot where minimal or no contention with other devices occurs and at the highest IRQ priority possible. For more information, refer to the *iHawk Optimization Guide*, publication number 0898011. The PDF file for this manual is available for viewing by clicking on the "Documents" icon on the desktop.
5. Insert the RCIM into the selected PCI slot and lock it down.
6. If this is to be part of an RCIM chain, attach the synchronization cable as required (see the *Real-Time Clock and Interrupt Module (RCIM) User's Guide* manual for complete details).
7. If you have the optional GPS module, attach the antenna lead and mount the antenna. The antenna should be mounted on the rooftop or in an open area.
8. Close the case and reconnect all power cords.
9. Power up the system and verify the system boots up properly. Refer to the RCIM guide for configuration options.

5.3.2 Verifying RCIM Operation

1. To check that the RCIM is operating correctly, perform the following:

   `cat /proc/driver/rcim/status`

   You should see an output similar to that shown below:

   ```
   RCIM-I board 0 is at revision 1 eeprom 1.0.
   This is a standalone (isolated) rcim.
   Has IRQ 17 and major number 252
   System time synchronized to RCIM.
   ```

   NOTE

   The RCIM board may be either RCIM-I, RCIM-II or RCIM-III and the revision, eeprom, and IRQ numbers may differ.

2. To view the current RCIM configuration, issue the following command:

   `cat /proc/driver/rcim/config`

   You should see output similar to the following:

```
h/Not_Configured,sync,pig0|out0,pig1|out1,pig2|
out2,pig3|out3,none|di0,none|di1,none|di2,none|
di3,none|di4,none|di5,none|di6,none|di7,di0/f,
i1/f,di2/f,di3/f,di4/f,di5/f,di6/f,di7/f,eti0/f,
eti1/f,eti2/f,eti3/f
```

5.3.3 Installing NTP Updates for GPS Module

If you installed an RCIM model equipped with the optional GPS module, the **ccur-ntp** RPM is required for its use. If you do not have an RCIM board with GPS module, skip to section **5.4**.

1. The ccur-ntp RPM is available via Concurrent's Software Repository system. Execute the following steps to install this RPM:

   a. Remove (erase) the existing ntp. Because there are some rpms which are normally installed which depend on ntp, use the **`--nodeps`** option to bypass the dependency test. This will save the existing **`/etc/ntp.conf`** file as **`/etc/ntp.conf.rpmsave`**.

      # **`rpm --nodeps -e ntp`**

   b. Download the ccur-ntp RPM from Concurrent's web site as described in section **9.2.2** on page 63.

   c. Install the RPM via the following command:

      # **`rpm -ivh ccur-ntp*.rpm`**

2. After your system is completely installed, refer to the *Real-Time Clock and Interrupt Module (RCIM) User's Guide* for instructions for configuring ntp.

5.4 Installing Frequency-Based Scheduler (FBS) Software

FBS is an optional package to RedHawk Linux. If you will be using FBS, install it at this time by performing the following steps:

1. With RedHawk Linux Version 5.1 running, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk Linux Frequency-Based Scheduler Version 5.1" appropriate to your system's architecture and insert it into the CD-ROM drive.
3. If the cdrom device does not mount automatically, execute the **`mount`** command; for example:

   `mount /media/cdrom`

4. To install, execute the following commands:

   ```
   cd /media/cdrom
   ./install-fbs
   ```

   Follow the on-screen instructions until the installation script completes.

5. When the installation completes, execute the following commands:

   ```
   cd /
   umount /media/cdrom
   eject
   ```

6. Remove the disc from the CD-ROM drive and store.
7. Exit single-user mode (Ctrl-D).

5.5 Installing the PCI-to-VME Bridge

If you are using the PCI-to-VME bridge supplied by Concurrent to connect your iHawk system with a VME system, it can be installed at this time. If you are not installing this feature, skip to to the next section.

5.5.1 Installing the Hardware

The hardware consists of a PCI adapter, a VMEbus adapter and a fiber optic cable.

NOTE: Concurrent Computer Corporation strongly recommends that you use an antistatic wrist strap and a conductive foam pad when installing and removing printed circuit boards.

1. To configure the jumpers on the VMEbus adapter card correctly, refer to Chapter 10 of the *SBS Technologies Model 618-3 Hardware Manual* included with the RedHawk Linux documentation. The PDF file for this manual is available for viewing by clicking on the "Documents" icon on the desktop.
2. To install the PCI adapter in your iHawk system:
   a. Power down the system.
   b. Locate a vacant PCI card slot in the chassis that supports a bus master.
   c. Remove the metal plate that covers the cable exit at the rear of the chassis.
   d. Insert the PCI adapter card into the connector.
   e. Fasten the adapter card in place with the mounting screw.
   f. Replace the cover.
3. To install the VMEbus adapter card:
   a. Ensure that the VMEbus chassis is powered down.
   b. Decide whether the VMEbus adapter card is the system controller. If so, it must be installed in slot 1. If not, locate an unoccupied 6U slot in the card cage for the adapter.
   c. Insert the card into the connector of the selected slot.
4. To install the fiber optic cable:
   a. Ensure that both systems are powered down.
   b. Remove the rubber boots on the fiber-optic transceivers as well as the ones on the fiber-optic cables. Be sure to replace these boots when cables are not in use.
   c. Plug one end of the fiber-optic cable into the PCI adapter card's transceiver.

d. Plug the other end of the cable into the VMEbus adapter card's transceiver.

e. Power on both systems.

f. Ensure that the READY LEDs on both adapter cards are lit.

5.5.2 Installing the Software

1. With RedHawk Linux Version 5.1 running on the iHawk system, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk Linux PCI-to-VME Bridge Library Version 5.1" appropriate to your system's architecture and insert it into the CD-ROM drive.
3. If the cdrom device does not mount automatically, execute the **mount** command; for example:

   ```
   mount /media/cdrom
   ```

4. To install, execute the following commands:

   ```
   cd /media/cdrom
   ./install-sbsvme
   ```

   Follow the on-screen instructions until the installation script completes.

5. When the installation completes, execute the following commands:

   ```
   cd /
   umount /media/cdrom
   eject
   ```

6. Remove the disc from the CD-ROM drive and store.
7. Exit single-user mode (Ctrl-D).

Refer to the "PCI-to-VME Support" chapter of the *RedHawk Linux User's Guide* for configuration and usage information. The PDF file for this manual is available for viewing by clicking on the "Documents" icon on the desktop.

5.6 Installing Additional RedHawk Products

If you will be installing any additional RedHawk products, do so now. Refer to the products' documentation for installation instructions.

5.7 Installing RedHawk Updates from the Updates Web Site

1. Additional updates may have been issued. They should be downloaded and installed now. Refer to "Software Updates" on page 63 for details.
2. When complete, or if no updates are available, reboot the system choosing the new RedHawk kernel and verify that the system boots properly.

6.0. NVIDIA Graphics Configuration

The Unified NVIDIA® Linux Display Driver (version 169.09) is included in the RedHawk Linux installation and built automatically for all pre-built and custom RedHawk Linux kernels. The RedHawk installer probes for the type and model of video adapters present in the system, and recommends installing the unified **ccur-nvidia-glx** rpm when it detects one or more NVIDIA video cards present.

Even if the Unified NVIDIA driver is installed, you still must modify **/etc/X11/xorg.conf** to actually use it. You must either change the "Driver" field in the "Device" section from "nv" to "nvidia", or copy one of the other **xorg.conf** example files, such as **xorg.conf.singledesktop**, into **xorg.conf**. Note that all of the Concurrent **xorg.conf** example files use "nvidia" for the "Driver" field.

Note that if you install an NVIDIA graphics card into your system after RedHawk Linux has been installed, you may need to install the unified NVIDIA user module, **ccur-nvidia-glx**, at that time. The rpm for this module can be found on the RedHawk Linux installation CD in the **nvidia** directory.

If you have a system with multiple video displays and you wish to configure a multiple display mode setup, refer to section 6.4.

6.1 The Unified and Legacy NVIDIA Display Drivers

The NVIDIA Linux Display Driver (version 169.09) that is included in the RedHawk Linux installation is referred to as the "unified" display driver. The unified driver supports the latest and more recent NVIDIA GeForce and Quadro based video cards.

In addition to the unified driver, an alternative "legacy" NVIDIA Linux Display Driver (version 96.43.05) is included that supports older NVIDIA GeForce and Quadro based Graphics Processor Units (GPUs).

Some NVIDIA GPUs are only supported in either the unified or the legacy driver, while other GPUs are supported in both drivers. **It is recommended that you use the unified NVIDIA Linux Display Driver whenever possible.**

NOTE

If you have an older NVIDIA GPU on your system that is only supported in the legacy driver, you ***must*** follow the steps described in section **6.3**, "Using the Legacy NVIDIA Driver", to manually remove the unified driver and then manually install the legacy driver. Attempting to use the unified driver on a system with a legacy NVIDIA video card may cause the system and/or X session to lock up.

Customers with recent NVIDIA graphics cards should be able to use the unified driver provided by default in RedHawk Linux. In this case, you will disregard section 6.3, "Using the Legacy NVIDIA Driver".

6.1.1 Supported NVIDIA GPUs

Table 6-1 lists the NVIDIA GPUs that are supported only in the unified (version 169.09) NVIDIA Linux Display Driver.

Table 6-2 lists the NVIDIA GPUs that are supported only in the legacy (version 96.43.05) NVIDIA Linux Display Driver.

GPUs that are not listed in these two tables are generally supported in both the unified and legacy drivers.

Older RIVA and RIVA TNT2 types of NVIDIA GPUs are not supported by either the unified or legacy drivers.

To view the list of all NVIDIA GPUs supported by each driver, see Appendix A in the following README files:

- If the unified driver (version 169.09) is installed on your system:

  `/usr/share/doc/ccur-nvidia-glx-169/README.txt`

- If the legacy driver (version 96.43.05) is installed on your system using the steps described in section **6.3**:

  `/usr/share/doc/NVIDIA_GLX-1.0/README.txt`

Table 6-1 GPUs Supported Only in the NVIDIA 'unified' Driver

NVIDIA Chip Name	Device PCI ID	NVIDIA Chip Name	Device PCI ID
GeForce 6800 LE	0x0042	GeForce 6150SE nForce 430	0x03D0
GeForce 7800 SLI	0x0095	GeForce 6100 nForce 405	0x03D1
GeForce 6700 XL	0x0147	GeForce 6100 nForce 400	0x03D2
Quadro FX 540M	0x014C	GeForce 6100 nForce 420	0x03D5
Quadro FX 550	0x014D	GeForce 8600 GTS	0x0400
GeForce 6200SE TurboCache™	0x0162	GeForce 8600 GT	0x0402
GeForce 6200 LE	0x0163	GeForce 8400 GS	0x0404
GeForce 6250	0x0169	GeForce 8600M GT	0x0407
GeForce 7100 GS	0x016A	GeForce 8700M GT	0x0409
GeForce 8800 GTX	0x0191	Quadro FX 370	0x040A
GeForce 8800 GTS	0x0193	Quadro NVS 320M	0x040B
GeForce 8800 Ultra	0x0194	Quadro FX 570M	0x040C
Tesla C870	0x0197	Quadro FX 1600M	0x040D
Quadro FX 5600	0x019D	Quadro FX 570	0x040E
Quadro FX 4600	0x019E	Quadro FX 1700	0x040F
GeForce 7350 LE	0x01D0	GeForce 8400 SE	0x0420
GeForce 7300 SE/7200 GS	0x01D3	GeForce 8500 GT	0x0421
GeForce Go 7200	0x01D6	GeForce 8400 GS	0x0422
GeForce 7500 LE	0x01DD	GeForce 8300 GS	0x0423
GeForce 6200	0x0221	GeForce 8600M GS	0x0425
GeForce 6200 A-LE	0x0222	GeForce 8400M GT	0x0426
GeForce Go 6150	0x0244	GeForce 8400M GS	0x0427
Quadro NVS 210S/NVIDIA GeForce 6150LE	0x0245	GeForce 8400M G	0x0428
GeForce Go 6100	0x0247	Quadro NVS 140M	0x0429
GeForce 7900 GT/GTO	0x0291	Quadro NVS 130M	0x042A
GeForce 7900 GS	0x0292	Quadro NVS 135M	0x042B
GeForce 7950 GX2	0x0293	Quadro FX 360M	0x042D
GeForce 7950 GX2	0x0294	Quadro NVS 290	0x042F
GeForce 7950 GT	0x0295	GeForce 7150M/nForce 630M	0x0531
GeForce Go 7950 GTX	0x0297	GeForce 7000M/nForce 610M	0x0533
GeForce 7600 GT	0x02E0	GeForce 7050 PV/NVIDIA nForce 630a	0x053A
GeForce 7900 GS	0x02E3	GeForce 7050 PV/NVIDIA nForce 630a	0x053B
GeForce 7950 GT	0x02E4	GeForce 7025 /NVIDIA nForce 630a	0x053E
Quadro FX 500/FX 600	0x032B	GeForce 8800 GTS 512	0x0600
GeForce 7650 GS	0x0390	GeForce 8800M GTS	0x0609
GeForce 7300 GT	0x0393	GeForce 8800M GTX	0x060C
GeForce 7600 LE	0x0394	GeForce 8800 GT	0x0611
GeForce 7300 GT	0x0395	Quadro FX 3700	0x061A
GeForce Go 7600 GT	0x0399		

Table 6-2 GPUs Supported Only in the NVIDIA 'legacy' Driver

NVIDIA Chip Name	Device PCI ID	NVIDIA Chip Name	Device PCI ID
GeForce PCX 4300	0x00FF	Quadro NVS with AGP8X	0x018A
GeForce2 MX/MX 400	0x0110	Quadro4 380 XGL	0x018B
GeForce2 MX 100/200	0x0111	Quadro NVS 50 PCI	0x018C
GeForce2 Go	0x0112	GeForce2 Integrated GPU	0x01A0
Quadro2 MXR/EX/Go	0x0113	GeForce4 MX Integrated GPU	0x01F0
Quadro FX 550	0x014C	GeForce3	0x0200
GeForce4 MX 460	0x0170	GeForce3 Ti 200	0x0201
GeForce4 MX 440	0x0171	GeForce3 Ti 500	0x0202
GeForce4 MX 420	0x0172	Quadro DCC	0x0203
GeForce4 MX 440-SE	0x0173	GeForce4 Ti 4600	0x0250
GeForce4 440 Go	0x0174	GeForce4 Ti 4400	0x0251
GeForce4 420 Go	0x0175	GeForce4 Ti 4200	0x0253
GeForce4 420 Go 32M	0x0176	Quadro4 900 XGL	0x0258
GeForce4 460 Go	0x0177	Quadro4 750 XGL	0x0259
Quadro4 550 XGL	0x0178	Quadro4 700 XGL	0x025B
GeForce4 440 Go 64M	0x0179	GeForce4 Ti 4800	0x0280
Quadro NVS	0x017A	GeForce4 Ti 4200 with AGP8X	0x0281
Quadro4 500 GoGL	0x017C	GeForce4 Ti 4800 SE	0x0282
GeForce4 410 Go 16M	0x017D	GeForce4 4200 Go	0x0286
GeForce4 MX 440 with AGP8X	0x0181	Quadro4 980 XGL	0x0288
GeForce4 MX 440SE with AGP8X	0x0182	Quadro4 780 XGL	0x0289
GeForce4 MX 420 with AGP8X	0x0183	Quadro4 700 GoGL	0x028C
GeForce4 MX 4000	0x0185	GeForce 7900 GT	0x0291
Quadro4 580 SGL	0x0188	Quadro FX 500/600 PCI	0x032B

6.2.1 Multiple NVIDIA Video Card Considerations

Note that you may **not** use the unified ***and*** legacy NVIDIA Linux Display Drivers together at the same time.

If you have a system with a mix of legacy-only and unified-only NVIDIA video cards, you will only be able to use one set of cards (legacy or unified) at any single point in time.

It is hightly recommended that you physically remove the legacy or unified video cards that you are not using so that they do not interfere with the card(s) that you are planning to use. Failure to do so may lock up your system or X11 session.

If you have a mixed video card system as described here, it is recommended that you upgrade your legacy NVIDIA video cards to unified video cards and use the unified driver as your display driver.

6.3 Using the Legacy NVIDIA Driver

If you have an NVIDIA graphics card with a GPU that is only supported on the legacy NVIDIA Linux Display Driver, use the steps in the following subsections to install and use the legacy driver.

This section also discusses how to remove the legacy driver and reinstall the unified driver.

6.3.1 Installing the NVIDIA Legacy Driver

1. It is recommended that you first uninstall the unified driver before installing the legacy driver.

   a. Exit any currently active X session.

   b. Login as root.

   c. Exit X11 state and enter full multiuser mode:

      `init 3`

   d. Issue the **`rpm -q ccur-nvidia-glx`** command to see if the unified driver RPM is currently installed on your system. If not, then skip ahead to step 2.

   e. Issue the following command to unload the nvidia kernel module:

      `modprobe -r nvidia`

   f. Issue the following command to remove the unified NVIDIA rpm:

      `rpm -e ccur-nvidia-glx`

2. Install the legacy driver. This step installs the user-space NVIDIA driver support and also builds and installs an nvidia kernel module into the currently executing kernel.

   This step assumes that you are still running as the root user in init state 3.

   a. You must first set up the kernel source configuration build environment so that it matches the currently executing pre-built kernel.

      If you are running a pre-built RedHawk kernel, issue the following commands:

      `cd /usr/src/`*redhawk-release-directory*
      `ccur-config -n -c`

      If you are running a custom built RedHawk kernel and the kernel source tree configuration does not match your kernel, use the **`ccur-config`** script to set up a matching kernel configuration. For example, if you created your original custom kernel configuration with the following commands:

      `cd /usr/src/`*redhawk-release-directory*
      `ccur-config -s -k custom trace-static`

      then use the following command to recreate a kernel configuration that matches your custom kernel:

      `cd /usr/src/`*redhawk-release-directory*
      `ccur-config -n custom`

See the **ccur-config(8)** man page and the "Configuring and Building the Kernel" chapter of the *RedHawk Linux User's Guide* for more information.

b. Insert the RedHawk Linux installation CD in the CD-ROM drive.

c. Move to the cdrom mount point/directory and issue a **cd nvidia** command.

d. Enter the following command:

```
sh NVIDIA-Linux-x86*-96.43.05-RedHawk-legacy-pkg*.run --ui=none -a -N -q
```

where the above options are:

`-ui=none`	use a simple user-interface
`-a`	accept the NVIDIA customer license
`-N`	do not attempt to connect to the NVIDIA web site
`-q`	do not ask any questions; use the default answers

At this point the legacy NVIDIA driver is installed and ready for use on this kernel. You may either start up an X session or return to 'init 5' (X11) system state.

NOTE

If you plan to use the legacy NVIDIA driver on additional pre-built or custom kernels, or if you plan to rebuild and reinstall kernel modules for the currently executing kernel, you may want to copy the **NVIDIA-Linux-x86*-96.43.05-RedHawk-legacy-pkg*.run** tarball to your local filesystem where it can be used for building future nvidia kernel modules without having to re-insert the RedHawk installation CD.

6.3.2 Building Additional Kernel Modules

If you boot with a kernel other than the one running when you installed the legacy driver in section **6.3.1**, you will need to build and install a new compatible kernel nvidia module in the currently executing kernel.

To build and install a legacy nvidia kernel module:

1. Boot the kernel and login as root.
2. Exit X11 state and enter full multiuser mode:

   ```
   init 3
   ```

3. Issue the following command to unload the nvidia kernel module (if loaded):

   ```
   modprobe -r nvidia
   ```

4. If you are using a RedHawk pre-built kernel, set up the proper build environment for the currently executing kernel:

   cd /usr/src/*redhawk-release-directory*
   ccur-config -n -c

   If you are running a custom built RedHawk kernel and the kernel source tree configuration does not match your kernel, use the **ccur-config** script to set up a matching kernel configuration. For example, if you created your original custom kernel configuration with the following commands:

   cd /usr/src/*redhawk-release-directory*
   ccur-config -s -k custom trace-static

   then use the following command to recreate a kernel configuration that matches your custom kernel:

   cd /usr/src/*redhawk-release-directory*
   ccur-config -n custom

   See the **ccur-config(8)** man page and the "Configuring and Building the Kernel" chapter of the *RedHawk Linux User's Guide* for more information.

5. See the **ccur-config(8)** man page and the "Configuring and Building the Kernel" chapter of the *RedHawk Linux User's Guide* for more information.

6. If you previously saved the NVIDIA***.run** tarball to your local filesystem, 'cd' to the directory where the tarball resides.

   Otherwise, insert the RedHawk Linux installation CD into the CD-ROM drive, move to the cdrom mount point/directory and issue a **cd nvidia** command.

7. Issue the following command to build a new legacy nvidia kernel module:

sh NVIDIA-Linux-x86*-96.43.05-RedHawk-legacy-pkg*.run --ui=none -a -N -q -K

where the above options are:

`-ui=none`	use a simple user-interface
`-a`	accept the NVIDIA customer license
`-N`	do not attempt to connect to the NVIDIA web site
`-q`	do not ask any questions; use the default answers
`-K`	only build the kernel module

At this point the legacy NVIDIA driver is installed and ready for use on this kernel. You may either start up an X session or return to 'init 5' (X11) system state.

6.3.3 Uninstalling the Legacy NVIDIA Driver

If you replace your legacy NVIDIA cards with newer ones that require the unified NVIDIA driver, you may uninstall the legacy driver and re-install the unified driver by following these steps:

1. Exit any currently active X session.

2. Login as root.
3. Exit X11 state and enter full multiuser mode:

   ```
   init 3
   ```

4. Unload the nvidia kernel module:

   ```
   modprobe -r nvidia
   ```

5. Uninstall the legacy nvidia driver:

   ```
   nvidia-installer --ui=none --uninstall
   ```

6. Insert the RedHawk Linux installation CD into the CD-ROM drive.
7. From the cdrom mount point/directory, issue the command:

   ```
   cd nvidia
   ```

8. Re-install the RedHawk unified driver:

   ```
   rpm -Uvh ccur-nvidia-glx-169.09.x86*.rpm
   ```

9. Update the kernel module dependencies for the NVIDIA driver:

   ```
   depmod
   ```

6.3.4 Restoring Unified NVIDIA Kernel Modules

The nvidia-installer normally restores the original unified **nvidia.ko** kernel driver module for all kernels as part of the uninstall process.

However, if the default unified **nvidia.ko** kernel module is accidentally removed, corrupted or is the wrong version for your currently executing kernel, you may recover from this situation by rebuilding and reinstalling the unified nvidia kernel module from the nvidia driver kernel source that resides within the RedHawk 4.2 kernel source tree.

For example, to restore the unified nvidia kernel module for a pre-built RedHawk Linux trace-static kernel:

1. Login as root.
2. Exit X11 state and enter full multiuser mode:

   ```
   init 3
   ```

3. Unload the nvidia kernel module (if loaded);

   ```
   modprobe -r nvidia
   ```

4. Rebuild the kernel modules and install them:

   ```
   cd /usr/src/redhawk-release-directory
   ccur-config -n -c
   make REDHAWKFLAVOR=-trace modules
   make REDHAWKFLAVOR=-trace modules_install
   depmod
   ```

6.4 Configuring Multiple Display Capability

There are several options for configuring two or more displays. The upper limit on the number of displays that can be driven effectively is application dependent. Theoretically, you can configure as many displays as you have video adapters.

It should be possible to use all the PCI Express (PCIe), PCI and AGP (if present) NVIDIA cards that are installed in your system, provided that all these cards are supported by either the unified or legacy NVIDIA driver (not a mix of unified and legacy cards). To maximize throughput, place the PCI cards in PCI slots that are not shared by other devices. Older iHawk systems may have one AGP slot, where the maximum number of displays run off the AGP card is limited by the number of heads on the card.

For configuring multiple adapters, execute the command **`lsirq -c 0300`** for a listing of all VGA compatible adapters. This will list the IRQ asignments for all VGA cards along with the busID (logical slot number) of each adapter. See **`lsirq(8)`** for more details about this command.

The following multiple display modes are available:

- One desktop spanned across two monitors.
  NVIDIA calls this TwinView. In this mode, the workspace size is twice as large as with a single monitor display. Windows can be dragged across displays, which can be positioned either vertically or horizontally. Switching between workspaces affects both displays simultaneously so it's easy to switch between data views that are linked together, for example, during demos and presentations. Multiple workspaces can be configured. See section **6.4.1** to configure this display mode.

- Separate desktops.
  This mode provides two or more separate desktops that act independently, for example a NightView RT debugger running on one monitor and an application running on another. Because the monitors work independently, displays are not synchronized as in TwinView mode and you cannot drag windows across monitors; however, it allows you more combinations for viewing data. See section **6.4.2** to configure this display mode.

The mode in use depends upon the settings in the X server configuration file when the X server is booted. You can easily switch between modes by maintaining more than one configuration file and copying the one you want to use to **`/etc/X11/xorg.conf`** before starting the X server.

Along with the original **`xorg.conf`** configuration file located in **`/etc/X11`**, there are additional sample configuration files. These samples illustrate configurations for TwinView, dual desktop, quad desktop and SLI. Use these as a reference when creating configuration files for use at your site.

NOTE

Parameters required for the multiple display modes are discussed below, but other fields in the configuration file may need to be modified depending upon other factors at your site. For complete information about the **`xorg.conf`** file, see the **`xorg.conf(5x)`** man page. The **`nvidia-xconfig`** program can be used to generate **`xorg.conf`** files; use the **`--help`** option or view the man page. Additional information can be found in the README file from the **http://www.nvidia.com** web site.

You will need to know the following information for each monitor (consult the documentation provided with the monitors):

- the range of possible refresh rates
- supported resolutions
- supported color depths

Note that entries for the "Driver" field in the "Device" section of the **xorg.conf** file have the following meanings:

"vesa" This is a generic Linux VGA adapter driver that will run any VGA compatible adapter. It is a rudimentary driver capable of running only low resolutions, color depths and refresh rates. It is most useful as a troubleshooting tool, providing some display capability when no other driver works.

"nv" This is the name of the generic NVIDIA driver supplied with the standard Red Hat X server. It runs any NVIDIA VGA adapter but does not drive a high performance NVIDIA graphics card to its full capabilities.

"nvidia" This tells the kernel to use the high performance NVIDIA driver installed by either the unified driver **ccur-nvidia-glx** rpm or the legacy driver NVIDIA run package. It must be entered manually in the "Driver" field of the **xorg.conf** file. This selection includes the NVIDIA X Server Settings tool that allows you to configure settings for special capabilities enabled on the adapter by the NVIDIA rpm and can be run under X by executing the following command: **/usr/bin/nvidia-settings &**.

6.4.1 Configuring One Desktop Across Two Monitors

To create a virtual desktop that displays across two monitors (for example, NVIDIA's TwinView), follow these steps:

1. Provide appropriate settings in the "Screen" section of the configuration file as shown below. The display mode and monitor characteristics are defined as "Options." Refresh rates and resolutions for each monitor and the location (orientation) of the second monitor used as the extended desktop are supplied. "TwinViewOrientation" can be "LeftOf", "RightOf", "Above" or "Below" to represent where the second monitor will be positioned relative to the first.

   Following is an example of the "Screen" section of the configuration file with possible settings for TwinView mode.

```
Section  "Screen"
         Identifier   "Screen0"
         Device       "Videocard0"
         Monitor      "Monitor0"
         DefaultDepth     24
         Option       "TwinView"
         Option       "SecondMonitorHorizSync" "30-97"
         Option       "SecondMonitorVertRefresh" "50-180"
         Option       "MetaModes" "1280x1024,1280x1024"
         Option       "TwinViewOrientation" "LeftOf"
         SubSection   "Display"
                Depth    16
                Modes "1280x1024" "1024x768" "800x600" "640x480"
         EndSubSection
         SubSection   "Display"
                Depth    24
```

```
                Modes "1280x1024" "1024x768" "800x600" "640x480"
        EndSubSection
EndSection
```

2. Verify that “nvidia” is specified for the “Driver” entry in the Device section.

6.4.2 Configuring Separate Desktops

For this mode, the Device, Screen, Monitor and ServerLayout sections in the configuration file are utilized. The number of displays available is limited only by hardware constraints. The examples below are for a dual display. For additional displays, add additional sections with appropriate designations (see **`xorg.conf.quaddesktop`** for examples). Follow the steps below to modify the configuration file.

1. First, update the ServerLayout section to reflect the multiple display configuration. Each screen to be used will be listed here and defined in Screen sections later in the file.

```
Section   "ServerLayout"
          Identifier   "nvidia dualdesktop"
          Screen  0    "Screen0" 0 0
          Screen  1    "Screen1" LeftOf "Screen0"
          InputDevice  "Mouse0" "CorePointer"
          InputDevice  "Mouse1" "SendCoreEvents"
          InputDevice  "Keyboard0" "CoreKeyboard"
EndSection
```

2. Then, create separate Device sections. In each section, provide the BusID (*bus*:*device*:*function*) of the graphics card associated with each display. This information displays when you use “**`lsirq -c 0300`**” or “**`cat /proc/pci`**”. In the case of a dual head adapter, specify separate sections (one for each head) using the same BusID. **Caution:** BusIDs can change when PCI cards are added or removed.

   The driver in each section in the example below is identified as "nvidia.” A separate screen is assigned in each section.

```
Section   "Device"
          Identifier   "nvidia0"
          Driver       "nvidia"
          BusID        "PCI:1:0:0"
          Screen       0
EndSection

Section   "Device"
          Identifier   "nvidia1"
          Driver       "nvidia"
          BusID        "PCI:1:0:0"
          Screen       1
EndSection
```

3. Next, create separate Screen sections, each using one of the devices identified in the Device sections; for example:

```
Section   "Screen"
          Identifier   "Screen0"
          Device       "nvidia0"
          Monitor      "Monitor0"
          DefaultDepth     24
          Subsection   "Display"
```

```
                Depth     24
                Modes "1600x1200" "1024x768" "800x600" "640x480"
        EndSubsection
EndSection

Section  "Screen"
        Identifier   "Screen1"
        Device       "nvidia1"
        Monitor      "Monitor1"
        DefaultDepth       24
        Subsection   "Display"
                Depth      24
                Modes "1600x1200" "1024x768" "800x600" "640x480"
        EndSubsection
EndSection
```

4. Finally, create separate Monitor sections, each describing the characteristics of one of the monitors to be used; for example:

```
Section  "Monitor"
        Identifier   "Monitor0"
        VendorName   "Viewsonic"
        ModelName    "G90f"
        DisplaySize  370   270
        HorizSync    30.0 - 97.0
        VertRefresh  50.0 - 180.0
        Option       "dpms"
EndSection
```

6.5 Connecting the Monitors and Booting the System

The graphics adapter card identifies multiple outputs as head 1, head 2, etc. With multiple monitors, the console output goes to head 1 during boot or when X is not running. The monitors on subsequent heads only display when X is running.

1. Connect monitor 1 to head 1 of the adapter, monitor 2 to head 2 of the adapter, etc.
2. To allow a recovery if there is a problem during the installation, set the system's default run level to enable the system to boot to console instead of starting up X.
   a. Edit **/etc/inittab**:

      vi /etc/inittab

   b. Locate the following line:

      `id:5:initdefault`

   c. Change the `5` to `3`, save and exit.

   d. Reboot the system with **init 3**.
3. Make a backup copy of the original configuration file, then copy the file you wish to use to **/etc/X11/xorg.conf**; for example:

   cp xorg.conf xorg.conf.orig
   cp xorg.conf.mydual xorg.conf

4. Startup X with **init 5**.

5. To test the success of the installation, run the following command:

   `startx`

   If successful, you will see the NVIDIA logo and then the desktop appears. Reset the system's default run level back to run level 5 (X11 startup mode).

   a. Using **`vi`** to edit **`/etc/inittab`**, locate the line reading:
      `id:3:initdefault`
      Change the `3` to `5`, save and exit.

   b. Execute the command **`init 6`**.

   If unsuccessful, the problem may be due to a misconfigured **`xorg.conf`** file. Review the **`/var/log/Xorg.#.log`** files along with the NVIDIA README document to debug your **`/etc/X11/xorg.conf`** file.

6. If you want to test the 3D performance of your card, perform the following:

   `glxgears`

   This should bring up three gears rotating at high speed with frames per second displayed in the associated window.

   If the setup is incorrect, you may get the following on-screen messages:

   ```
   Xlib: extension "GLX" missing on display ":0.0".
   Error: couldn't get an RGB, Double-buffered visual
   ```

   Refer to the NVIDIA README document for troubleshooting tips.

7.0. Additional Information

This section contains helpful information about issues that may be applicable at your site.

7.1 Installation/Configuration Issues

7.1.1 File System Backup Recommendations

All file systems should be archived or preserved prior to commencing a re-installation. This can be accomplished using normal backup methods. Concurrent recommends obtaining a new disk while keeping the original disk to ensure minimal data loss.

Table 7-1 contains a list of configuration files that may be useful in configuring a newly installed RedHawk system to match an existing configuration.

Table 7-1 Files Recommended for Backup

NOTE: This is not an all inclusive list and not all files are meant to be copied intact to a new system.

Path	Comments
/etc/hosts	To preserve hostnames on file.
/etc/fstab	To maintain existing mounts. Will not preserve mount points.
/etc/sysconfig/network/*	Network information.
/etc/sysconfig/network-scripts/ifcfg-*	Network device configuration.
/etc/resolv.conf	DNS configuration info.
/etc/nsswitch.conf	NIS/DNS configuration info.
/etc/ntp.conf /etc/ntp/step-tickers	NTP configuration. See ntpd(1).
/etc/rc.local /etc/rhosts	Startup info.
/etc/X11/xorg.conf or XF86Config	X11 server/device/screen configuration data.
/etc/modules.conf	For reference and comparison.
/etc/security/capability.conf	PAM configuration.
/etc/ssh/*	To preserve host keys.
/etc/inittab	For reference and comparison.
/etc/xinetd.d/	For reference and comparison. Capture any file customizations.
/var/spool/cron/	To preserve user crontab entries.
/sbin/chkconfig --list > <***save_file_name***>	For reference to configure new system.
/etc/*.conf	Customized configuration files.
/etc/sysconf/*	Customized configuration files.
/etc/pam.d/*	Customized PAM files.
rpm files	Those installed after previous system installation that need to be moved forward to the new system. Use **/bin/rpm -qal --last | more** to view.

7.1.2 File System Swap Size Guidelines

Table 7-2 supplies the recommended swap partition sizes for various sizes of main memory. The swap partition size is based on the size of the system's main memory plus an additional 5 percent. This swap space allocation should be adequate to accommodate system crash dumps.

Table 7-2 Swap Partition Size Guidelines

Main Memory Size (MB)	Recommended Swap Partition Size (MB)
256	269
1024	1075
2048	2150
3072	3225
4096	4300
5120	5376
6144	6451
7168	7526
8192	8601

7.1.3 Configuring Greater Than 4 GB of Memory on i386 Systems

If memory is added to an iHawk i386 system increasing the total physical memory to greater than 4 GB, the kernel configuration options that support large memory systems must be enabled in the kernel. Select the following options from the Processor Type and Features -> High Memory Support (4GB) submenu of the kernel Configuration GUI:

- 64GB
- Allocate 3rd-level pagetables from highmem

7.1.4 RCIM Connection Modes

Before installing an RCIM on an iHawk system, determine the connection mode. It is easier to connect the synchronization cable to the input connector before the RCIM is installed. An RCIM can be connected in one of the following four modes:

Isolated mode	There are no connections to any other RCIM.
Master mode	The RCIM is at the head of a chain of RCIMs. There is no cable connection going into this RCIM, only a cable connection going out. The RCIM master is unique in that it controls the synchronized clocks.
Pass-through Slave mode	The RCIM is connected to two other RCIMs. There is an input cable connection coming from the previous RCIM in the chain, and an output cable connection going to the next RCIM in the chain.
Final Slave mode	The RCIM is connected to one other RCIM. There is an input cable connection going into a final slave RCIM but no output cable connection coming out of it.

7.2 Considerations

7.2.1 Compiler Requirements

If you plan to build a kernel from source residing on your iHawk system on a different system, that system must be using **gcc-4.1.2-14** or later.

7.2.2 Hyper-thread Activation

Hyper-threading is configured by default in each of the standard RedHawk Linux kernels running on iHawk Series 860 systems. It can be disabled on a per-CPU basis using the **`cpu(1)`** command. However, hyper-threading is also defined in the system BIOS. The setting in the BIOS will supersede any configuration of this feature in the kernel. Therefore, if hyper-threading is turned off in the BIOS, it will not be available on your system even if it is configured in the kernel.

If your system is not behaving as expected in regards to hyper-threading based on the kernel configuration, check the BIOS setting and modify it as needed. Refer to your hardware documentation to determine which BIOS setting is involved.

Note that hyper-threading is not supported for the model 868 8-way iHawk system and should be disabled in the BIOS prior to booting a Version 5.1 RedHawk Linux kernel on this type of system.

8.0. Known Issues

Special consideration should be given to the following areas.

BIOS "Console Redirection"

The "Console Redirection" BIOS feature, when enabled, has been observed to interfere with the proper operation of the integrated VGA video and the XFree86 X server with some iHawk platforms, such as the Dell PowerEdge™ 6650 (iHawk Model HQ665).

NMI Button

Using the NMI button on the processor enters **kdb** on the console when configured. It cannot be used, however, to step repeatedly through an error condition.

kdb and USB Keyboards

While **kdb** does have some support for USB keyboards, it is found not to be reliable and is not supported in the RedHawk debug kernel. On systems with a USB keyboard, it is recommended that a serial console be configured and **kdb** used over the serial port. See Appendix H of the *RedHawk Linux User's Guide* for instructions for setting up the serial console.

PAM Capabilities

- Note that in order to use the **pam_capability(8)** feature with **ssh**, the **/etc/ssh/sshd_config** file must have the following option set:

  ```
  UsePrivilegeSeparation no
  ```

- Using Kerberos telnet services with PAM is not recommended. If krb5-telnet is on and Kerberos is not properly configured, the following error will occur upon logging in via telnet:

  ```
  login: Cannot resolve network address for KDC in requested realm
  while getting initial credentials
  ```

  The krb5-telnet xinetd service should be disabled:

  ```
  chkconfig krb5-telnet off
  ```

sadc(8) Issue

When a system configuration change is made that affects the number of CPUs (for example when booting with a uniprocessor kernel or changing hyper-threading configuration), the **sadc(8)** (sar data collector) program cannot successfully write data in the daily data file **/var/log/sa/sa**?? (?? is the day of the month). This results in **crond** sending an email every ten minutes to root@localhost with the following message:

```
Cannot append data to that file
```

To eliminate the email, remove or move the **/var/log/sa/sa**?? file for the current date.

Model 868 8-way iHawk Restriction

Hyper-threading is not currently supported on a model 868 8-way iHawk system. Hyper-threading should be disabled in the BIOS before booting RedHawk Linux on this system. Refer to the hardware documentation for the appropriate BIOS setting.

irqbalance

The **irqbalance** feature is disabled on RedHawk Linux. This Red Hat feature is meant to distribute interrupts equally across CPUs. It does not honor the IRQ affinity masks set up in **/proc/irq/***irq#***/smp_affinity** and therefore will cause interrupts to be routed to shielded CPUs.

This feature can be enabled/disabled at boot time with:

```
chkconfig irqbalance {on|off}
```

as well as started/stopped while the system is running with:

```
service irqbalance {start|stop}
```

Boot Problems with Monitor-embedded USB Hub and Flash Memory Reader

An iHawk system with a USB hub embedded within its attached monitor that includes a flash memory reader will fail to boot if the flash media is not installed. Once the flash media is installed, the system will boot.

Compatibility Issues for Ada programs

RedHawk 5.1 is based on Red Hat Enterprise Linux 5.1. Binary and source compatibility issues exist for MAXAda programs between RHEL 3.0 and RHEL 5.1.

Please see the Compatibility section in the *MAXAda for RedHawk Linux Version 3.5.1 Release Notes*, publication number 0898357-3.5.1, for details.

Inactive Virtual Terminal State After Changing Run Levels

Note that on some systems, changing from run-level 5 to run-level 3 may place the VGA console onto an inactive virtual terminal (e.g. /dev/tty8). If this happens, switching to an active virtual terminal will allow normal system operation to continue; press Ctrl-Alt-1 to switch to virtual terminal 1.

Instability with NVIDIA 8174 Driver and NVS400 Graphics Card

Some stability problems have been observed when using the Model 8174 NVIDIA driver with the NVIDIA NVS400 graphics card in Dell 6800 systems. If you experience system instability with this configuration, you will likely need to downgrade to version 7676 of the NVIDIA driver. Contact Concurrent technical support for additional information and assistance.

Excessive ksoftirqd Runs Affecting Determinism

In RedHawk 4.1, the IP route cache table size was changed from a fixed size of 4K entries to a dynamic size based on the amount of available memory. With 4 GB of memory, this table is sized at 128K entries. Every 10 minutes, a flush of that table is initiated. The approximate time to flush 4K entries is 1.5ms every 10 minutes. If the table is sized at 128K entries, this time can be in the area of 10ms to 15ms every 10 minutes. This can be problematic if network determininsm is at issue, especially in a single CPU system.

If this is an issue, the IP route cache table can be set to a fixed size using the grub command `rhash_entries=`*n*, where *n* is the number of table entries; e.g., `rhash_entries=4096` for 4K entries.

Pings to Multicast Addresses Disabled by Default

In RedHawk 4.1, a kernel.org change was imported that modified the default setting for the **sysctl** flag that allows pinging to a multicast address. Where previously it was set to enable broadcast and multicast ICMP echo (pings) and timestamp requests, it was changed to disable that functionality.

There are two methods in which to change this flag if you need to ping multicast addresses:

- **The sysctl(8)** utility changes the value in a running kernel and takes effect immediately; no rebuild or reboot is needed:

  ```
  # sysctl -w net.ipv4.icmp_echo_ignore_broadcasts=0
  ```

- To initialize the parameter to the desired value on every reboot, add the following command and **sysctl** parameter to **/etc/sysctl.conf**:

  ```
  # Controls broadcast and multicast ICMP echo and timestamp requests

     net.ipv4.icmp_echo_ignore_broadcasts = 0
  ```

Grub Option Holdovers from Previous Kernel Installations

Be aware that when kernel packages are installed, they add a grub entry with options associated with that kernel (e.g., debug kernels add “crashkernel=64@16M,” other kernels add “quiet,” etc.). In addition, all kernel packages also copy the default grub options into their grub entries. These default grub options are taken from the kernel entry that is currently marked as the default kernel to boot in **grub.conf**.

Depending upon which kernel is designated the default booting kernel and which kernels are installed subsequently, you may find that grub entries that are inherited from the default kernel and previously installed kernels are not suitable for the currently running kernel.

vdso and glibc 2.3.2

Changes to the location of the vdso page in the user address space in RedHawk Version 4.2 may result in a segfault if using **glibc** version 2.3.2 (RedHawk 5.1uses **glibc** 2.5).

If you are experiencing these segfaults, vdso can be disabled using the boot command **vdso=0** or **echo 0 > /proc/sys/vm/vdso_enabled**. To determine whether you have vdso enabled, **cat /proc/1/maps** and look for a vdso entry; for example:

```
b7f31000-b7f32000 r-xp b7f31000 00:00 0 [vdso]
```

USB Ports Non-functional on Supermicro Boards

Some Supermicro boards (Model X6DA8-G2, CCUR part number 820-2010483-913) may have non-functional USB ports accompanied by the following message:

```
USB 1-1: new high speed USB device using ehci_hsd and address 2 USB 1-1:
device descriptor read/64, error -110
```

This can be corrected by changing the BIOS setting for both SATA [0/1] IDE Interface and SCSI interface controllers to "ENABLED".

Configure Quad Port Ethernet

If your system has a quad port Ethernet, you may need to edit **/etc/modprobe.conf** and configure the ports to the appropriate driver. The example below assigns the quad port to eth0 through eth3 and uses the Intel e1000 module (Gigabit Network driver). For readability, convention has these lines inserted at the top of **/etc/modprobe.conf**.

```
alias eth0 e1000
alias eth1 e1000
alias eth2 e1000
alias eth3 e1000
```

Start haldaemon to burn files onto CD/DVDs

The **haldaemon** service has been shown to interfere with real-time performance and is turned off by default. However, it must be running in order to burn a file (e.g. iso) onto a CD or DVD from the context menu of the file. To burn a file onto a CD or DVD, first start the **haldaemon** service:

```
$ service haldaemon start
```

After the copying process is complete, stop the service:

```
$ service haldaemon stop
```

Issues with Tigon3 Network Adapter

Significant reliability and operability problems have been detected with the Tigon3 (BCM95700A6) network adapter. As yet, there are no solutions for these problems under RedHawk 5.1.

Deadline Support in FBS

Scheduling with deadlines in kernel update 5.1.3 is not supported. Full support is included in RedHawk version 5.2.1.

9.0. Software Updates and Support

9.1 Direct Software Support

Software support is available from a central source. If you need assistance or information about your system, please contact the Concurrent Software Support Center at our toll free number 1-800-245-6453. For calls outside the continental United States, the number is 1-954-283-1822. The Software Support Center operates Monday through Friday from 8 a.m. to 5 p.m., Eastern Standard Time.

Calling the Software Support Center gives you immediate access to a broad range of skilled personnel and guarantees you a prompt response from the person most qualified to assist you. If you have a question requiring on-site assistance or consultation, the Software Support Center staff will arrange for a field analyst to return your call and schedule a visit.

You may also submit a request for assistance at any time by using the Concurrent Computer Corporation web site at **http://www.ccur.com/isd_support_contact.asp**.

9.2 Software Updates

Updates to Concurrent RedHawk software can be obtained via Concurrent's Software Portal. There are three ways of installing product updates:

- Using the Network Update Utility (NUU) installed on your RedHawk system
- Manual installation after browsing and downloading individual RPMs from Concurrent's software repositories
- Building a customized Update CD using Concurrent's web site for immediate download

9.2.1 Updating via NUU

NUU supports installation and updating of software products from Concurrent software repositories over a network. NUU utilizes Yum and the RPM subsystems to install and update software.

NUU is installed automatically with RedHawk, however, you should configure it to include all of the Concurrent software products installed on your system.

Clicking on the "Updates (NUU)" icon on your desktop launches NUU to check to see if Concurrent updates are available for your system.

Instructions for configuring NUU can be found in the **QuickStart.pdf** document available at **http://redhawk.ccur.com/updates/QuickStart.pdf.**

The first time you invoke NUU you will need to specify your redhawk.ccur.com Login ID and Password that were provided in the shipping documents accompanying your system. If you require assistance, refer to "Direct Software Support" on page 63.

9.2.2 Installation of RPMs Manually Downloaded

You can browse Concurrent's software repositories to locate updated RPMs and download them for manual installation.

Access the RedHawk Updates web site (**http://redhawk.ccur.com**) by clicking on the "Concurrent Software Portal" icon on the desktop. Accessing this web site displays the following screen:

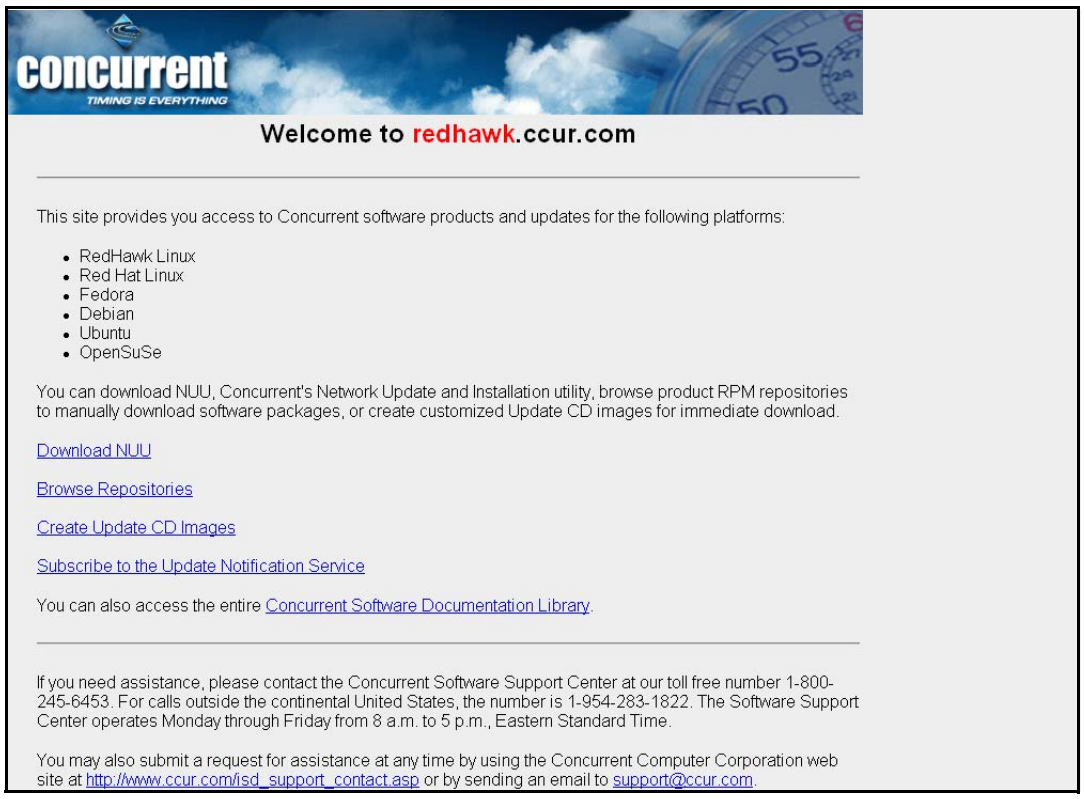

Clicking on the **Browse Repositories** link takes you to an authentication page.

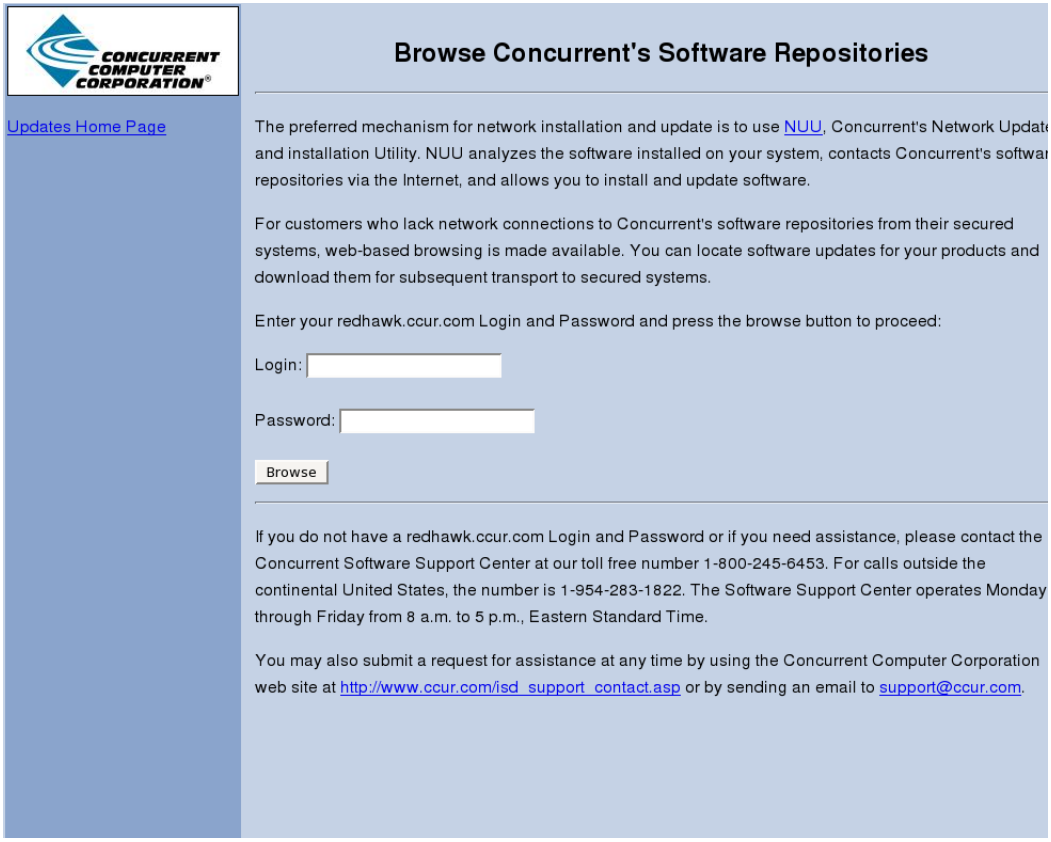

Enter your redhawk.ccur.com Login ID and Password and click the Browse button.

Select the products of interest and architecture from the following pages to see the list of RPMs in the product software repository.

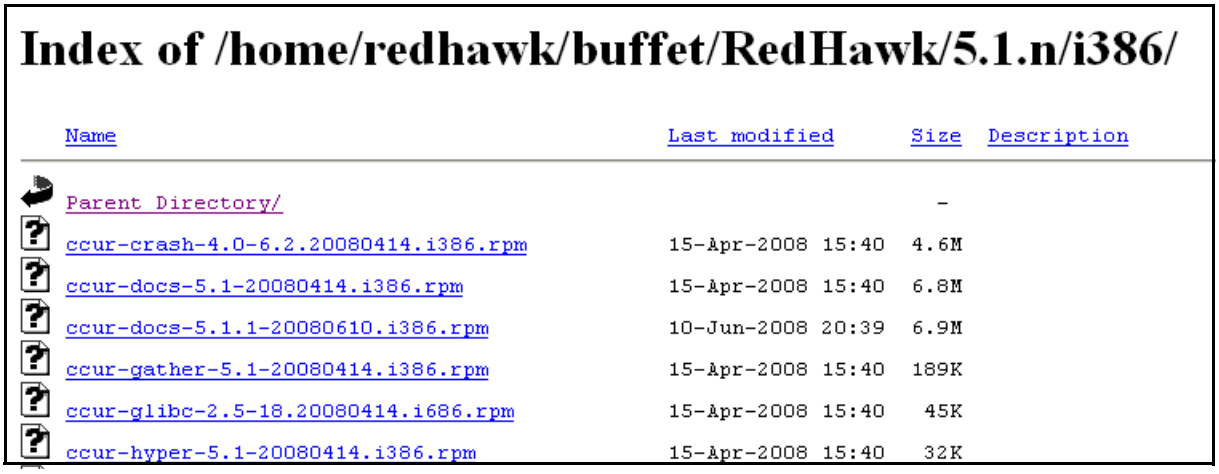

To locate the latest RPMs in the repository, click on the Last modified column header to sort by date. You may need to click twice to set the sort order to newest-to-oldest.

After locating the RPMs of interest and downloading them to your system, you can manually install them.

To install newly downloaded packages, follow these steps:

1. Log in as root and take the system down to single-user mode:
   a. Right click on the desktop and select Open Terminal.
   b. At the system prompt, type **`init 1`**.
2. Change directory to the location of the updates and issue the following command:

   ```
   rpm -Uvh *.rpm
   ```

   The time it takes to install will vary depending on the number of updates being installed.
3. When complete, exit single-user mode (Ctrl-D).

NOTE

If you have installed an update containing new RedHawk kernels on a system that has post-installation loadable modules present, those modules must be recompiled for the new kernel; for example, an NVIDIA driver that is a later version than the one included with RedHawk or any third party package that uses loadable modules.

9.2.3 Customized Update CDs

You can use Concurrent's Software Portal to create a customized Update CD for your system which you can then download and burn onto physical media, or simply mount as an ISO9660 image.

Update CDs have customized copies of product software repositories and a simple graphical interface for selecting packages for update and installation. These CDs use NUU (described above) to talk to the CD to obtain packages -- no network access is required during installation via Update CDs.

Access the RedHawk Updates web site (**http://redhawk.ccur.com**) by clicking on the "Concurrent Software Portal" icon on the desktop, then click on Create Update CD images.

You will need to enter your redhawk.ccur.com Login ID and Password and then you can select the products to update. A CD image is built as part of the web session. At the end of the session, you can immediately download it for subsequent installation.

9.2.4 Automatic Notification of Available Updates

When updates, incremental releases or major releases occur, those who have subscribed to Concurrent's automatic notification service are automatically sent an email with a link to the redhawk.ccur.com site describing the updates relative to your company. Multiple email addresses can be subscribed per customer site.

To subscribe, unsubscribe or change notification options, click on the “Subscribe to the Update Notification Service” link on the **redhawk.ccur.com** site.

9.3 Documentation Updates

For the latest documents, including updated Release Notes, User Guides and FAQ, go to Concurrent’s documentation web site at **http://redhawk.ccur.com/docs**.