Linux® OS Version 4.2.4 Release Notes

June 2008

0898003-4.2.4

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 4.2, including kernel update versions 4.2.1 through 4.2.4.

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 4 distribution. The RedHawk installation CDs 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 35 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 where interrupt response of less than 30 microseconds is guaranteed.

RedHawk Linux has at least the same level of POSIX® conformance as 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 for less intrusive application control. 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 that are delivered with RedHawk Linux Version 4.2 for each Concurrent system.

WARNING

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

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, 000]** **iHawk 880 Series [HQHS2]** **ImaGen [HQ0G1]**
Operating System CDs:	
Red Hat Enterprise Linux 4.0 Update 4 (x86) (9 discs)	
Red Hat Enterprise Linux 4.0 with Update 4 Updates for i386 & EM64T Systems [LXUPDATE WA-IA32 Version 4.2]	
RedHawk Linux OS for i386 & EM64T Systems [WA-IA32 Version 4.2]	
RedHawk Linux Documentation CD [LXDOCKIT i386 & EM64T WA-IA32 Version 4.2]	
Software Updates for RedHawk	
Optional Product CDs:	
RedHawk Linux Frequency-Based Scheduler for i386 & EM64T Systems [WU1210-JL Version 4.2]	
NightStar RT for RedHawk [WU1220-LDV 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 4.2	
RedHawk Linux Cluster Manager for i386 & EM64T Systems [WA9017-L] Version 4.2	
RedHawk Global File System for i386 & EM64T Systems [WA9018-L] Version 4.2	
RedHawk High Availability NFS for i386 & EM64T Systems [WA9019-L] Version 4.2	
RedHawkOpenFabrics Enterprise Distribution for i386 & EM64T Systems [WA9020-L] Version 4.2	

Processor Type:	**AMD Opteron™ (32-bit and 64-bit) and EM64T x86_64 (64-bit)**
Model:	**iHawk 860 Series [HQ685, 680, 285, 280, 185, 067, 047, 000]** **iHawk 870 Series [HR844, 430, 221, 210, 00W, 00T]** **iHawk 880 Series [HQHS2]** **ImaGen [HR0G1]**
CDs:	
Red Hat Enterprise Linux 4.0 Update 4 (AMD64) (10 discs) **Red Hat Enterprise Linux 4.0 with Update 4 Updates for x86_64 Systems** [LXUPDATE WA-AMD64 Version 4.2] **RedHawk Linux OS for x86_64 Systems** [WA-AMD64 Version 4.2] **RedHawk Linux Documentation CD** [LXDOCKIT x86_64 WA-AMD64 Version 4.2] **Software Updates for RedHawk**	
Optional CDs:	
RedHawk Linux Frequency-Based Scheduler for x86_64 Systems [WU1210-JA Version 4.2] **NightStar RT for RedHawk** [WU1220-ADV 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 4.2 **RedHawk Linux Cluster Manager for x86_64 Systems** [WA9017-L] Version 4.2 **RedHawk Global File System for x86_64 Systems** [WA9018-L]Version 4.2 **RedHawk High Availability NFS for x86_64 Systems** [WA9019-L] Version 4.2 **RedHawk OpenFabrics Enterprise Distribution for x86_64 Systems** [WA9020-L] Version 4.2	

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, 2.0 is the initial release of major version 2, and 2.1 is a subsequent release to version 2.0. Both 2.1 and 2.2 constitute complete product releases and do not require a previous release to be currently installed on the system; both are available on CD 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 4.2.1 kernel update can only be installed on a system currently running RedHawk Linux version 4.2.

1.4 System Updates

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

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 53.

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 accessible.

2.0. Documentation

2.1 RedHawk Linux Documentation

Table 2-1 lists the documentation provided with RedHawk Version 4.2 as PDF files. After installing the RedHawk Linux Documentation CD, the files are available for viewing:

- by clicking on the “Documents” icon on the desktop
- by clicking on the red document entries in Table 2-1 (optional product documentation is available only if the product has been installed)
- from Concurrent’s documentation web site at **redhawk.ccur.com/docs**

Table 2-1 RedHawk Linux Documentation

RedHawk Linux Operating System Documentation	Document Number
RedHawk Linux OS Online Documentation Roadmap	0898002-474
RedHawk Linux OS Version 4.2.4 Release Notes	0898003-4.2.4
RedHawk Linux User’s Guide	0898004-520
Real-Time Clock & Interrupt Module (RCIM) PCI Form Factor User’s Guide	0898007-320
iHawk Optimization Guide	0898011-100
RedHawk Linux FAQ	N/A
Optional RedHawk Product Guides	
RedHawk Linux Frequency-Based Scheduler (FBS) User’s Guide	0898005-260
RedHawk High Availability NFS Installation Guide	0898018-010
RedHawk OpenFabrics Enterprise Distribution Installation Guide	0898019-000
Guide to SNARE for Linux	N/A
Optional RedHawk Product Documentation Sets	
RedHawk Linux Cluster Manager	
RedHawk Global File System	
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 and User Guides, can be accessed from Concurrent’s documentation web site at **http://redhawk.ccur.com/docs**. An updated RedHawk FAQ can be accessed at **http://www.ccur.com/isdfaq/RedHawk-FAQ.pdf**.

Printed copies can be ordered by contacting the Concurrent Software Support Center. Refer to “Direct Software Support” on page 53 for details.

2.2 Vendor Documentation

Red Hat Enterprise Linux operating system documentation is included as PDF files on the Red Hat CDs 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 4 with updates
- RedHawk Linux operating system with updates
- 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 Global File System (optional)
- RedHawk High Availability NFS (optional)
- RedHawk OpenFabrics Enterprise Distribution (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 Versions 4.2 and any subsequent kernel updates.

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 4.2. See the section "Software Updates" on page 53 for instructions for downloading and installing kernel updates.

Version 4.2.4 This kernel update includes:

- Integrated Intel® e1000 driver version 8.0.1
- Integrated Intel e1000e driver version 0.2.9.5
- Integrated Intel igb driver version 1.2.24
- Build Advanced Linux Sound Architecture (ALSA) as a module
- Support for Intel HD Audio
- Upgrades to **`/usr/share/hwdata/pci.ids`**, **`lsirq(8)`** and addition of **`update-pciids`** script

Version 4.2.3 This kernel update includes:

- Added driver support for 3ware 9690SA SAS PCIe RAID Controller
- Bug fixes for the 3ware 9650SE SATA-II PCIe RAID Controller
- Fix for Dell T5400 reboot problem; added T5400 to reboot_dmi_table
- Fix kfree loop latency; open a preemption window in rcu callback cleanups
- Fix **`vmsplice(2)`** security vunerability
- Fix to rcim kernel headers and init script

Version 4.2.2 Added driver support for 3ware 9650SE SATA-II PCIe RAID Controller.

Version 4.2.1 This kernel update includes:

- fix for the "ntfilter warning: bad timestamp" bug
- completed support for all RCIM II cable features
- improved NTP synchronization on systems with poor hardware clocks
- fix for Dell 490 reboot problem; added 490 to reboot_dmi_table
- backport of some NFS changes to fix potential races and panics
- e1000 driver updated to 7.6.9 for Intel PRO/1000 PT quad port server adapter
- kdump kernel config synchronized with static configs
- increase maximum number of FBSs from 4 to 63

4.2 Enhancements in Version 4.2

4.2.1 RedHawk based on RHEL 4 Update 4

RedHawk Version 4.2 uses Red Hat Enterprise Linux 4 Update 4 as its base with a merge of kernel.org version 2.6.18.8.

4.2.2 32-bit Support on AMD Opteron Processors

RedHawk for x86_64 Systems (WA-AMD64) now includes support for running the 32-bit version of RedHawk on AMD Opteron based systems in addition to the 64-bit version already supported. This capability benefits those who want the advanced AMD architecture, while requiring their OS to run in native 32-bit mode for specific device driver compatibility.

4.2.3 Unified and Legacy NVIDIA Display Drivers

The NVIDIA Linux Display Driver (version 1.0-9746), called the “unified” driver, is included in the RedHawk Linux installation. The unified driver supports the latest and more recent NVIDIA GeForce and Quadro based video cards. During RedHawk installation, the type and model of video adapters present in the system is detected and the unified driver is built for all pre-built and custom RedHawk Linux kernels when one or more NVIDIA video cards is present.

In addition to the unified driver, an alternative “legacy” NVIDIA Linux Display Driver (version 1.0-9631) is included that supports older NVIDIA GeForce and Quadro based Graphics Processor Units (GPUs).

For complete information about these NVIDIA display drivers, see section “7.0. NVIDIA Graphics Configuration”.

4.2.4 SNARE Upgrade

Support for SNARE for Linux in RedHawk has been upgraded to version 0.9.8. This upgrade no longer includes a GUI interface; however, all previous features are included as well as the addition of report generation.

User documentation for SNARE for Linux is provided by the *Guide to SNARE for Linux*. Configuration of SNARE in the RedHawk 4.2 kernel is provided in the RedHawk FAQ. Both are included with RedHawk user documentation.

4.2.5 New Functionality for kexec-based Crash Dumps (kdump)

A number of new features has been added to **`kdump`** for taking crash dumps:

- **`kdump`** is off by default; it must be turned on to use.
- New command line arguments to **`kdump`** have been added to turn **`kdump`** on and off, query conditions and display current values.
- New configuration options in **`/etc/sysconfig/kdump`** have been added.
- New scripts are available for customizing the actions of **`kdump`**.
- The **`namelist`** directory is renamed to **`kerninfo`**.
- The required boot parameter **`crashkernel=64M@16M`** can now be set with a size of 32M, 64M or 128M.
- Dumps to NFS roots are now supported.
- A delayed start for **`kdump`** can be configured for crash dumps in diskless/cluster environments.
- **`kdump`** crash dumps are now supported on x86_64 systems; LKCD is no longer supported.
- The **`crash(8)`** utility has been upgraded.

Details can be found in the **kdump(8)**, **kexec(8)** and **crash(8)** man pages, through comments in the **/etc/sysconfig/kdump** configuration file and in the "Kernel Debugging" chapter of the *RedHawk Linux User's Guide*.

4.2.6 kdb Upgrade

An upgrade to **kdb** is included in RedHawk Version 4.2. Changes include a new out of memory (oom) feature through **/proc/sys/vm/panic_on_oom**.

Documentation for **kdb** is included online at **/***kernel-source***/Documentation/kdb** and in the Kernel Debugging chapter of the *RedHawk Linux User's Guide*.

4.2.7 InfiniBand Support through RedHawk OFED

Kernel support for OFED has been added to RedHawk Version 4.2.

RedHawk OpenFabrics Enterprise Distribution (OFED) is a collection of software components intended for use on a computer cluster constructed as an InfiniBand network.

InfiniBand (IB) is a low-latency, high-performance, serial I/O interconnect. InfiniBand is deployed primarily in server clusters ranging from two to thousands of nodes. In addition to connecting servers, InfiniBand can also be used to connect communications and storage fabrics in data centers.

RedHawk OFED is purchased separately. See your Concurrent representative for details.

4.2.8 GFS Support through RedHawk GFS

Kernel support for GFS (Global File System) has been added to RedHawk Version 4.2.

RedHawk Global File System (GFS) is a common storage file system that allows all computers on a network to locally maintain an identical, shared block device. GFS reads and writes to the block device like a local filesystem. Changes made to the filesystem on one computer are immediately written to all other computers in the cluster. GFS uses a lock module to allow the computers to maintain I/O consistency. GFS provides for fast local reads at the expense of write performance.

RedHawk GFS is purchased separately. See your Concurrent representative for details.

4.2.9 High Availability NFS Support through RedHawk HA-NFS

Kernel support for High Availability NFS is included in RedHawk Version 4.2.

RedHawk High Availability NFS (HA-NFS) combines all the components needed for a high availability cluster under RedHawk Linux. An HA-NFS cluster is a pair of computers acting as a single dual-redundant file system. In normal operation, each node acts as a primary supplier for a subset of the total filesystem resources. Upon failure of either node, the resources of the failed node are acquired and made primary on the working node.

RedHawk HA-NFS is purchased separately. See your Concurrent representative for details.

4.2.10 memexact Utility for Physical Memory Reservation

The **memexact(1)** utility processes command options and produces appropriate memory reservation commands based on the contents of **/proc/iomem** or a file specified on the command line.

Refer to the man page for details. An example of its use can be found in the *RedHawk Linux User's Guide* in the "Real-time Performance" Chapter.

4.3 Modifications in Version 4.2

4.3.1 Changes to System Timekeeping

Previously, RedHawk used a time base and time interpolator for system timekeeping. Through enhancements in the 2.6.18 kernel upon which RedHawk Version 4.2 is based and the addition of Thomas Gleixner's **2.6.18-hrt-dyntick2** patch, a clocksource mechanism is now used. RedHawk takes this a step further by using a TSC based clock to satisfy most timekeeping requests and disciplining it to a second clocksource to improve the stability of the clock. When an RCIM is present in the system, the RCIM is used as the second clocksource; otherwise, the HPET or PM timer are used.

Reading the file **/sys/devices/system/clocksource/clocksource0/current_clocksource** displays the current secondary clocksource. Writing the name of another clocksource to this file using **echo(1)** changes the assignment.

The global timer handling also changed. Previously this function was driven by the PIT timer at a fixed period. It is now emulated by using the local timer interrupt. Local timer shielding controls the set of processors which perform the global timer function.

See the txt files in the **/***kernel-source***/Documentation/hrtimers** to learn more about the new timekeeping features.

4.3.2 High Resolution POSIX Clocks and Timers Now in rt Library

The POSIX clock and timer functions provided in **libccur_rt** are being deprecated as they are now also available in **librt** through high resolution enhancements in the 2.6.18 kernel upon which RedHawk Version 4.2 is based.

Through the standard **librt** POSIX timers, the CLOCK_PROCESS_CPUTIME_ID and CLOCK_THREAD_CPUTIME_ID clocks are available for use. However, note that these clocks cannot be used with local timer shielding.

The files in **/proc/sys/kernel/posix-timers** previously used to control expiration rates are no longer needed and have been removed.

4.3.3 POSIX Message Queue Routines Removed From ccur_rt Library

The POSIX message queue routines that were previously included in **libccur_rt** are now only available in **librt**. Binaries that linked to the message queue routines in **libccur_rt** will need to be recompiled. Refer to the man pages for information.

4.3.4 Kernel Level Priority Inheritance Support Removed

Kernel level support for priority inheritance, previously included through the PRIO_INHERIT kernel configuration parameter, is no longer included in RedHawk. All user level priority inheritance mutex services are still available as documented in Chapter 5 of the *RedHawk Linux User's Guide*.

4.3.5 FBS Modifications

- When establishing a real-time clock using **`fbssetrtc`** or the **`stc`** command, the number of clock counts per minor cycle now ranges from 2 to 65535. Although a count of 1 cannot be used, a timing interval equal to a resolution value can be set by using the next lower resolution value and a count of 10; e.g., a resolution of 1,000 and count value of 10 results in a timing interval of 10,000 microseconds.
- Limits have been set for two user-defined values that are used by the fbsched driver in calculating memory to be kmalloc'ed. The number of tasks allowed to be scheduled is now limited to 2048 tasks and the number of cycles allowed is limited to CYCLES * TASKS <= 0x100000.

5.0. Installation Procedures

Use the installation procedures provided here under the following circumstances:

- generating or regenerating a system
- upgrading a system currently running Version 1.x to Version 4.2. Refer to "File System Backup Recommendations" on page 46 before starting the installation.

If upgrading from Version 2.x or 4.1 to Version 4.2, follow the steps in "Upgrade Procedures" on page 27.

Review the following points before installing the software:

- Read through the entire installation procedure once before actually performing the installation.
- It is **strongly recommended** that installation and updates of software be performed in single-user mode to guard against the possibility of corrupting a running application.
- 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 plan to configure networking using fixed numbers (recommended), 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.
  - HPET timers should be disabled. See the section "Known Issues" for more information.

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 completing the steps below. (If you are uncertain which version is applicable, refer to the "Product Media" section on page 2.) During installation, refer to Red Hat's online help if needed.

1. Power on the system, if needed, to open the CD-ROM drawer.
2. Insert Red Hat Enterprise Linux Installation Disc 1 in the CD-ROM 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. If you wish to check the Red Hat media, type `linux mediacheck` at the `boot:` prompt (optional). Otherwise, press Enter (or allow this screen to timeout after a minute) 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, it may be available on the Red Hat installation CD, or it may be available for download from the RedHawk Updates web site. If you need to download the driver image file from the RedHawk Updates web site, refer to section "Creating a Driver Disk for Red Hat Installation" on page 48.

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 CD 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. At the Welcome screen, no action is needed. Proceed to the next screen.

   - For language selection, choose English language. Proceed to the next screen.

   - When using the GUI interface:

     - for keyboard configuration, accept keyboard defaults. Proceed to the next screen.

     - If a previous installation of RedHawk or Red Hat Enterprise Linux existed on the hard drive, you may be prompted to choose whether to perform an upgrade or an install. Choose "Install Red Hat Enterprise Linux WS".

   - For disk partitioning, choose "Manually Partition with Disk Druid" and proceed to the next screen to partition the disk manually and create the partition sizes as shown in Table 5-1 below. Note that this table shows recommended partition sizes for a 36 GB disk or larger.

     NOTES

     All partitions on SCSI drives must be placed on **/dev/sda**; all partitions on IDE drives must be placed on **/dev/hda**.

     If the system has multiple drives, you must deselect all other drives.

     Standard RedHawk does not support software RAID and Concurrent does not recommend its use.

   - At the Partitioning screen, first delete all existing partitions by highlighting each line and selecting DELETE. This must be done one partition at a time.

     To add partitions, highlight the disk drive and select NEW. Provide the mount point, filesystem type and size. Select the Force to be primary partition box at the bottom of each partition setup page. Select OK to continue to set up the next partition. **Note:** If the system has multiple drives, the additional drives listed under Allowable Drives must be unchecked each time a new partition is added.

Table 5-1 Recommended Disk Partition Sizes

Mount Point	Type	Filesystem Type	Partition Size
/boot	primary	ext3	500 MB
/	primary	ext3	16384 MB (16 GB)
swap (See Note below)	primary	swap	Size of system's main memory **plus** 5 percent additional. (See Table 8-2 on page 47 for suggested swap sizes)
/home	primary	ext3	Select "Fill to maximum allowable size."
Note: Select the **swap** file system from the list of filesystems.			

- Proceed to the next screen when finished with partitioning.
- For configuring the Boot Loader, accept GRUB boot loader defaults. Enter a GRUB password if desired (not recommended). Proceed to the next screen.
- For Network configuration, configure networking as appropriate (fixed network numbers are recommended - you will need address information if using fixed numbers). Click Edit for the selected network device, e.g., eth0. In the popup box that appears, deselect "Configure using DHCP", and supply the host IP address and netmask. Click OK to return to the Network Configuration screen.
- For Hostname specification, select manual configuration. Supply the hostname for your system and specify the gateway and DNS servers appropriate to your site. Proceed to the next screen after completion.
- For Firewall configuration, choose "No firewall". Also, it is highly recommended that SELinux be disabled; click on the SELinux selection button and select "Disabled" from the choices presented. Proceed to the next screen.

NOTES

You will be presented with a dialog box warning you that choosing "No firewall" is not recommended by Red Hat. Click "Proceed" to continue (if necessary you can alter the firewall configuration after RedHawk installation is completed).

If you leave SELinux enabled (set to the Active choice) you may experience problems mounting disk media in the CD/DVD drive. Disabling SELinux until all product installations have been completed is highly recommended.

- For Additional Language Support, choose any additional languages as necessary. Proceed to the next screen.
- For Time Zone Selection, select the appropriate time zone, then proceed to the next screen.
- To set the Root Password, enter the root password and confirmation, then proceed to the next screen.
- For selecting Packages to be installed, select "Customize the set of packages to be installed", then proceed to the next screen.
- At the list of Package Groups to install, scroll down to the bottom of the list and select "Everything." Proceed to the next screen.
- A screen preparing you for the installation of Red Hat Enterprise Linux now appears. Click Next to continue.
- A list of the required install media is displayed. Verify that you have the CDs listed and click Continue to proceed.

  At this point in the installation, the disk will format and packages on Red Hat Installation disc 1 will be installed. When disc 1 is completed, you will be prompted to insert disc 2. Remove disc 1 and

insert disc 2, then click OK. Take the same action when prompted for the remaining Installation discs and finally disc 1 again.

6. Remove all media from the drives and store appropriately. Make the selection to reboot the system. Verify that Red Hat boots up properly.

7. When using the GUI installation interface, once the system boots up successfully, proceed as follows:

   - At the Welcome screen:
     Click Next to continue.

   - At the License Agreement screen:
     Click "Yes, I agree to the License Agreement", then click Next.

   - At the Date and Time screen:
     Set the appropriate date and time. Enable NTP if desired. Click Next to continue.

   - At the Display screen:
     Select monitor type and resolution, if needed, then click Next to continue.

   - At the Red Hat Login screen:
     Select "Tell me why I need to register and provide a Red Hat login." and click Next. At the Why Register? screen, select "I cannot complete registration at this time. Remind me later." and click Next.

   - At the System User screen:
     Enter a username and password for each account you wish to set up at this time (this can be done after the full installation). Click Next to continue. If no accounts are added, a warning box will be displayed with a recommendation that an account other than 'root' should be provided. Just click Continue in the warning box to continue to the next screen.

   - If your system has a sound card installed, the Sound Card screen will appear. At the Sound Card screen:
     Click Play test sound if desired. Click Next.

   - At the Install Additional CDs screen:
     You are presented with the option of installing additional Red Hat CDs. If you wish to install any of them now, click on the Install button for the desired CD and follow the instructions that display. Click Next to continue.

   - At the Finished Setup! screen:
     Click Next to continue.

8. Initial Red Hat installation is now complete. The login window should appear at this point. Remove any media from respective drives. Continue with the steps in section **5.2** to install Red Hat software updates.

5.2 Installing Red Hat Software Updates

The disc labeled “Red Hat Enterprise Linux 4.0 with Update 4 Updates LXUPDATE Version 4.2” contains the latest updates released by Red Hat and validated against RedHawk Linux. These updates are important to the proper operation of RedHawk Linux.

To install the Red Hat software updates, perform the following steps.

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 you used a driver diskette to install Red Hat Enterprise Linux,** type **`export NO_KERNEL_UPDATE=1`**. If you did not use a driver diskette during installation, skip to the next step.
3. Locate the disc labeled “Red Hat Enterprise Linux 4.0 with update 4 Updates” appropriate to your system’s architecture and insert it into the CD-ROM drive.
4. If the cdrom device does not mount automatically, execute the **`mount`** command with the mount point that is designated in **`/etc/fstab`**.

   NOTE: **`/media/cdrom`** is used in the examples in this section; e.g.:

   ```
   mount /media/cdrom
   ```

   If the cdrom entry is not in **`/etc/fstab`**, executing the following command will mount the cdrom device at **`/media/cdrom`**:

   ```
   mount /dev/cdrom /media/cdrom
   ```
5. Execute the following commands to install the update software:

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

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

   NOTE: This process may take some time to complete.
6. When the installation completes, execute the following commands:

   ```
   cd /
   umount /media/cdrom
   eject
   ```
7. Remove the disc from the CD-ROM drive and store.
8. **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.

9. Type “reboot” at the system prompt and verify that the new Red Hat kernel boots.
10. Continue with the steps in the next section to install RedHawk Linux.

5.3 Installing RedHawk Linux

After installing Red Hat and its updates, 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 4.2" appropriate to your system's architecture and insert it into the CD-ROM drive.

4. If the cdrom device does not mount automatically, execute the **`mount`** command; for example:

   ```
   mount /media/cdrom
   ```

5. Execute the following commands to install RedHawk Linux:

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

   NOTES: 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 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-ROM drive and store.

8. Locate the disc labeled "RedHawk Linux Documentation CD Version 4.2" appropriate to your system's architecture and insert it into the CD-ROM drive. Use **`mount`** if needed to mount the cdrom device.

9. Execute the following commands to install RedHawk documentation:

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

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

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

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

11. Exit single-user mode (Ctrl-D).

12. 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.

13. 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 **8.1.3** on page 47 for details.

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

5.4 Installing the RCIM

The following serves as an installation checklist for installing a PCI RCIM 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) PCI Form Factor 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.

5.4.1 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 47).

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 PCI slot 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 RCIM 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 *Real-Time Clock and Interrupt Module (RCIM) PCI Form Factor User's Guide* for configuration options.

5.4.2 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.
   ```

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
   ```

5.4.3 Downloading NTP Updates for RCIM II with GPS Module

If you installed the RCIM II model equipped with the optional GPS module, access the RedHawk Updates web site to download the ccur-ntp update required for its use if it has not already been downloaded. If you do not have the RCIM II board with GPS module, skip to section **5.5**.

1. To download the **ccur-ntp** update, refer to “Software Updates” on page 53 for details.
2. To install the package:
   a. Change directory to the location of the update and unpack the tarball; for example:

      ```
      # tar -xvf ccur-ntp-4.2.0.a.20040617-4*.tar
      ```

   b. Determine if ntp or ccur-ntp is already installed:

      ```
      # rpm -q ntp
      ntp-4.1.2-4.EL3.1
      ```

   c. 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
      ```

   d. Install the ccur-ntp update:

      ```
      # rpm -ivh ccur-ntp-4.2.0.a.20040617-4*.rpm
      ```

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

5.5 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 4.2 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 4.2” 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).

5.6 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.

5.6.1 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.

5.6.2 Installing the Software

1. With RedHawk Linux Version 4.2 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 4.2” 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.

5.7 Installing Additional RedHawk Products

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

5.8 Installing RedHawk Updates from CD

Discs labeled “Software Updates for RedHawk” include updates to RedHawk and its products that should be installed now.

1. Ensure that the system is in multi-user mode.
2. Locate the disc labeled “Software Updates for RedHawk” appropriate to your system’s architecture, 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. Execute the following commands to install the RedHawk updates included on the CD:

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

   When the installation is compete, the Main window for Concurrent's Network Update Utility (NUU) displays, listing product updates that complete your installation of RedHawk. Depending on the installation state of your Concurrent products, the NUU Main window may show software modules which are out of date.

   Click on the Select Update of All Out-of-Date button, then click on the Apply Actions button to install all listed product updates.

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. If there are more update discs, repeat these steps to install all updates.

8. Check for additional RedHawk updates available through the RedHawk Updates web site as outlined in the following section.

5.9 Installing RedHawk Updates from the Updates Web Site

1. Additional updates may have been issued that are not included on the "Software Updates for RedHawk" disc(s). They should be downloaded and installed now. Refer to "Software Updates" on page 53 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. Upgrade Procedures

Use the upgrade procedures provided here only under the following circumstances:

- upgrading from Version 2.x or 4.1 to Version 4.2

If generating or regenerating a system, upgrading a system currently running Version 1.x to Version 4.2, or converting a system from 32-bit to 64-bit, follow the steps in “Installation Procedures” on page 14.

Review the following points before upgrading the software:

- It is recommended that the installer first read through the entire procedure once before actually performing the upgrade.
- It is **strongly recommended** that installation and updates of software be performed in single-user mode to guard against the possibility of corrupting a running application.
- Install all available upgrades. No quality claims are made by Concurrent unless all appropriate version levels of all RedHawk Linux software are installed.
- The 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 upgrade, remove any floppy media that may have been left in the machine.
- The Red Hat kernel should be running when upgrading to RedHawk Linux Version 4.2.

6.1 Pre-Upgrade Procedure

Before upgrading a RedHawk Version 2.x or 4.1 system to Version 4.2, perform the following steps.

1. Ensure that the Red Hat SMP kernel is the default kernel by editing **`/etc/grub.conf`** and changing the `default=#` line to the Red Hat SMP kernel’s grub entry number. Note that grub entries are numbered starting from 0.

   In addition, for upgrades from Version 4.1 to 4.2, back up **`/etc/grub.conf`** to **`/etc/grub.conf.bck`** as a precaution because you will be creating a new boot loader configuration in section `6.2`.

2. Edit the **`/etc/X11/xorg.conf`** file (on 2.3 systems, this file is called **`/etc/X11/XF86Config`**) 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. Make a backup of the **/etc/X11/xorg.conf** file (**/etc/X11/XF86Config** on 2.3 systems). Between Red Hat Enterprise Linux 3 (RHEL3) and RHEL4, Red Hat changed from using the XFree86 X Server to the X.org X Server. It is possible that the new **/etc/X11/xorg.conf** file created by RHEL4 may be migrated incorrectly and you will want to compare the old and new files.

   Note that you cannot simply rename the original **XF86Config** file to **xorg.conf**. In order to migrate the original file you must:

   a. Change the name of the keyboard driver from "keyboard" to "kbd". The driver definition is located in the following code section:

   ```
   Section "InputDevice"
       ...
       Identifier "Keyboard0"
       ...
       Driver     "kbd"
       ...
   EndSection
   ```

   b. Ensure once again that there is no `Option "dri"` line in the **xorg.conf** file.

6.2 Upgrading to RHEL4

RedHawk Version 2.x was based on Red Hat Enterprise Linux 3. The system must first be upgraded to Red Hat Enterprise Linux 4.

1. Reboot the system with the disc labeled "Red Hat Enterprise Linux 4.0 Update 4 Disk #1" in the CD-ROM drive.
2. After a short delay, a screen containing the `boot:` prompt should appear. Press Enter to start the upgrade.
3. At the Media Check screen, select Skip (unless you wish to truly verify the installation media).
4. At the Welcome screen, no action is needed. Select Next.
5. At the Language Selection screen, choose U.S.English. Select Next.
6. At the Keyboard configuration screen, choose Default. Select Next.
7. At the Upgrade Examine screen, the default is to "Upgrade an Existing Installation". Select Next.
8. For a Version 2.3 to 4.2 upgrade:

   At the Upgrade Boot Loader Configuration screen, the default is to "Update boot loader configuration". Select Next.

   For a Version 4.1 to 4.2 upgrade:

   a. At the Upgrade Boot Loader Configuration screen, the default is to "Skip boot loader updating." Select "Create new boot loader configuration" instead. Select Next.

   b. In the "Boot Loader Configuration" screen, deselect "Configure advanced boot loader option." Select Next.

9. At the About to Upgrade screen, select Next. A Required Install Media popup is displayed regarding the CDs required. Select Continue.
10. The Installing Packages screen displays. Follow the instructions on the screen to install RHEL 4.0 Update 4 from the CDs. You will install from discs 1 through 5, then 1 again.
11. At the Installation Completed screen, remove any media from the drives, store appropriately, then select Reboot.
12. During reboot, select the Red Hat Enterprise Linux SMP kernel when prompted.
13. A Migrated Old Desktop information box is displayed. Select OK.
14. Continue with the steps in the following section.

6.3 Updating Red Hat Software

The disc labeled "Red Hat Enterprise Linux 4.0 with Update 4 Updates LXUPDATE Version 4.2" contains the latest updates released by Red Hat and validated against RedHawk Linux. These updates are important to the proper operation of RedHawk Linux.

To update the installed Red Hat software, perform the following steps.

1. If not already running **Red Hat Enterprise Linux**, boot the system with the Red Hat kernel.
2. 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`**.
3. Locate the disc labeled "Red Hat Enterprise Linux 4.0 with Update 4 Updates Version 4.2" appropriate to your system's architecture and insert it into the CD-ROM drive.
4. If the cdrom device does not mount automatically, execute the **`mount`** command with the mount point that is designated in **`/etc/fstab`**.

   NOTE: `/media/cdrom` is used in the examples in this section; e.g.:

   ```
   mount /media/cdrom
   ```

   If the cdrom entry is not in **`/etc/fstab`**, executing the following command will mount the cdrom device at **`/media/cdrom`**:

   ```
   mount /dev/cdrom /media/cdrom
   ```
5. Execute the following commands to install the update software:

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

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

   NOTE: This process may take some time to complete.
6. When the installation completes, execute the following commands:

   ```
   cd /
   umount /media/cdrom
   eject
   ```
7. Remove the disc from the CD-ROM drive and store.

8. **Perform this step ONLY if you are installing on a single processor system.** If installing on a multiprocessor system, skip to the next step.

   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.

9. Type “reboot” at the system prompt, then select the updated Red Hat Enterprise Linux kernel and verify that the updated kernel boots. Continue with the steps in the following section.

6.4 Upgrading to RedHawk Linux Version 4.2

To upgrade to RedHawk Linux Version 4.2, perform the following steps:

1. With the Red Hat Enterprise Linux kernel updated in the previous step running on the system, 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. Locate the disc labeled “RedHawk Linux OS Version 4.2” 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. Execute the following commands to install RedHawk Linux Version 4.2 upgrade software:

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

   NOTE: 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.

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. Locate the disc labeled “RedHawk Linux Documentation CD Version 4.2” appropriate to your system’s architecture and insert it into the CD-ROM drive.
8. If the cdrom device does not mount automatically, execute the **`mount`** command; for example:

   ```
   mount /media/cdrom
   ```

9. Execute the following commands to install RedHawk Linux documentation:

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

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

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

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

11. Remove the disc from the CD-ROM drive and store.

12. Exit single-user mode (Ctrl-D).

13. 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 **8.1.3** on page 47 for details.

14. Continue to the next step to upgrade FBS.

6.5 Upgrading to Frequency-Based Scheduler (FBS) Version 4.2

If you are using FBS, upgrade to FBS Version 4.2 at this time by performing the following steps. If you are not using FBS, proceed to "Upgrading to PCI-to-VME Bridge Version 4.2" on page 32.

1. With RedHawk Linux Version 4.2 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 4.2" 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. Execute the following commands to install the FBS software:

   ```
   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).

6.6 Upgrading to PCI-to-VME Bridge Version 4.2

If you are using the PCI-to-VME bridge supplied by Concurrent to connect your iHawk system with a VME system, upgrade the software to Version 4.2 at this time by performing the following steps. If you are not using this bridge, proceed to "Installing RedHawk Updates from CD" on page 32.

1. With RedHawk Linux Version 4.2 running, 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 4.2" 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. Execute the following commands to install the PCI-to-VME software:

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

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

5. Upon completion, 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).

6.7 Installing RedHawk Updates from CD

Discs labeled "Software Updates for RedHawk" include updates to RedHawk and its products that should be installed now.

1. Ensure that the system is in multi-user mode.

2. Locate the disc labeled "Software Updates for RedHawk" appropriate to your system's architecture, 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. Execute the following commands to install the RedHawk updates included on the CD:

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

   When the installation is compete, the Main window for Concurrent's Network Update Utility (NUU) displays, listing product updates that complete your installation of RedHawk. Depending on the installation state of your Concurrent products, the NUU Main window may show software modules which are out of date.

   Click on the Select Update of All Out-of-Date button, then click on the Apply Actions button to install all listed product updates.

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. If there are more update discs, repeat these steps to install all updates.

8. Check for additional RedHawk updates available through the RedHawk Updates web site as outlined in the following section.

6.8 Installing RedHawk Updates from the Updates Web Site

1. Additional updates may have been issued that are not included on the "Software Updates for RedHawk" disc(s). They should be downloaded and installed now. Refer to "Software Updates" on page 53 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.

6.9 Additional Requirements

If you are using custom drivers with a pre-defined RedHawk Linux kernel, it will be necessary to run **`ccur-config`** after upgrading to a new release of RedHawk in order to configure the kernel source tree to match the currently running kernel. Follow these steps:

1. Change directory (**`cd`**) to **`/usr/src/`***redhawk-release-directory.*

2. Run the **`ccur-config(8)`** tool as follows:

   ```
   ./ccur-config -n -c
   ```

   Wait for on-screen messages to complete.

If you are building additional modules out of the kernel source tree, it will be necessary to run **`ccur-config`** after upgrading to a new release of RedHawk in order to configure the drivers in the kernel. Follow these steps:

1. Change directory (**`cd`**) to **`/usr/src/`***redhawk-release-directory.*

2. Run the **`ccur-config(8)`** tool as follows:

   ```
   ./ccur-config -c
   ```

Refer to the man page or *RedHawk Linux User's Guide* for more details, if needed.

7.0. NVIDIA Graphics Configuration

The Unified NVIDIA® Linux Display Driver (version 1.0-9746) 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.

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 7.4.

7.1 The Unified and Legacy NVIDIA Display Drivers

The NVIDIA Linux Display Driver (version 1.0-9746) 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 1.0-9631) 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 **7.3**, “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 7.3, “Using the Legacy NVIDIA Driver”.

7.1.1 Supported NVIDIA GPUs

Table 7-1 lists the NVIDIA GPUs that are supported only in the unified (version 1.0-9746) NVIDIA Linux Display Driver.

Table 7-2 lists the NVIDIA GPUs that are supported only in the legacy (version 1.0-9631) 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 the section "Appendix A. Supported NVIDIA Graphics Chips" in the following README files:

- If the unified driver (version 1.0-9746) is installed on your system:

  `/usr/share/doc/ccur-nvidia-glx-1.0/README.txt`

- If the legacy driver (version 1.0-9631) is installed on your system using the steps described in section **7.3**:

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

Table 7-1 GPUs Supported Only in the NVIDIA 'unified' Driver

NVIDIA Chip Name	Device PCI ID
GeForce 6700 XL	0x0147
GeForce 6200 LE	0x0163
GeForce 8800 GTX	0x0191
GeForce 8800 GTS	0x0193
GeForce 7300 SE	0x01D3
GeForce 6200	0x0221
GeForce Go 6100	0x0247
GeForce 7900 GS	0x0292
GeForce 7950 GX2	0x0294
GeForce 7300 GT	0x0393

Table 7-2 GPUs Supported Only in the NVIDIA 'legacy' Driver

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

7.2.1 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.

In addition, you must ensure that your **`/etc/X11/xorg.conf`** configuration file does not attempt to make use of the NVIDIA video card(s) that are incompatible with your currently installed driver software. Failure to do so may lockup 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.

7.3 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.

7.3.1 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*-1.0-9631-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*-1.0-9631-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.

7.3.2 Building Additional Kernel Modules

If you boot with a kernel other than the one running when you installed the legacy driver in section **7.3.1**, 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*-1.0-9631-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.

7.3.3 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-1.0-*.rpm
   ```

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

   ```
   depmod
   ```

7.3.4 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
```

7.4 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.

There is only one high performance AGP slot in an iHawk graphics system, so the maximum number of displays that can run off AGP depends on the number of heads on the installed card. It should be possible to mix and match as many PCI and PCI Express (PCIe) cards as there are PCI and PCIe slots, and PCI cards with AGP. Systems designed with AGP are usually limited to two physical CPUs. If more than two physical processors are needed in your system, utilize PCI and PCIe adapters. To maximize throughput, place the PCI cards in PCI slots that are not shared by other devices.

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 **7.4.1** 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 **7.4.2** 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 **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 &**.

7.4.1 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.

7.4.2 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
```

7.5 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. If you will be using separate desktops and wish to use **`gnome`**, edit **`/etc/X11/gdm/gdm.conf`** and uncomment the following line:

```
1=Standard
```

Note that **`kde`** works without modifications.

4. 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
```

5. Startup X with **`init 5`**.

6. 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.

7. 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.

8.0. Additional Information

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

8.1 Installation/Configuration Issues

8.1.1 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 8-1 contains a list of configuration files that may be useful in configuring a newly installed RedHawk system to match an existing configuration.

Table 8-1 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.

8.1.1 File System Swap Size Guidelines

Table 8-2 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 8-2 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

8.1.3 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

8.1.4 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.

8.1.5 Creating a Driver Disk for Red Hat Installation

If your system contains disk drivers or devices that are not supported by the Red Hat installation program, you will need to use a driver diskette with the needed driver image file to install Red Hat Enterprise Linux. If this diskette is not supplied with your system, it can be created by copying the driver image file to a floppy.

Look for the file on the RedHawk Updates web site (**redhawk.ccur.com**). It will have the suffix "**.img**." Follow the instructions in the section "Software Updates" on page 53 to download the file. On any working system, copy this file to a formatted floppy disk.

To copy the file to disk on a Linux system, use the following command line as root, where `fd0` is the system's floppy drive.

dd if=*imagefilename* **of=/dev/fd0**

If working on a Windows system, use the DOS **rawrite** command. Note that the **rawrite** utility only accepts filenames having eight characters, a period, and a three-character filename extension. Therefore, you may need to rename this file to meet that convention before running **rawrite**.

In DOS, from the same directory containing the driver diskette image, type the following, where **d** is the cdrom drive letter and **a** is the floppy drive:

d:\dosutils\rawrite *imagefilename* **a:**

Use this diskette as you follow the instructions under "Installation Procedures."

8.2 Considerations

8.2.1 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-2.96-98** or later.

8.2.2 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 4.2 RedHawk Linux kernel on this type of system.

9.0. Known Issues

Special consideration should be given to the following areas.

Networking Issues

- There are sporadic problems with password files that are distributed to multiple systems via NIS. This can prevent valid users from logging into the system when their accounts are made known to that system via NIS.
- NFS mounts of large partitions will sporadically have problems in seeing all of the files that should be mounted in the partition.
- An incompatibility exists between various versions of the nfs-utils package and RedHawk Linux. Concurrent strongly discourages installing separate updates to this package.

Linux Test Project

When running the Linux Test Project test suite LTP-2002(*), the user must remove the **Fork07** test case prior to building the test suite to prevent a system hang on RedHawk systems with greater than 1GB of memory. This is a known bug in the official Linux kernel. For more information, refer to the Linux Test Project web site at **ltp.sourceforge.net**.

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.

Kernel Crash Dumps not Supported on RAID

Kernel crash dumps are not currently supported on Model HQ074 and HQ058 iHawk systems configured with RAID. When configured with SCSI disks, kernel crash dumps are supported.

lspci(8) Output

Under earlier versions of RedHawk Linux, output from **lspci(8)** displays the RCIM peripheral as "PLX Technology, Inc.: Unknown device." To correct this so that the RCIM is correctly identified, update **/usr/share/hwdata/pci.ids** with the version of this file supplied in the RedHawk kernel tree. Contact Concurrent Software Support (see page 53 for details) if you require assistance.

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**}

Performance Issue Caused by Enabled HPET Timer

Having the HPET timer enabled results in intermittent RTC interrupt errors and a corresponding increase in process dispatch latency of applications that wait on the RTC interrupt, as well as a general system slowdown. If a system boots in a very slow and jerky manner, it might be because of the HPET being enabled.

It is recommended that the HPET timer functionality be disabled in the BIOS. Refer to your hardware documentation to determine which BIOS setting is involved; it may not be labeled HPET Timer. For example, on an Opteron iHawk system using a Tyan Thunder K8W S2885 board, it is known as Multimedia Timer and is accessed in the BIOS under Advanced ->ACPI configuration->ACPI Advanced Configuration->Multimedia Timer.

Periodic Interrupts on IBM Blade System

It has been observed that the IBM HS20 Blade system experiences periodic SMI hardware interrupts. These interrupts last approximately 6 microseconds and appear to occur with a very regular frequency of nearly exactly one hour intervals.

Concurrent is working with IBM to determine the cause of this potential source of non-determinism, and it is hoped that the issue will be addressed in a future BIOS revision which can be installed at the customer site.Additional information and BIOS upgrade instructions will be announced and made available as soon as a resolution to the issue is available from IBM.

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.

Opteron NUMA Restriction

NUMA-enabled x86_64 kernels may not boot up on some systems where one or more NUMA nodes are configured without physical memory. Ensure that all NUMA nodes are configured with physical memory before attempting to boot a NUMA-enabled kernel.

Compatibility Issues for Ada programs

RedHawk 4.1 are based on Red Hat Enterprise 4.0. Binary and source compatibility issues exist for MAXAda programs between RHEL 3.0 and RHEL 4.0.

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 “watchdog,” 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 4.2 uses **glibc** 2.3.4).

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
```

10.0. Software Updates and Support

10.1 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/support_supportservices_CustomerAssistance_rt.aspx**.

10.2 Software Updates

Updates to Concurrent RedHawk software can be obtained via Concurrent's Updates web site, **http://redhawk.ccur.com**. This site can be accessed by clicking on the "Updates" icon on the desktop once the RedHawk Linux Documentation CD has been installed. There are three ways of installing product updates:

- Using the Network Update Utility (NUU) available from the Updates web site
- 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

10.2.1 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.

If you have not already downloaded NUU, access the Updates web site and click on Download NUU. 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 53.

10.2.2 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 "Updates" icon on the desktop. Accessing this web site displays the following screen:

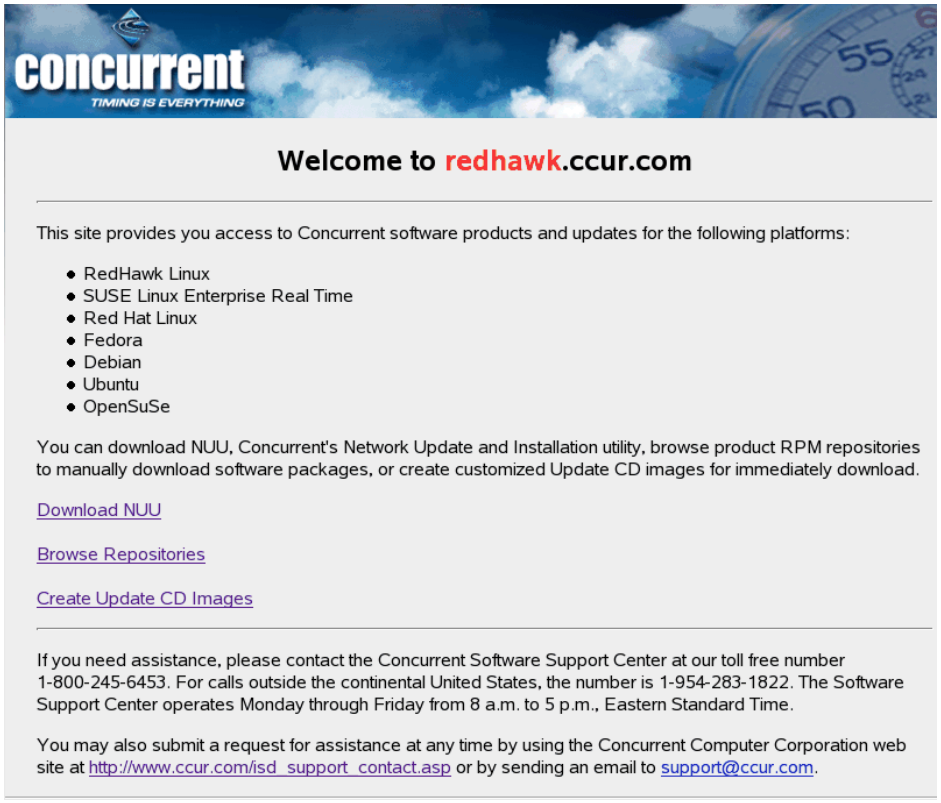

Clicking on the **Browse Repositories** link takes you to an authentication page.

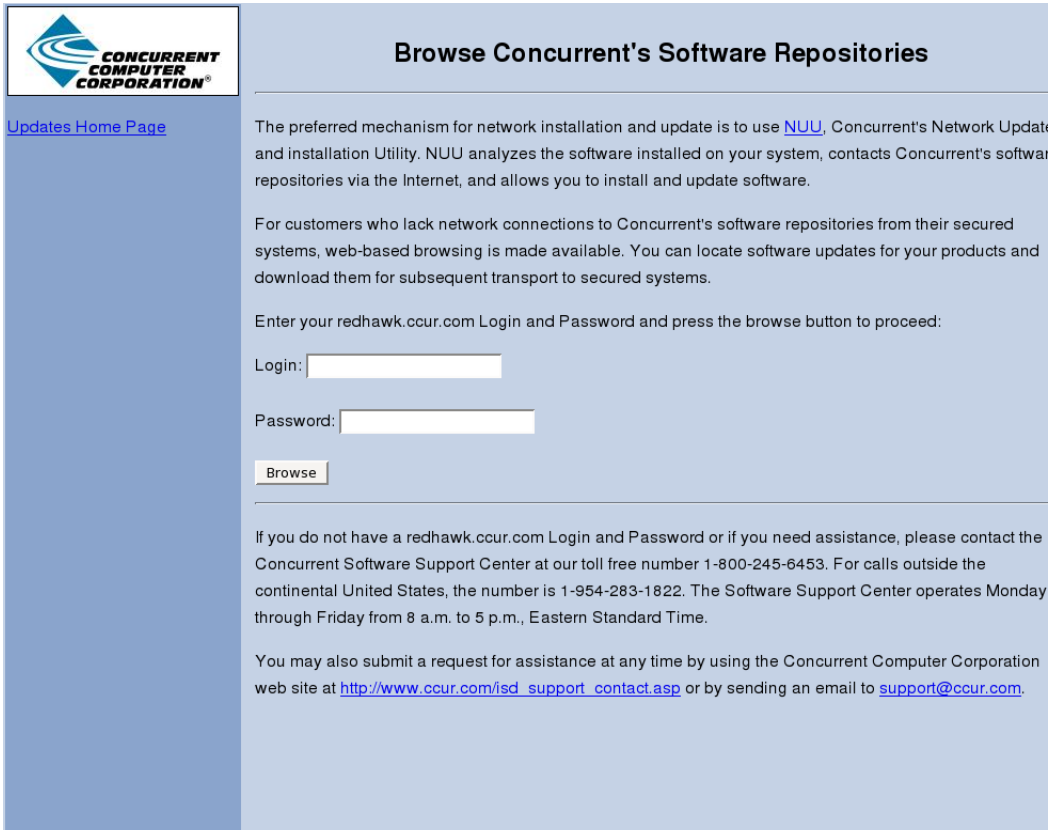

Enter your redhawk.ccur.com Login ID and Password and click the Browse button.

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

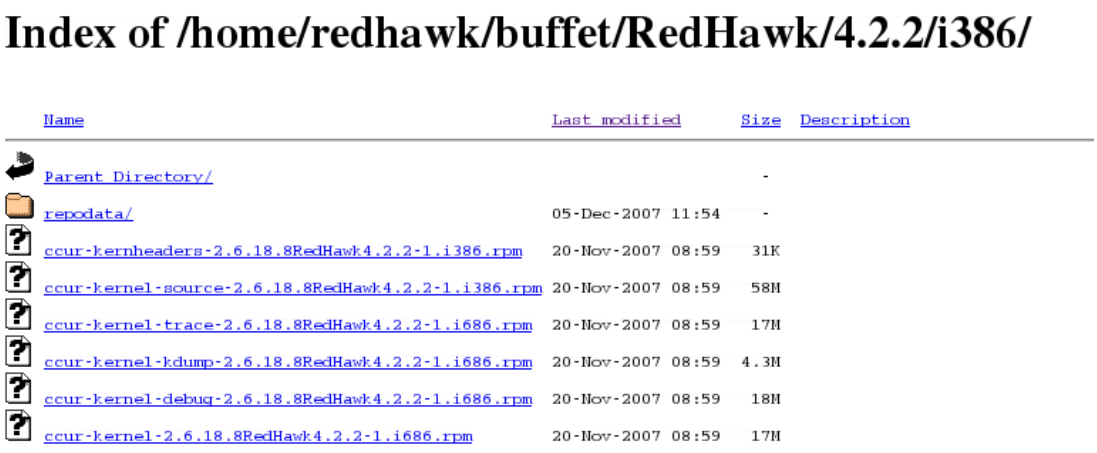

To locate the latest RPMs in the repository, click on the Last modified column header to sort by date. You may need to click twice to set the sort order to newest-to-oldest.

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.

10.2.3 Customized Update CDs

You can use the RedHawk Updates web site 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 "Updates" 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.

10.3 Documentation Updates

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

For the latest RedHawk FAQ, go to **http://www.ccur.com/isdfaq/RedHawk-FAQ.pdf**.