



Version 6.0 Release Notes

December 2011

0898003-6.0a



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

Duplication of this manual without the written consent of Concurrent Computer Corporation is prohibited. Any copy of this manual reproduced with permission must include the Concurrent Computer Corporation copyright notice.

Trademark Acknowledgments

Concurrent Computer Corporation and its logo are registered trademarks of Concurrent Computer Corporation. All other Concurrent product names are trademarks of Concurrent while all other product names are trademarks or registered trademarks of their respective owners. Linux® is used pursuant to a sublicense from the Linux Mark Institute.

Contents

1.0	Introduction	1
1.1	Product Description	1
1.2	Product Media	2
1.3	Version Numbers	3
1.4	System Updates	4
1.5	Syntax Notation	4
2.0	Documentation	5
2.1	RedHawk Linux Documentation	5
2.2	Vendor Documentation	6
3.0	Prerequisites	7
3.1	Software	7
3.2	Hardware	7
4.0	Changes in this Release	8
4.1	Enhancements in Version 6.0	8
4.1.1	Based on Red Hat Enterprise Linux 6.0	8
4.1.2	Based on kernel.org 2.6.36	8
4.1.3	Updated NVIDIA & CUDA	8
4.1.4	Updated ATI	8
4.1.5	More Than 64 CPUs	8
4.1.6	Performance Events	9
4.1.7	Multiple CPU Architecture Virtualization	9
4.1.8	Logical Volume Manager	9
4.1.9	Unified Extensible Firmware Interface	9
4.1.10	Page Locking Enhancements	9
4.1.11	New RCIM III Time Synchronization	9
4.1.12	RAM Disk Memory Shielding	9
4.1.13	New shield Command Feature	10
4.1.14	New run Command Features	10
4.2	Modifications in Version 6.0	10
4.2.1	Installation Method	10
4.2.2	Initial RAM Filesystems	10
4.2.3	NUMA Library Changes	10
4.2.4	Ksoftirqd Default Priority Changed	11
4.2.5	Workqueue Daemon Default Priority Changed	11
4.2.6	Swap Space	11
5.0	Installation Procedures	12
5.1	Installing Red Hat Enterprise Linux Software	12
5.1.1	Installing RHEL	13
5.1.2	Installing Red Hat Updates	22
5.2	Installing RedHawk Linux	25
5.3	Installing the RCIM	28
5.3.1	Quick Hardware Installation Checklist	28
5.3.2	Verifying RCIM Operation	28
5.3.3	Installing NTP Updates for GPS Module	29
5.4	Installing Frequency-Based Scheduler (FBS) Software	29
5.5	Installing the PCI-to-VME Bridge	30
5.5.1	Installing the Hardware	30
5.5.2	Installing the Software	31

5.6	Installing the PCIe-to-VME Bridge	32
5.6.1	Installing the Hardware	32
5.6.2	Installing the Software	32
5.6.3	Configuring the PCIe-to-VME Windows	33
5.7	Installing Additional RedHawk Products	34
5.8	Installing RedHawk Updates from Disc	34
5.9	Installing RedHawk Updates from the Updates Web Site	35
6.0	NVIDIA Graphics Configuration	36
6.1	Supported NVIDIA GPUs	36
6.2	Multiple NVIDIA Video Card Considerations	37
6.3	Display Configurations	37
6.3.1	Single	37
6.3.2	Xinerama	37
6.3.3	MultiDesktop	38
6.3.4	Twinview	38
6.3.5	Twinview-Xinerama	39
6.3.6	Twinview-MultiDesktop	39
6.4	Using nvidia-xconfig	39
6.5	Using nvidia-settings	40
6.6	Modified xorg.conf file generated by nvidia-xconfig	47
6.7	Troubleshooting Problems	48
7.0	ATI Graphics Configuration	51
7.1	ATI Graphics Configuration	51
7.2	Supported ATI Cards	51
7.3	ATI and NVIDIA Configuration Considerations	51
7.4	Display Configurations	51
7.4.1	Single Head	51
7.4.2	Dual Head	52
7.5	MultiView	52
7.6	Using Catalyst Control Center	52
7.7	Known Issues	53
7.8	Supporting Other ATI Graphics Controllers	53
8.0	Additional Information	55
8.1	Installation/Configuration Issues	55
8.1.1	File System Backup Recommendations	55
8.1.2	File System Swap Size Guidelines	56
8.1.3	RCIM Connection Modes	56
8.1.4	How to Use Linux Rescue	56
8.1.5	How to Install RHEL Remotely Over VNC	58
8.1.6	If the Installation GUI Fails to Start	60
8.2	Considerations	61
8.2.1	Compiler Requirements	61
8.2.2	Hyper-thread Activation	61
9.0	Known Issues	62
10.0	Software Updates and Support	65
10.1	Direct Software Support	65
10.2	Software Updates	65
10.2.1	Updating via NUU	65
10.2.2	Installing Manually Downloaded RPMs	66
10.2.3	Customized Update Discs	68
10.3	Documentation Updates	68

1.0. Introduction

This document provides release information and installation instructions for Concurrent Computer Corporation's RedHawk™ Linux® operating system Version 6.0.

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® 6.0 distribution. The RedHawk installation media provide additional real-time kernels and libraries for accessing RedHawk specific kernel features. Optionally, the NightStar™ RT development tool set is available for developing real-time applications, and the Frequency-Based Scheduler can be used to schedule cyclic applications driven from a real-time clock or other external interrupt source. The Performance Monitor module allows you to get detailed CPU utilization of each process on such a scheduler.

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 44 years experience developing real-time operating systems.

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

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

RedHawk Linux matches the level of POSIX® conformance of other Linux distributions based on the 2.6 series of kernels. Linux on the Intel® x86 architecture has defined a defacto binary standard of its own which allows shrink-wrapped applications that are designed to run on the Linux/Intel x86 platform to run on Concurrent's iHawk platform.

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

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

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

1.2 Product Media

The tables below list the optical media discs that are delivered with RedHawk Linux Version 6.0 for each Concurrent system.

WARNING

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

32-Bit Product Media

Processor Type:	Intel® Xeon™ i386 and EM64T 32-bit
Model:	iHawk 860 Series [HQ685, 680, 665, 660, 645, 460, 295, 290, 285, 280, 265, 260, 195, 185, 175, 160, 075, 074, 069, 067, 065, 060, 058, 053, 049, 047 045, 036, 034, T74, T54, R54, R90, 000] iHawk 880 Series [HQHS2] ImaGen [HQ0G1, HR0G1]
Operating System Discs:	
Red Hat Enterprise Linux 6.0 (32-bit) (1 DVD) Red Hat Enterprise Linux 6.0 Source (32-bit) (2 DVDs) Red Hat Enterprise Linux 6.0 Updates (32-bit) (1 DVD) [LXUPDATE WA-IA32 Version 6.0] RedHawk Linux OS (32-bit) (1 DVD) [WA-IA32 Version 6.0] Software Updates for RedHawk (1 DV D or 1 CD, if present)	
Optional Product Discs:	
RedHawk Linux Frequency-Based Scheduler (32-bit) [WU1210-JL Version 6.0] NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version x.x (x.x=latest version)] RedHawk Linux PCI-to-VME Bridge Software Library for i386 & EM64T Systems [WU-HS630-LIB (WCS-SBS-620)] Version 6.0 RedHawk Linux PCIe-to-VME Bridge Software Library for i386 & EM64T Systems [WU-HS630-C2243] Version 6.0 RedHawk Linux Cluster Manager (32-bit) [WA9017-L] Version 6.0	

64-Bit Product Media

Processor Type:	AMD Opteron™ 64-bit and EM64T x86_64 64-bit
Model:	iHawk 860 Series [HQ685, 680, 285, 280, 185, 069, 067, 049, 047, T74, T54, R54, R90, 000] iHawk 870 Series [HR844, 430, 221, 210, 00W, 00T] iHawk 880 Series [HQHS2] ImaGen [HQ0G1, HR0G1]
Operating System Discs:	
Red Hat Enterprise Linux 6.0 (64-bit) (1 DVD) Red Hat Enterprise Linux 6.0 Source (64-bit) (2 DVDs) Red Hat Enterprise Linux 6.0 Updates (64-bit) (1 DVD) [LXUPDATE WA-AMD64 Version 6.0] RedHawk Linux OS (64-bit) (1 DVD) [WA-EM64T-AMD64 Version 6.0] Software Updates for RedHawk (1 DVD or 1 CD, if present)	
Optional Discs:	
RedHawk Linux Frequency-Based Scheduler for x86_64 Systems [WU1210-JA Version 6.0] NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version x.x (x.x=latest version)] RedHawk Linux PCI-to-VME Bridge Software Library for x86_64 Systems [WU-HS630-LIB (WCS-SBS-620)] Version 6.0 RedHawk Linux PCIe-to-VME Bridge Software Library for x86_64 Systems [WU-HS630-C2243] Version 6.0 RedHawk Linux Cluster Manager (64-bit) [WA9017-L] Version 6.0	

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, 6.0 is the initial release of major version 6, and 6.1 is a subsequent release to version 6.0. Both 6.0 and 6.1 constitute complete product releases and do not require a previous release to be currently installed on the system; both are available on optical media 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 6.0.1 kernel update can only be installed on a system currently running RedHawk Linux version 6.0.

1.4 System Updates

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

NOTE

Concurrent does not recommend downloading RHEL updates.

The RedHawk Linux kernel replaces the standard RHEL kernel and is likely to work with any version of the RHEL 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 65.

1.5 Syntax Notation

The following notation is used throughout this document:

<i>italic</i>	Books, reference cards, and items that the user must specify appear in <i>italic</i> type. Special terms may also appear in <i>italic</i> .
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 <i>blue</i> type will launch your web browser and display the web site. Clicking on publication names and numbers in <i>red</i> type will display the corresponding manual PDF, if the file is available and accessible.

2.0. Documentation

2.1 RedHawk Linux Documentation

Table 2-1 lists the documentation provided with RedHawk Version 6.0 as PDF files. After installing RedHawk Linux documentation, the files are available for viewing:

- by clicking on the “Documents” icon on the desktop
- from Concurrent’s documentation web site at <http://redhawk.ccur.com/docs>

Table 2-1 RedHawk Linux Documentation

RedHawk Linux Operating System Documentation	Document Number
<i>RedHawk Linux Version 6.0 Release Notes</i>	0898003-6.0
<i>RedHawk Linux User’s Guide</i>	0898004-630
<i>Real-Time Clock & Interrupt Module (RCIM) User’s Guide</i>	0898007-400
<i>RedHawk Linux FAQ</i>	N/A
Optional RedHawk Product Guides	
<i>RedHawk Linux Frequency-Based Scheduler (FBS) User’s Guide</i>	0898005-340
Optional RedHawk Product Documentation Sets	
<i>RedHawk Linux Cluster Manager</i>	(not available)
<i>SBS PCI-to-VME Bridge Series 600 Hardware Manual</i>	85851150
<i>SBS PCI-to-VME Bridge Series 800 Hardware Manual</i>	85913163
<i>SBS PCI-to-VME Bridge Software Manual</i>	85221990

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.

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://redhawk.ccur.com/docs/root/1redhawk/FAQ/RedHawk-FAQ.pdf>.

Printed copies can be ordered by contacting the Concurrent Software Support Center. Refer to “Direct Software Support” on page 65 for details.

2.2 Vendor Documentation

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

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

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

3.0. Prerequisites

3.1 Software

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

- Red Hat Enterprise Linux 6.0 with updates
- RedHawk Linux operating system (with optional updates)
- NightStar RT for RedHawk tools (optional)
- RedHawk Linux Frequency-Based Scheduler (FBS) (optional)
- RedHawk Linux PCI-to-VME bridge software (optional)
- RedHawk Linux PCIe-to-VME bridge software (optional)
- RedHawk Linux Cluster Manager (optional)

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

NOTE

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

3.2 Hardware

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

4.0. Changes in this Release

This section describes enhancements and other changes introduced in Version 6.0 (and any subsequent kernel updates if applicable).

4.1 Enhancements in Version 6.0

4.1.1 Based on Red Hat Enterprise Linux 6.0

RedHawk 6.0 uses Red Hat Enterprise Linux 6.0 as its base. RHEL 6.0 is the world-wide standard enterprise-class Linux distribution that provides a rock-solid foundation.

4.1.2 Based on kernel.org 2.6.36

RedHawk 6.0 uses the `kernel.org` Linux 2.6.36 kernel as the basis for its core kernel. This version includes many exciting enhancements. For more information examine the kernel change summaries at the <http://kernelnewbies.org> website.

4.1.3 Updated NVIDIA & CUDA

RedHawk 6.0 includes the NVIDIA Linux Display Driver version 270.40. Legacy driver support is no longer included but is available for download from the NVIDIA website. Refer to "NVIDIA Graphics Configuration" on page 36 for details.

On 64-bit systems RedHawk 6.0 includes pre-installed NVIDIA CUDA 4.0 libraries and header files. In addition, example code is provided. See `/usr/share/doc/nvidia/cuda/sdk` for more information.

Note that when using the NVIDIA Linux Display Driver on a 32-bit system it may be necessary to add `vmalloc=192M` to the kernel's grub command line in the system's `/etc/grub.conf` file in order for graphics to start properly.

4.1.4 Updated ATI

RedHawk 6.0 includes the ATI Linux FirePro Display Driver version 8.773.1.1. This driver includes support for FirePro, FireStream and FireGL graphics cards.

Note that ATI Stream support is no longer available as a standard product in this release.

Note that when using the ATI Linux FirePro Display Driver on a 32-bit system it may be necessary to add `vmalloc=192M` to the kernel's grub command line in the system's `/etc/grub.conf` file in order for graphics to start properly.

4.1.5 More Than 64 CPUs

RedHawk 6.0 is the first release of the new Kernel Trace support that allows more than 64 CPUs to be traced. This update allows RedHawk to fully support as many as 1024 CPUs on a single system.

4.1.6 Performance Events

RedHawk 6.0 provides kernel support for the new “Performance Events” subsystem. This subsystem provides a generic mechanism for monitoring key hardware and software events that occur during a specified range in time. Tools and documentation for Performance Events are included with the RedHawk 6.0 product. See the **perf(1)** and **ccur_perf_event_open(3)** man pages for more information.

RedHawk 6.0 has also instrumented the Performance Events subsystem with a new kernel tracepoint that is useful with the NightTrace tool suite. See the NightTrace kernel trace documentation for more information.

4.1.7 Multiple CPU Architecture Virtualization

RedHawk 6.0 includes the most recent version of QEMU and provides virtualization support for multiple CPU architectures, including: i386, x86_64, ARM, PPC, and MIPS.

4.1.8 Logical Volume Manager

RedHawk 6.0 now supports the Logical Volume Manager (LVM). LVM allows allocation of space on mass-storage devices that is much more flexible than conventional partitioning schemes. Multiple mass-storage devices may be used to create a single logical volume.

4.1.9 Unified Extensible Firmware Interface

RedHawk 6.0 now supports the Unified Extensible Firmware Interface (UEFI) specification. UEFI is a replacement for the much older BIOS firmware that is present in previous iHawk system motherboards. UEFI provides many enhancements including the ability to boot from disks larger than 2TB.

4.1.10 Page Locking Enhancements

RedHawk 6.0 can now optionally send a SIGBUS signal to a task when a page fault occurs. This can be used to determine the cause of unexpected page faults within an application.

RedHawk 6.0 has also extended the **mlockall** system service to support memory locking pages in another process. This feature can dramatically improve the real-time performance of an application without requiring any modifications to the application’s source code.

See the **mlockall(2)** and **mlockall_pid(2)** man pages for more information.

4.1.11 New RCIM III Time Synchronization

RedHawk 6.0 can now quickly and accurately synchornize the system time on one or more RCIM III slave systems to an RCIM III master system. This time synchronization mechanism utilizes the low-latency RCIM cables connecting the machines, and therefore no network connection between the machines is used or required. See the **rcimdate(1)** man page for more information.

4.1.12 RAM Disk Memory Shielding

RedHawk 6.0 now provides a mechanism for ensuring that all memory associated with a RAM disk is contained within a specified set of NUMA nodes. For more information refer to the `ramdisk-numa.txt` file that is installed in the following directory:

```
/usr/src/linux-2.6.36.*RedHawk6.0/Documentation/blockdev
```

4.1.13 New shield Command Feature

The **shield** command has been enhanced to support a short-hand notation for specifying all CPUs in a NUMA node. In this new alternate format, a leading **n** or **c** character specifies whether the value is a NUMA node ID or a CPU ID, respectively. See the **shield(1)** man page for more information.

4.1.14 New run Command Features

The **run** command has been enhanced to support the `SCHED_BATCH` and `SCHED_IDLE` scheduling policies and the `SCHED_RESET_ON_FORK` scheduling policy attribute.

The **run** command has also been enhanced to support a short-hand notation for specifying all CPUs in a NUMA node. In this new alternate format, a leading **n** or **c** character indicates whether the value is a NUMA node ID or a CPU ID, respectively.

The ability to externally memory lock and unlock other process's address space has also been added. Lastly, a new memory policy `cpulist` option has been added which displays the CPUs belonging to each NUMA node in a CPU list format instead of a bitmask format.

See the **run(1)** man page and the *RedHawk Linux User's Guide* for more information.

4.2 Modifications in Version 6.0

4.2.1 Installation Method

The method of installing RedHawk has been significantly updated and improved. It is no longer necessary to ensure that all base distribution packages are installed before installing from the RedHawk media. This allows existing systems to be converted to RedHawk systems with minimal effort. See "Installation Procedures" on page 12 for more information.

4.2.2 Initial RAM Filesystems

RedHawk 6.0 now requires the use of an initial RAM filesystem (*initramfs*) to boot the kernel successfully. The use of an *initramfs* allows many new features to be supported by RedHawk including the Logical Volume Manager. See the **initrd(4)** man page for more information.

By default, the RedHawk pre-built kernels and also user-created custom kernels have *initramfs* kernel modules stripped of all debugging symbols upon installation. If this is not desired for a custom kernel, install the modules of the custom kernel using the following command:

```
make INSTALL_MOD_STRIP=0 modules_install
```

4.2.3 NUMA Library Changes

The `libnuma` library version 1 included with RedHawk 5.4 has been updated to `libnuma` version 2 in RedHawk 6.0.

Unfortunately, many of the NUMA library function calls have been modified or replaced with new function calls, and the `nodemask_t` pointers used as function call arguments have been replaced in most instances with bitmask structure pointer arguments.

Binaries that were compiled for `libnuma` version 1 do not need to be re-compiled to run with `libnuma` version 2. Source code previously written for `libnuma` version 1 may be re-compiled without change by using the `-DNUMA_VERSION1_COMPATIBILITY` compilation flag.

Please refer to the **numa (3)** manual page for more details.

4.2.4 Ksoftirqd Default Priority Changed

The default priority of the **ksoftirqd** kernel daemon has changed from 95 to 98 for RedHawk 6.0. The actual priority of **ksoftirqd** can be changed after the system has booted if desired. In addition, the default priority can be permanently changed by adjusting the `CONFIG_SOFTIRQ_PRI` kernel tunable and building a custom kernel.

4.2.5 Workqueue Daemon Default Priority Changed

The default policy and priority of the various per-CPU and global kernel workqueue daemons has changed from `SCHED_OTHER` to `SCHED_RR/priority 5`. This is to accommodate real-time applications that need to have threads running above and/or below these daemons. For example, a thread doing disk I/O should run below the workqueue daemons in order to avoid potential priority inversion issues with the new workqueue components of the SCSI and ATA disk drivers.

If desired this priority can be changed back to `SCHED_OTHER` by setting the `CONFIG_WORK_PRI` tunable to 0 in a custom kernel.

4.2.6 Swap Space

Previous versions of RedHawk recommended swap space partition sizes that did not scale well to very large memory systems. The swap space partition size guidelines have now been revised to accommodate a much wider variety of possible memory configurations.

5.0. Installation Procedures

Use the installation procedures provided here to generate or regenerate a system.

NOTE: Upgrades from earlier versions of RedHawk to RedHawk 6.0 are *not* supported; only a full installation of RedHawk 6.0 is supported.

Review the following points before installing the software:

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

Before you start:

- **Main Memory Size** — You need to know the system’s main memory size in order to properly set up the file system swap size. If you are not sure of the main memory size, access the system’s BIOS to ascertain this information.
- **Network Addresses** — If you do not plan to configure networking via DHCP, you will need to have the following information for your system before beginning the installation:
 - IP address
 - gateway address
 - netmask 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.
 - Ensure that the drive to be used for installations and updates is bootable and that the boot order is correctly set in the BIOS so that you will boot to the installer from a disc.

5.1 Installing Red Hat Enterprise Linux Software

The version of Red Hat Enterprise Linux (RHEL) that is appropriate to your system’s architecture must be installed on the system first by *specifically completing the steps below*.

Installation of RHEL is accomplished in two parts:

- Installing RHEL
- Installing Red Hat Updates

If you already have a RHEL system at the proper version level you can skip directly to "Installing Red Hat Updates" on page 22.

5.1.1 Installing RHEL

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

1. Power on the system, if needed, to open the DVD drawer.
2. Insert RHEL Installation Disc in the drive and power down the system.
3. Power on the system with the installation disc installed to boot from the disc. After a short delay, the following screen should appear.



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

NOTES

If the installation GUI does not launch:

If you wish to use GUI mode but it does not launch, refer to section **8.1.6** on page 60 for some boot options that may cause the GUI to launch. Some NVIDIA cards require specifying the display resolution before the install GUI will start. If successful, continue to follow the installation directions below.

In addition to GUI mode, RHEL can be installed in text mode or by using VNC.

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 in the steps 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.

If you wish to use VNC for the installation instead of the GUI or text modes, refer to section **8.1.5** on page 58 for full instructions.

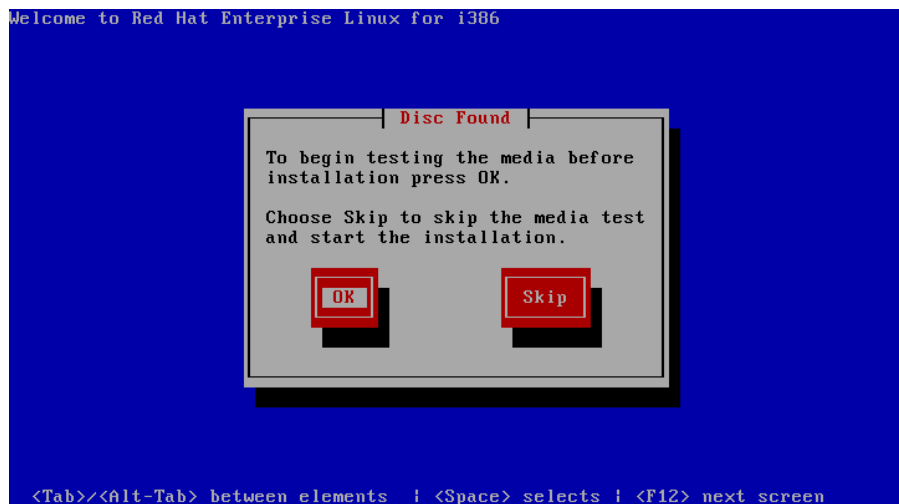
If your system contains disk drives or devices that are not supported by the RHEL installation program, you will need to obtain the necessary driver(s) in order to install Red Hat Enterprise Linux. A diskette containing the driver may be included with your system or it may be available on the RHEL installation media. If not available, contact the driver manufacturer for an updated driver or call Concurrent Support (see "Direct Software Support" on page 65).

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

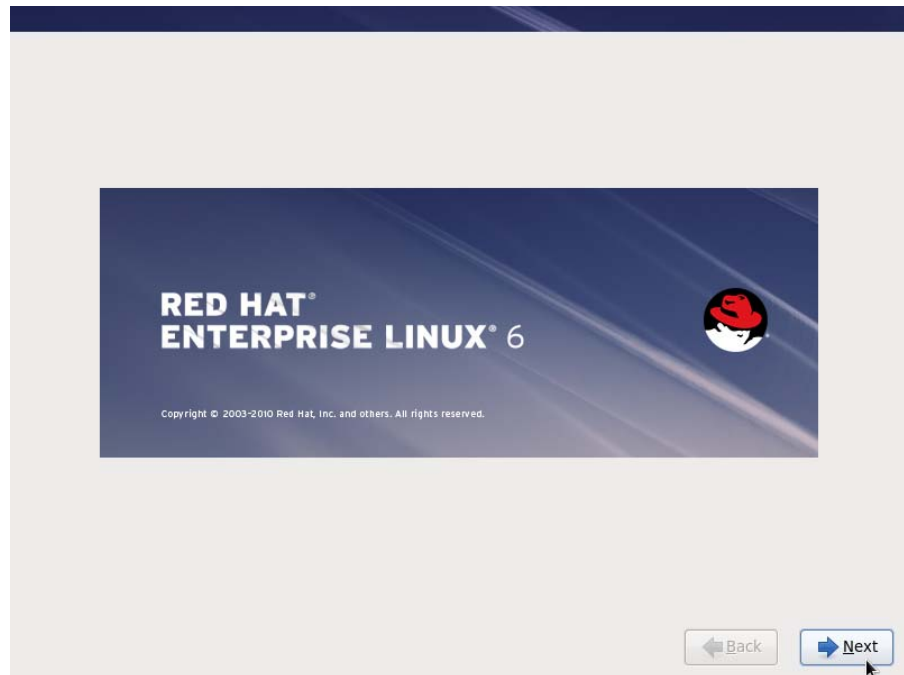
If your hard drive is not recognized and no driver diskette is supplied, a screen appears stating that no hard drives are found, would you like to select drivers. The installation disc contains a number of device drivers. Select **Yes**, then **Add Device** to display a list of drivers. Select the desired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

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

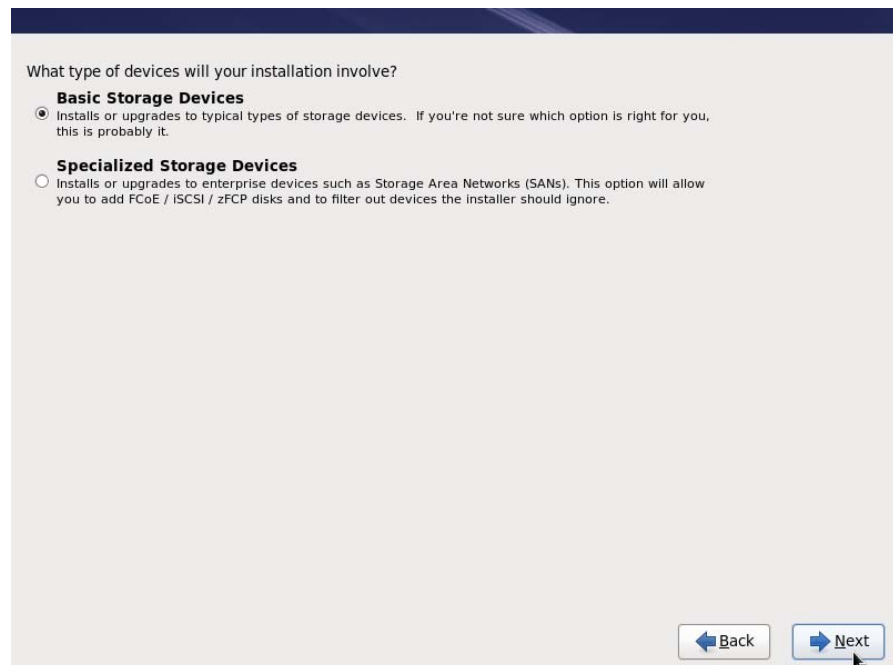
NOTE: The Source DVD media cannot be checked using this procedure because RHEL does not include the required checksum in the DVD content.



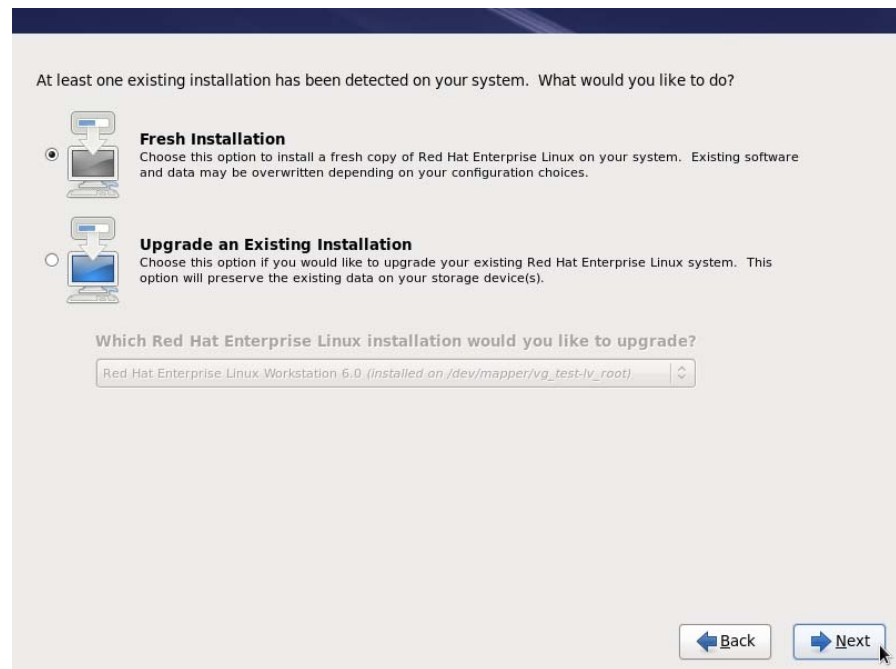
6. At the Welcome screen, no action is needed. Press Enter to proceed.



7. Select your desired language and keyboard layout on the next two screens, pressing **Next** at each to proceed.
8. At the next screen you are prompted to choose what type of storage devices your installation will be using. For almost all cases, select **Basic Storage Devices** and click **Next** to continue. Configuration of **Specialized Storage Devices** is beyond the scope of these instructions.



9. If you have previously installed RHEL 6.0 on the system, the dialog as shown below will appear. RedHawk 6.0 does not support upgrades so you should choose the **Fresh Installation** option and click **Next**.



If you have not previously installed RHEL 6.0 you will not see this dialog and you can skip to the next step.

Note that if you are using a new disk you will be asked to confirm initialization of the disk before you can continue with the next steps.

10. In the next screens you are prompted to enter a hostname, optionally configure networking, choose a time zone, and enter a root password.

Note that by default RHEL 6.0 does not configure networking to start at boot time; to change this configure networking and enable selected interfaces to **Connect automatically** at boot time.

Enter the desired information on each screen and click **Next**.

11. The next screen requires you to choose the type of installation you wish to perform. It is recommended that you choose **Use All Space** unless your system administrator requires a custom layout. Custom layouts are beyond the scope of this document.

Click **Next** to continue.

Which type of installation would you like?

Use All Space
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.

Replace Existing Linux System(s)
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.

Shrink Current System
Shrinks existing partitions to create free space for the default layout.

Use Free Space
Retains your current data and partitions and uses only the unpartitioned space on the selected device(s), assuming you have enough free space available.

Create Custom Layout
Manually create your own custom layout on the selected device(s) using our partitioning tool.

Encrypt system
 Review and modify partitioning layout

12. A confirmation dialog will now appear to convey that any data currently on the hard drive will be lost. If you wish to continue with the installation click **Write changes to disk** to continue. Disk partitions will now be created and formatted to prepare for the installation.
13. Next you must select the software package set to install. Choose **Software Development Workstation** and then click **Next** to start package installation. Note that more RHEL packages will be updated and installed during a separate future installation step.

The default installation of Red Hat Enterprise Linux is a desktop install. You can optionally select a different set of software now.

Desktop
 Minimal Desktop
 Web Development Workstation
 Software Development Workstation
 Minimal

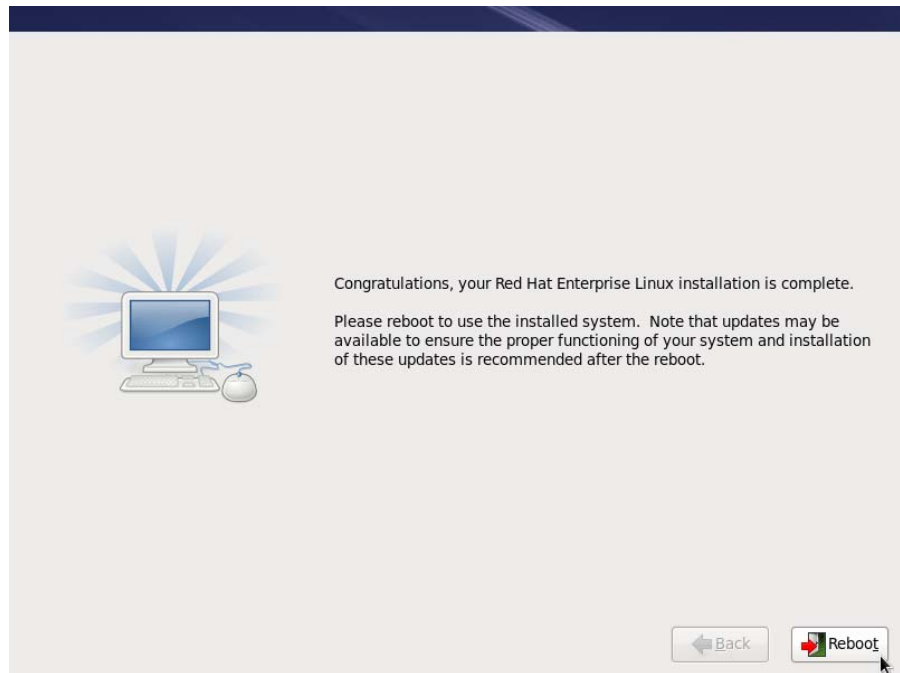
Please select any additional repositories that you want to use for software installation.

Red Hat Enterprise Linux

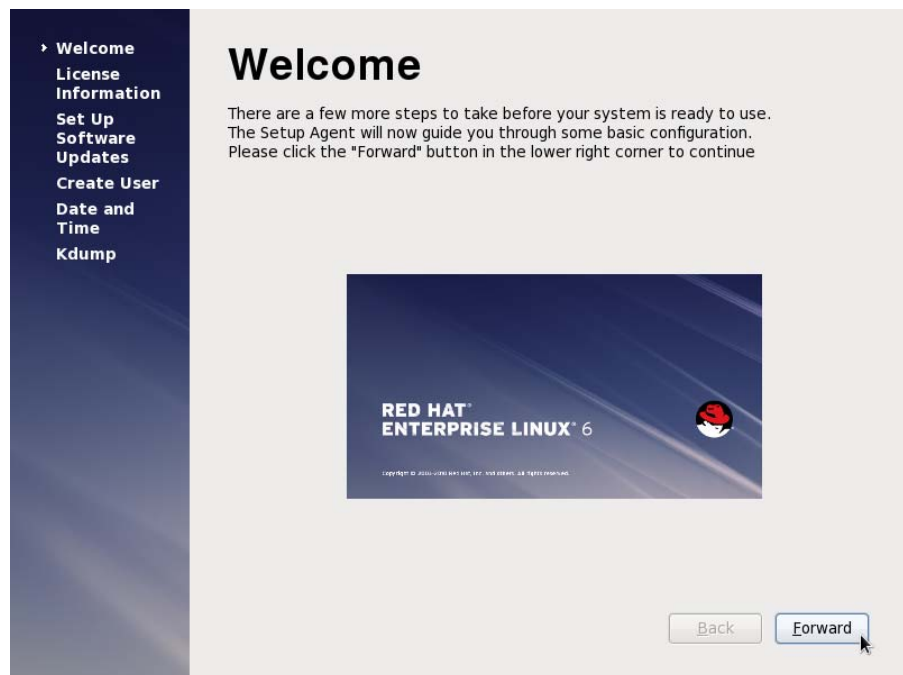
You can further customize the software selection now, or after install via the software management application.

Customize later Customize now

14. After all packages have been installed you will see the following screen. Remove the DVD install media and press **Reboot** to continue with the installation.



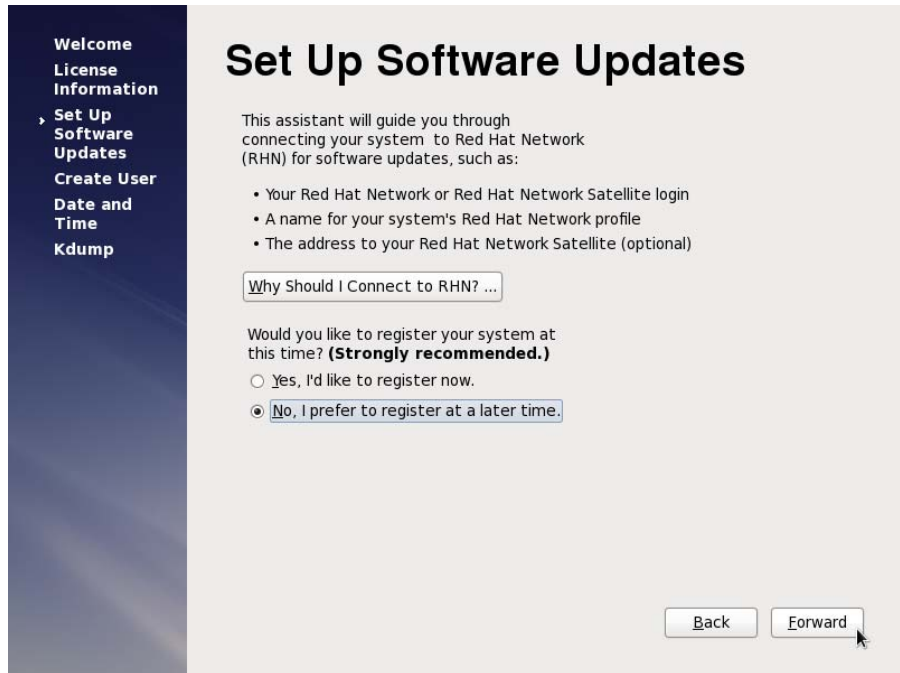
15. After the system reboots you are presented with the Welcome screen. Click **Forward** to proceed.



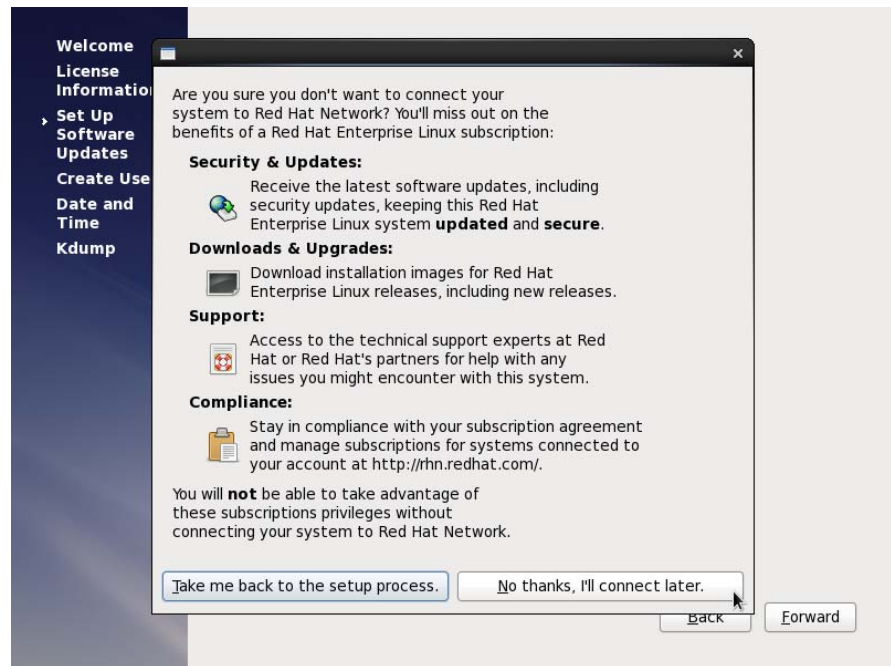
16. The License Agreement screen appears. Click **Yes, I agree to the License Agreement**, and then click **Forward**.

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

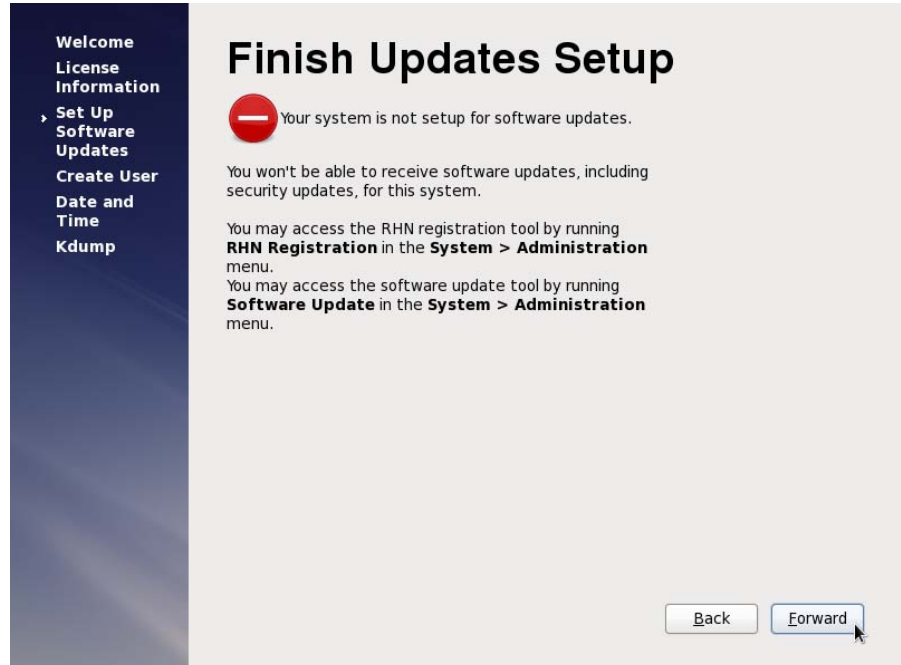
If you still wish to register with RHEL you may do so after all the software installation steps are completed. Click **Forward** to continue.



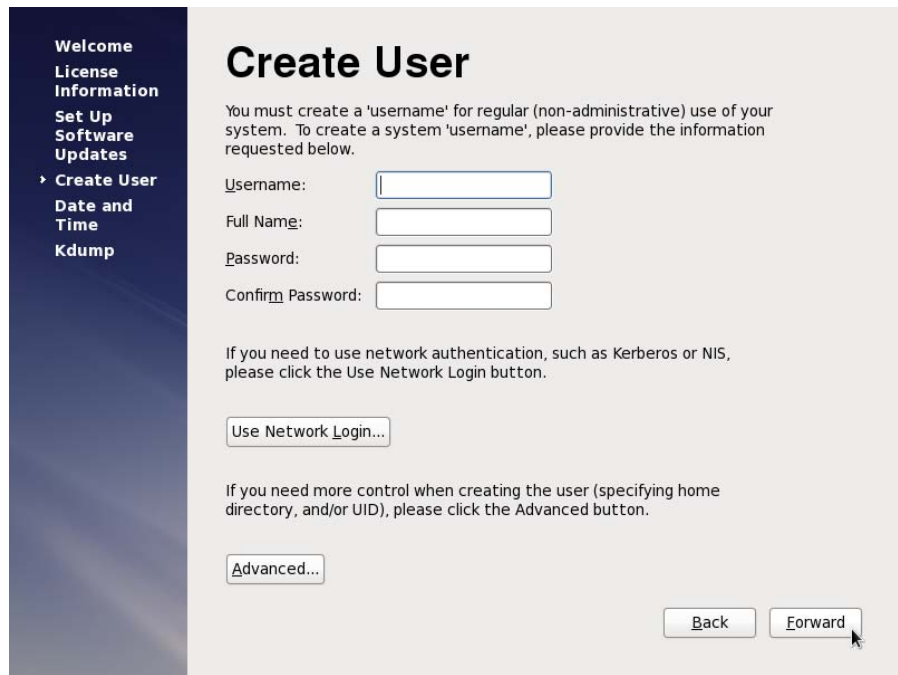
18. A pop-up window will now display asking for confirmation of the software updates choice. Click **No thanks, I'll connect later** to continue.



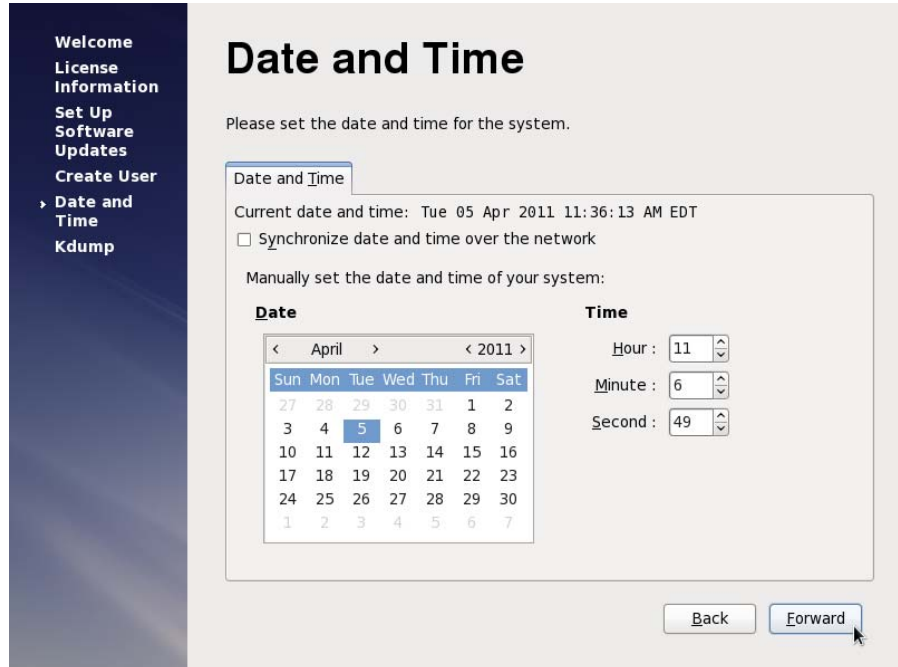
19. The next screen simply reminds you that you have chosen to defer software update registration. Click Forward to continue.



20. The next screen requires you to create a non-root user. Fill in the name and password fields as appropriate. Click Forward to continue.



- On the next screen set the date and time and select time synchronization via the Network Time Protocol if desired. Click **Forward** to continue.



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

Note that instead of a checkbox you may simply be presented with a dialog

stating there is insufficient memory in the system to configure kdump. This warning can be safely ignored.

Click **Forward** and the login screen will appear.



Note that you may also be presented with a dialog asking you to reboot the system in order to “reallocate system memory” for kdump. This dialog can be skipped as there is no need to reboot the system at this time.

After the login screen appears, preliminary installation of RHEL is complete. Proceed to the next section to complete the RHEL installation and update.

5.1.2 Installing Red Hat Updates

Install the disc labeled RHEL 6.0 Updates. This contains the latest updates released by Red Hat which have been validated by Concurrent on RedHawk Linux. These updates are important to the proper operation of RedHawk Linux.

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

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

NOTES

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

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

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

If these files are not present, you should examine the boot messages in `/var/log/messages` and look for your optical media device and find the associated `/dev` file.

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

```
[ -d /media/dvd ] || mkdir -p /media/dvd
mount device-name /media/dvd -o ro
```

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

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

```
/media/dvd
/media/dvdrw
```

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

```
/dev/hda /media/dvd auto pamconsole,exec,noauto,managed 0 0
/dev/hdc /media/dvdrw auto pamconsole,exec,noauto,managed 0 0
```

- Execute the following commands to install the update software, substituting the actual mount point used for your optical media disc device for `/media/dvd`, if necessary:

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

NOTE

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

Run the `mount` command (without arguments) and check the options. If `noexec` is present, remount the optical media using:

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

where *mount-point* is the directory name in the output of the previous `mount` command for that optical media. For example:

```
> mount | fgrep noexec
/dev/scd0 on /media/RHEL6.0-Updates type iso9660 (ro,noexec)

> mount -o remount,exec /media/RHEL6.0-Updates
```

- The installation script will now ask if you wish to install the full set of updated RHEL 6.0 packages. Type `Y` and press `ENTER` to choose the default action to install and update all packages at this time.

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

- The next stage involves reading in package header information from each package on the Updates DVD and performing a package dependency transaction check. No action is required of you during this step.

```

root@test:/media/dvd
File Edit View Search Terminal Help
[root@test dvd]# ./install-updates
Loaded plugins: refresh-packagekit, rhnplugin
This system is not registered with RHN.
RHN support will be disabled.
rhel6.0-update-media          | 1.9 kB    00:00 ...
rhel6.0-update-media/primary | 1.3 MB    00:00 ...
rhel6.0-update-media          |           2912/2912
Setting up Update Process
Resolving Dependencies
--> Running transaction check
--> Package NetworkManager.i686 1:0.8.1-5.el6_0.1 set to be updated
--> Package NetworkManager-glib.i686 1:0.8.1-5.el6_0.1 set to be updated
--> Package NetworkManager-gnome.i686 1:0.8.1-5.el6_0.1 set to be updated
--> Package apr-util.i686 0:1.3.9-3.el6_0.1 set to be updated
--> Package apr-util-ldap.i686 0:1.3.9-3.el6_0.1 set to be updated
--> Package autofs.i686 1:5.0.5-23.el6_0.1 set to be updated
--> Package bind-libs.i686 32:9.7.0-5.P2.el6_0.1 set to be updated
--> Package bind-utils.i686 32:9.7.0-5.P2.el6_0.1 set to be updated
--> Package bzip2.i686 0:1.0.5-7.el6_0 set to be updated
--> Package bzip2-devel.i686 0:1.0.5-7.el6_0 set to be updated
--> Package bzip2-libs.i686 0:1.0.5-7.el6_0 set to be updated
--> Package cifs-utils.i686 0:4.4-5.el6_0.2 set to be updated
--> Package cups.i686 1:1.4.2-35.el6_0.1 set to be updated
--> Package cups-libs.i686 1:1.4.2-35.el6_0.1 set to be updated
    
```

Once the package dependency transaction check has completed, updating of packages is executed. No action is required of you during this step.

```

root@test:/media/dvd
File Edit View Search Terminal Help
Transaction Summary
-----
Install      2 Package(s)
Upgrade     161 Package(s)

Total download size: 229 M
Downloading Packages:
-----
Total                               17 MB/s | 229 MB    00:13
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID fd431d51: NOK
EY
rhel6.0-update-media/gpgkey          | 6.3 kB    00:00 ...
Importing GPG key 0xFD431D51 "Red Hat, Inc. (release key 2) <security@redhat.com
>" from /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Importing GPG key 0x2FA658E0 "Red Hat, Inc. (auxiliary key) <security@redhat.com
>" from /etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Updating      : kernel-headers-2.6.32-71.24.1.el6.i686                1/324
  Updating      : hplip-common-3.9.8-33.el6_0.1.i686                  2/324
  Updating      : tzdata-2011d-3.el6.noarch                            3/324
    
```

NOTE

While the messages displayed during package updates appear to indicate that package *downloading* is occurring, no actual network access is performed when using the Updates DVD.

- Once all package updating has completed, several new packages are then installed. Upon completion of installation the entire update has completed.

```

root@test:/media/dvd
File Edit View Search Terminal Help
Total download size: 1.4 M
Installed size: 5.4 M
Downloading Packages:
-----
Total                               6.5 MB/s | 1.4 MB    00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing      : numactl-2.0.3-9.el6.i686                1/6
  Installing      : compat-libstdc++-33-3.2.3-69.el6.i686   2/6
  Installing      : compat-libcap1-1.10-1.i686             3/6
  Installing      : compat-expat1-1.95.8-8.el6.i686        4/6
  Installing      : numactl-devel-2.0.3-9.el6.i686         5/6
  Installing      : glibc-static-2.12-1.7.el6_0.5.i686     6/6

Installed:
compat-expat1.i686 0:1.95.8-8.el6      compat-libcap1.i686 0:1.10-1
compat-libstdc++-33.i686 0:3.2.3-69.el6  glibc-static.i686 0:2.12-1.7.el6_0.5
numactl.i686 0:2.0.3-9.el6           numactl-devel.i686 0:2.0.3-9.el6

Complete!
[root@test dvd]#

```

- When the update completes, execute the following commands:

```
cd ..
eject
```

NOTE: If there is more than one media device, you must specify the device to be ejected; e.g., **eject /dev/hda**.

- Remove the disc from the optical media drive and store.
- Type “reboot” at the system prompt and verify that the new Red Hat kernel boots.
- Continue with the steps in the next section to install RedHawk Linux.

5.2 Installing RedHawk Linux

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

- Log in as root and take the system down to single-user mode:
 - If in GUI mode, right click on the desktop and select Open Terminal.
 - At the system prompt, type **init 1**.

2. Locate the disc labeled “RedHawk Linux OS” appropriate to your system’s architecture and insert it into the optical media drive.
3. If the device does not mount automatically, execute the **mount** command appropriate for your optical media device.
4. Execute the following commands to install RedHawk Linux, substituting the actual mount point used for your optical media device for **/media/dvd** if necessary:

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

```

root@test:/media/dvd
File Edit View Search Terminal Help
[root@test dvd]# ls
Documentation  install-redhawk  Packages  uninstall-redhawk
extras        launch.in        restore   yum.repos.d
install       launch.un        uninstall
[root@test dvd]# ./install-redhawk

Installing RedHawk Linux 6.0 (Glacier)

This script will install RedHawk Linux 6.0 on the current
system (test).

Do you want to install RedHawk Linux 6.0? [y/n] y

-----
Concurrent End-User License Agreement (the EULA).

This license should be read before acceptance of its terms.  If the terms
are not acceptable then you must reject the license and terminate the
installation process.

A copy of this license may be found in /usr/share/doc/ccur/RedHawk-EULA
after installation.

-----
View, Accept, or Reject the terms of this license? [v/a/r] █

```

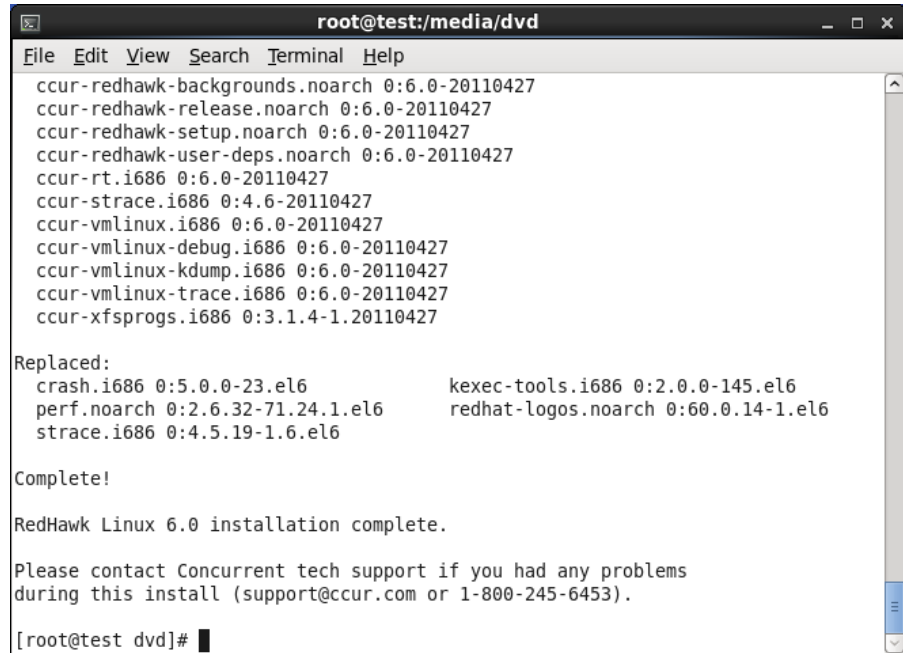
NOTES

- NOTE 1: Special instructions are displayed during the installation that may require appropriate action, such as video card installation and configuration. Follow the on-screen instructions until the installation script completes.
- NOTE 2: If you choose to install an optional graphics package on a 32-bit system, it may be necessary to add **vmalloc=192M** to the kernel’s grub command line in the system’s **/etc/grub.conf** file in order for graphics to start properly.
- NOTE 3: The installation script will prompt you to accept or reject the End User License Agreement. You can view the agreement by entering “v” at the prompt.
- NOTE 4: While the messages displayed during RedHawk installation appear to indicate that package *downloading* is occurring, no actual network access is performed when using the RedHawk DVD.

NOTE 5: It is normal for the installation the script to appear to be stalled near the end of installation. It is performing several background operations that are required and the script must not be interrupted.

NOTE 6: Near the end of installation you may be presented with dialogs indicating that the system cannot get an exclusive lock on the packaging backend. These dialogs are erroneous and may be closed and ignored.

- The following messages will be displayed once all RedHawk package installation has completed successfully.



```

root@test:/media/dvd
File Edit View Search Terminal Help
ccur-redhawk-backgrounds.noarch 0:6.0-20110427
ccur-redhawk-release.noarch 0:6.0-20110427
ccur-redhawk-setup.noarch 0:6.0-20110427
ccur-redhawk-user-deps.noarch 0:6.0-20110427
ccur-rt.i686 0:6.0-20110427
ccur-strace.i686 0:4.6-20110427
ccur-vmlinux.i686 0:6.0-20110427
ccur-vmlinux-debug.i686 0:6.0-20110427
ccur-vmlinux-kdump.i686 0:6.0-20110427
ccur-vmlinux-trace.i686 0:6.0-20110427
ccur-xfspgms.i686 0:3.1.4-1.20110427

Replaced:
  crash.i686 0:5.0.0-23.el6          kexec-tools.i686 0:2.0.0-145.el6
  perf.noarch 0:2.6.32-71.24.1.el6   redhat-logos.noarch 0:60.0.14-1.el6
  strace.i686 0:4.5.19-1.6.el6

Complete!

RedHawk Linux 6.0 installation complete.

Please contact Concurrent tech support if you had any problems
during this install (support@ccur.com or 1-800-245-6453).

[root@test dvd]#

```

- When the installation completes, execute the following commands:

```
cd ..
eject
```

- Remove the disc from the drive and store.
- If you installed in text mode, edit `/etc/inittab` and change the first line:

```
from: id:3:initdefault:
to:   id:5:initdefault:

to re-establish X11 mode.
```

- Reboot the system choosing a new RedHawk kernel and verify that the system boots properly.
- Continue with the next section if an RCIM is to be installed. Otherwise, proceed to the following section to install Frequency-Based Scheduler (FBS).

5.3 Installing the RCIM

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

5.3.1 Quick Hardware Installation Checklist

1. Before installing the RCIM, determine if you will be using the RCIM to accept or deliver external interrupts and the mode in which the RCIM will run (see “RCIM Connection Modes” on page 56).
2. Verify that the `ccur-rcim` RPM has been installed by executing the command `rpm -q ccur-rcim`. The system will inform you if it is not installed. It is a standard package in the RedHawk installation.
3. Power down the system and remove all power cords.

NOTE

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

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

5.3.2 Verifying RCIM Operation

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

```
cat /proc/driver/rcim/status
```

You should see an output similar to that shown below:

```
RCIM-III board 0 is at revision 1 eeprom 2.0.
```


This is a standalone (isolated) rcim.
 Has IRQ 108 and major number 253
 GPS present, RCIM firmware version 10

NOTE

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

- To view the current RCIM configuration, issue the following command:

```
cat /proc/driver/rcim/config
```

You should see output similar to the following:

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

5.3.3 Installing NTP Updates for GPS Module

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

- The **ccur-ntp** RPM is available via Concurrent's Software Repository system and it is also included in the *extras* directory on the RedHawk product media. Execute the following steps to install this RPM:
 - Mount the RedHawk product media and `cd` to the *extras* directory or download the **ccur-ntp** RPM from Concurrent's web site as described in section 10.2.2 on page 66.
 - Install the RPM via the following command:

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

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

5.4 Installing Frequency-Based Scheduler (FBS) Software

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

- With RedHawk Linux Version 6.0 running, log in as root and take the system down to single-user mode.
- Locate the disc labeled "RedHawk Linux Frequency-Based Scheduler Version 6.0" appropriate to your system's architecture and insert it into the optical media drive.

3. If the optical media disc device does not mount automatically, execute the **mount** command; for example:

```
mount /media/dvd
```

4. To install, execute the following commands:

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

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

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

```
cd ..
eject
```

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

5.5 Installing the PCI-to-VME Bridge

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

5.5.1 Installing the Hardware

The hardware consists of a PCI adapter, a VME 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 VME 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 VME adapter card:
 - a. Ensure that the VME chassis is powered down.

- b. Decide whether the VME 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 VME adapter card's transceiver.
 - e. Power on both systems.
 - f. Ensure that the READY LEDs on both adapter cards are lit.

5.5.2 Installing the Software

1. With RedHawk Linux Version 6.0 running on the 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 6.0" appropriate to your system's architecture and insert it into the optical media drive.
3. If the optical media device does not mount automatically, execute the **mount** command; for example:


```
mount /media/dvd
```
4. To install, execute the following commands:


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

 Follow the on-screen instructions until the installation script completes.
5. When the installation completes, execute the following commands:


```
cd ..  
eject
```
6. Remove the disc from the optical media drive and store.
7. Exit single-user mode (Ctrl-D).

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

5.6 Installing the PCIe-to-VME Bridge

If you are using the Concurrent PCIe-to-VME Bridge adapter to connect your iHawk system to VME cards, it can be installed at this time. If you are not installing this feature, skip to the next section.

5.6.1 Installing the Hardware

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

The Concurrent PCIe-to-VME bridge hardware consists of:

- A PCIe host adapter card that should be installed in the iHawk host. The PCIe host adapter card contains no jumpers or switches that require configuration.
 - A VME bridge card that should be installed in the VME chassis. The VME bridge card contains 3.3V/5V key posts for the two on-board PMC sites. These posts are used to select the PCI signaling voltage to be used. Both posts must be set to the same position.
 - A PCIe bridge cable that is used to connect these two boards.
1. To install the PCIe host adapter card in your iHawk system:
 - a. Power down the system.
 - b. Locate a vacant PCIe x4 (or greater) card slot in the chassis that supports a bus master.
 - c. Remove the metal plate that covers the card slot at the rear of the chassis.
 - d. Insert the PCIe adapter card into the slot.
 - e. Fasten the adapter card in place with the mounting screw.
 2. To install the VME bridge card:
 - a. Ensure that the VME chassis is powered down.
 - b. The VME bridge card should be installed in the left-most (system controller) slot in your VME chassis.
 3. To install the cable:
 - a. Ensure that both the iHawk system and VME chassis are still powered down.
 - b. Plug one end of the cable into the PCI host adapter card's transceiver.
 - c. Plug the other end of the cable into the VME bridge card's transceiver.
 - d. Power on the iHawk system and VME chassis.

5.6.2 Installing the Software

1. With RedHawk Linux Version 6.0 running on the system, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk Linux PCIe-to-VME Bridge Version 6.0" appropriate to your system's architecture and insert it into the optical media driver.

3. If the optical media device does not mount automatically, execute the mount command; for example:

```
mount /media/dvd
```

4. To install, execute the following commands:

```
cd /media/dvd
./install-pcievme
```

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

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

```
cd ..
eject
```

6. Remove the disc from the optical media drive and store.
7. Optionally configure the VME windows as described in the next section and then exit single-user mode (Ctrl-D).

5.6.3 Configuring the PCIe-to-VME Windows

The PCIe-to-VME support provides up to eight configurable VME windows that VME cards installed in the VME chassis may use to respond to their respective RedHawk VME driver's access requests.

The `/etc/sysconfig/pcievme-vmemap` configuration file controls the creation of these VME windows. This file is used by the `pcievme` init service to automatically configure the Concurrent PCIe-to-VME bridge card during system bootup, or during a `service pcievme start` or `service pcievme restart` invocation.

Note that the `pcievme-config` utility may also be used to manually reconfigure the VME windows without using the `pcievme` init service. See the `pcievme-config(1)` manual page for more details.

The `pcievme-vmemap` file allows you to specify the following attributes for each window:

- VME start and end addresses of the window
- VME address type: A16, A24 or A32
- supervisor or user VME AM code
- enable or disable VME posted writes

Refer to and modify the `/etc/sysconfig/pcievme-vmemap` file to setup the VME windows that are required for your specific VME devices and corresponding RedHawk VME drivers.

After the `pcievme` init service has successfully setup the VME windows, the `/proc/driver/pcievme/tundra_universe_map` file may be used to view the current VME window configuration. For example, the `tundra_universe_map` file may contain text similar to the following listing.

```
## Universe driver 1.2-chthv2 06Dec2010, ccur_version 3.0.1

0:BAR0: 0xF8400000 - 0xF8400FFF
```

```

0:SLAVES: 8

## Enabled PCI Target Image Registers (LSI0-LSI7)
_PCI:0:L0: CTL=0xC0801000 BS=0xF8401000 BD=0xF8411000 TO=0x07BEF000
_PCI:0:L1: CTL=0xC0820000 BS=0xF8430000 BD=0xF84C0000 TO=0xE7BD0000
_PCI:0:L2: CTL=0xC0820000 BS=0xF84C0000 BD=0xF8500000 TO=0xE8B40000
_PCI:0:L4: CTL=0xC0810000 BS=0xF8411000 BD=0xF8421000 TO=0x079EF000

## Enabled VMEbus Slave Image Registers (VSI0-VSI7)
_PCI:0:V0: CTL=0xA0F70000 BS=0x02000000 BD=0x02FF0000 TO=0x00000000

## Board:Code:Width:AddrSpace:VMEstart:VMEend:PCIstart:PCIend:Size
_VME:0:S:32:A16:0xFFFF0000:0xFFFFFFFF:0xF8401000:0xF8410FFF:64K
_VME:0:U:32:A32:0xE0000000:0xE008FFFF:0xF8430000:0xF84BFFFF:576K
_VME:0:U:32:A32:0xE1000000:0xE103FFFF:0xF84C0000:0xF84FFFFF:256K
_VME:0:U:32:A24:0xFFE00000:0xFFE0FFFF:0xF8411000:0xF8420FFF:64K

```

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 Disc

Discs labeled "Software Updates for RedHawk" include updates to RedHawk and its products. If a RedHawk Updates disc is provided, install it 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 drive.
3. If the disc device does not mount automatically, execute the **mount** command; for example:

```
mount /media/dvd
```

4. Execute the following commands to install the RedHawk updates included on the disc:

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

When the installation is complete, 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.

Before installing any modules, check for NUU updates separately. Apply any NUU updates and then restart NUU before applying any other updates.

With NUU up-to-date, from the NUU Main window, 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 ..  
eject
```

6. Remove the disc from the optical media drive and store.
7. If there are more update discs, repeat these steps to install all updates.
8. Reboot the system choosing the updated RedHawk kernel and verify that the system boots properly.
9. 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 any “Software Updates for RedHawk” disc you may have received). They should be downloaded and installed now. Refer to “Software Updates” on page 65 for details.
2. When complete, or if no updates are available, reboot the system choosing the new RedHawk kernel and verify that the system boots properly.

6.0. NVIDIA Graphics Configuration

The Unified NVIDIA Linux Display Driver (version 270.40) 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.

See sections **6.1 Supported NVIDIA GPUs** and **6.2 Multiple NVIDIA Video Card Considerations** for more information on NVIDIA driver compatibility.

After you have installed the correct driver for your NVIDIA card, you must create a new `/etc/X11/xorg.conf` to ensure the proper operation of the driver under X Window System Version 7.1.1 and RedHawk Linux Version 6.0.

It is required to specify that X load the “nvidia” driver along with the appropriate modules in `xorg.conf`. These entries will not be present in the default `xorg.conf` file created by the installer. Follow the process described in section **6.4 Using nvidia-xconfig** to generate the new `xorg.conf` file.

After you have generated the new `xorg.conf`, follow the process described in section **6.5 Using nvidia-settings** to configure your displays, whether single or multiple displays. The types of display configurations are described in **6.3 Display Configurations**.

You should create all new display configurations based on these processes. Failure to do so may prevent the proper operation of your system.

Do not base any configuration on the `xorg.conf` file created by the install process or on any other previous configuration you might have been using.

After you have completed the basic process for your display configuration, you can edit the resulting `/etc/X11/xorg.conf` for final adjustments.

6.1 Supported NVIDIA GPUs

If the `ccur-nvidia-glx` rpm has been installed on your system, you can view the list of all NVIDIA GPUs supported by the 190.42 driver in “Appendix A” in the following file:

```
/usr/share/doc/ccur-nvidia-glx-270/README.txt
```

Alternatively, you can examine the list of supported GPUs for driver 190.42 at the NVIDIA web site which is located at:

<http://www.nvidia.com>

Note that if your NVIDIA GPU is not supported by the 190.42 driver, you will have to do one of the following:

- replace the existing older NVIDIA card in your system with a newer NVIDIA card that is supported by the 190.42 driver
- go to the NVIDIA web site mentioned above and download a legacy NVIDIA driver that supports your NVIDIA card, and follow the installation instructions described on that driver’s web page

NOTE

The NVIDIA driver that is provided with RedHawk installations has been modified to reduce jitter on shielded CPUs in a real-time environment. Therefore, if you choose to download and use a legacy NVIDIA driver directly from the NVIDIA web site, you may observe interference on shielded CPUs during graphics activity.

6.2 Multiple NVIDIA Video Card Considerations

Note that you may not use the unified 190.42 Driver and a legacy NVIDIA Linux Display Driver together at the same time.

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

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

If you have a mixed video card system as described here, it is recommended that you upgrade your legacy NVIDIA video cards to unified video cards and use the unified 190.42 driver as your display driver.

6.3 Display Configurations

Displays can be configured as a single monitor with one or more controlling keyboards and mice, or together as multiple monitors with one or more controlling keyboards and mice.

Screens can be configured left, right, above, or below each other in any fashion. The differences in how the screens operate and interact are described in this section.

Use the **nvidia-settings** utility to create all multiple display configurations. See section **6.5 Using nvidia-settings**.

6.3.1 Single

Single mode is a single X Screen displayed as a single Desktop on a single monitor.

Use the **nvidia-xconfig** utility to create a single NVIDIA display. See section **6.4 Using nvidia-xconfig**.

6.3.2 Xinerama

Xinerama mode enables multiple X Screens displayed as a single Desktop spanning multiple monitors. This mode has the following characteristics:

- A separate X Screen for each display is used.
- The mouse moves freely between displays.
- If you maximize a window, it fills a single display.
- You can drag objects from one display to another.

- This mode causes a little extra overhead for X.

Xinerama requires entries in **xorg.conf** to describe the layout of the screens to the X server and the setting of the server flag, “Xinerama” “1”.

Xinerama is done by the X server and does not require the use of NVIDIA, but only configurations based on NVIDIA and the **nvidia-settings** utility are covered in this document.

6.3.3 MultiDesktop

MultiDesktop mode enables multiple X Screens displayed as separate Desktops on each monitor. This mode has the following characteristics:

- A separate X Screen for each display is used.
- The mouse moves freely between displays.
- If you maximize a window, it fills a single display.
- You can *not* drag objects from one display to another.
- This causes a little extra overhead for X.

This requires entries in **xorg.conf** to describe the layout of the screens to the X server and the setting of the server flag, “Xinerama” “0”, or by the omission of the Xinerama server flag.

MultiDesktop is done by the X server and does not require the use of NVIDIA, but only configurations based on NVIDIA and the **nvidia-settings** utility are covered in this document.

6.3.4 Twinview

Twinview mode allows you to display a single Desktop that spans two monitors or to clone the output and display a duplicated Desktop on each monitor. This mode has the following characteristics:

- The mouse moves freely between displays.
- If you maximize a window, it expands across both displays unless you:
 - set a special flag which emulates the behavior of Xinerama
 - operate in clone mode
- In clone mode, the monitors display exactly the same image.
- You can drag objects from one display to another unless you are operating in clone mode.
- Twinview mode causes no extra overhead for X.

The NVIDIA driver conceals all information about multiple display devices from the X server; as far as X is concerned, there is only one X Screen.

All display devices share one frame buffer. Thus, all the functionality present on a single display (e.g., accelerated OpenGL) is available with Twinview.

Twinview only works on a single GPU with multiple heads. It will not work between separate cards or heads with different BusIDs (like Xinerama can).

The “Twinview” and “metamodes” options must be set in the **xorg.conf** Screen section.

See `/usr/share/doc/ccur-nvidia-glx-270/README.txt` for a detailed discussion on Twinview.

6.3.5 Twinview-Xinerama

This mode enables multiple X Screens to join Twinview Screens as a single Desktop on three or more monitors.

The “Twinview” and “metamodes” options must be set in the **xorg.conf** Screen section for each supported device.

This also requires entries in **xorg.conf** to describe the layout of the Screens to the X server and the setting of the server flag, “Xinerama” “1”.

The definition of the Xinerama server flag to X will override the ability of the Twinview Screen to emulate Xinerama.

If you maximize a window in the Twinview Screen, it expands across both displays in the Twinview Screen if not in clone mode.

If you maximize a window in a non-Twinview Screen, it fills that display.

6.3.6 Twinview-MultiDesktop

This mode enables multiple X Screens to join Twinview Screens as separate Desktops on three or more monitors.

The “Twinview” and “metamodes” options must be set in the **xorg.conf** Screen section for each supported device.

This also requires entries in **xorg.conf** to describe the layout of the Screens to the X server and the setting of the server flag, “Xinerama” “0”, or by the omission of the Xinerama server flag.

The definition of the Xinerama server flag to X will override the ability of the Twinview Screen to emulate Xinerama.

If you maximize a window in the Twinview Screen, it will expand across both displays in the Twinview Screen if not in clone mode.

If you maximize a window in a non-Twinview Screen, it will fill that display.

6.4 Using nvidia-xconfig

You must create `/etc/X11/xorg.conf` using the **nvidia-xconfig** utility before using your system with the nvidia driver.

The **nvidia-xconfig** utility will generate some superfluous entries which you may choose to remove. See the example **xorg.conf** file in section **6.6** for reference.

Follow this procedure exactly to generate the file:

1. Read sections **6.0** - **6.3** and complete any actions required.
2. Reboot to runlevel 3.
3. Move `/etc/X11/xorg.conf` to `/etc/X11/xorg.conf.orig`:

```
mv -i /etc/X11/xorg.conf /etc/X11/xorg.conf.orig
```

4. Run the `nvidia-xconfig` utility (`/usr/bin/nvidia-xconfig`):

```
nvidia-xconfig
```

5. Reboot to runlevel 5.

See the `nvidia-xconfig(1)` man page for details about using `nvidia-xconfig`.

6.5 Using nvidia-settings

The `nvidia-settings` utility should be used to configure any of the multidisplay configurations described in section 6.3, and is the preferred method for tuning single display settings.


Before using `nvidia-settings`, complete the `nvidia-xconfig` procedure outlined in section 6.4. Failure to do so may result in poor or unstable performance.

See the `nvidia-settings(1)` man page for details about using `nvidia-settings`.

1. To start the `nvidia-settings` utility, issue following the command:

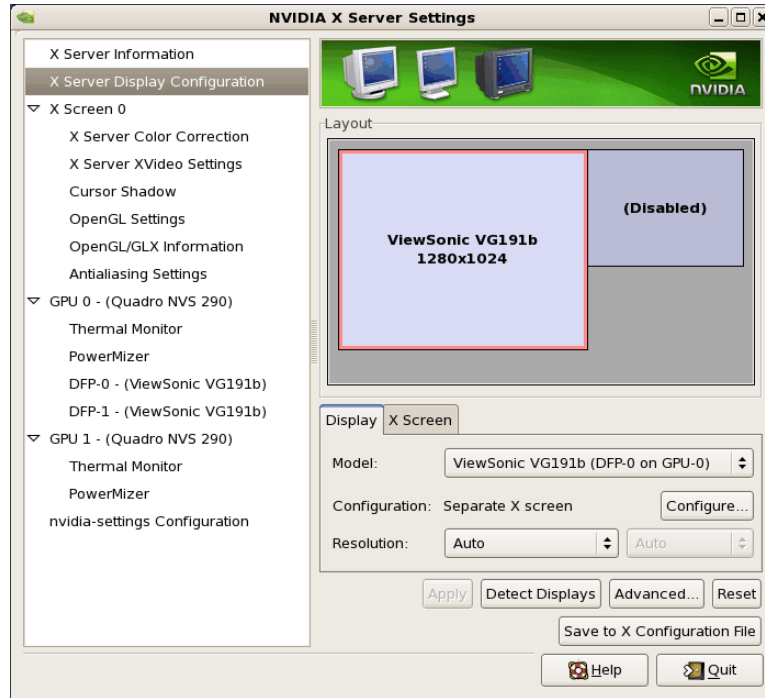
```
nvidia-settings
```

The GUI will start:



Here we see that we are viewing “X Server Information”, as indicated by the list of options on the left.

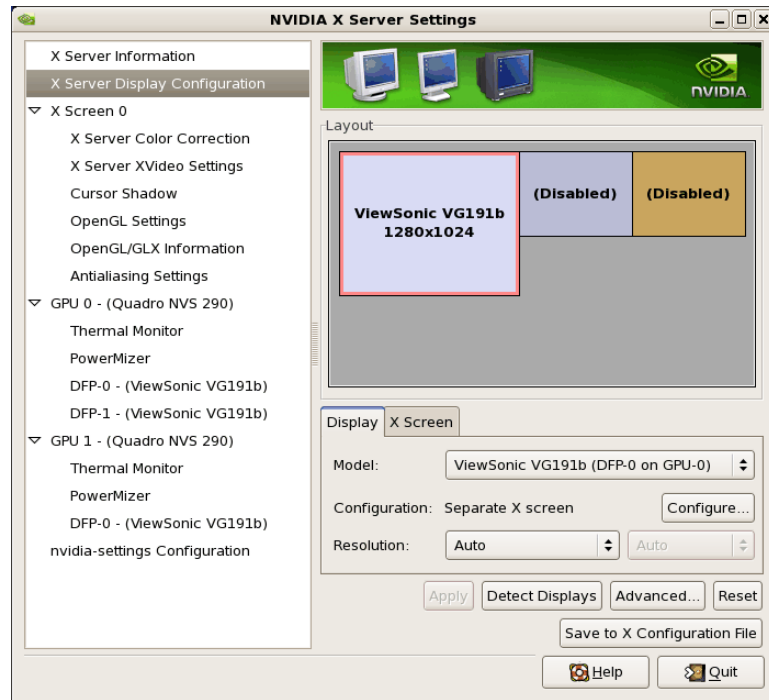
- To create a new “X Server Display Configuration” select that option from the list on the left as shown below:



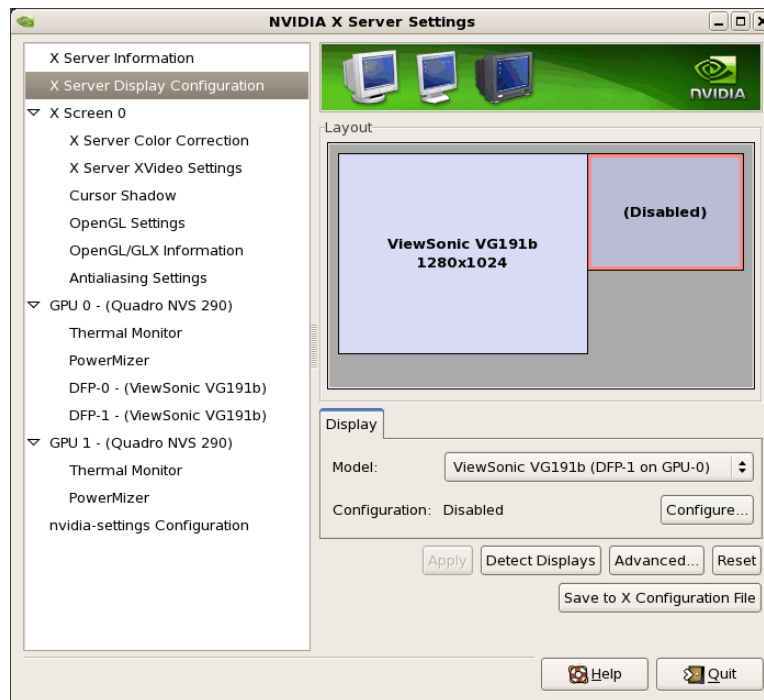
In the Layout window, we see a graphical depiction of the number of displays detected and configured. Notice that the large box is the primary monitor already configured by `nvidia-xconfig`. The smaller box to the right (Disabled) is a display that has not yet been configured.

Click on the Detect Displays button if there are other monitors connected that are not shown in the Layout window.


Below is what the Layout window would look like if we detected a third monitor:



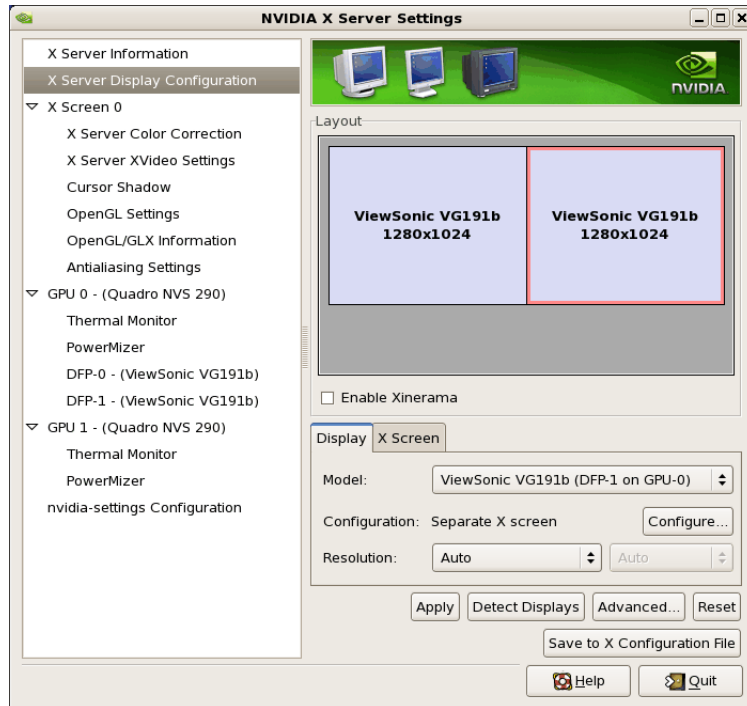
3. To configure two detected displays as a MultiDesktop, we would select the second monitor in the Layout window by clicking on it. We can see the second monitor is selected for configuration by the red outline around the box:




- Click on the **Configure** button below the Layout window and select the “Separate X screen” option, then click **OK**:



- If we wanted to enable Xinerama we could do so at this point by selecting the “Enable Xinerama” check box below the Layout window:

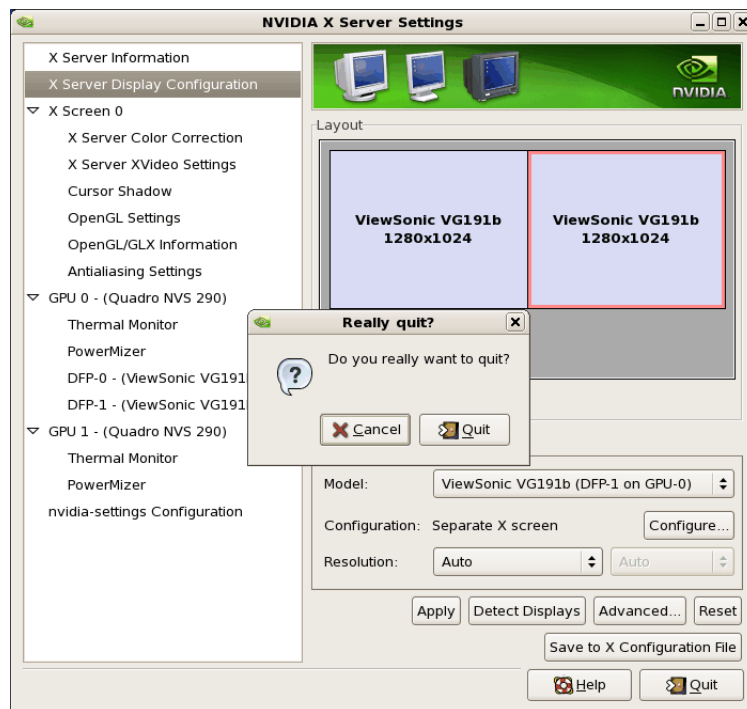


- To save the configuration to **xorg.conf**, click on the Save to X Configuration File button near the bottom of the GUI:



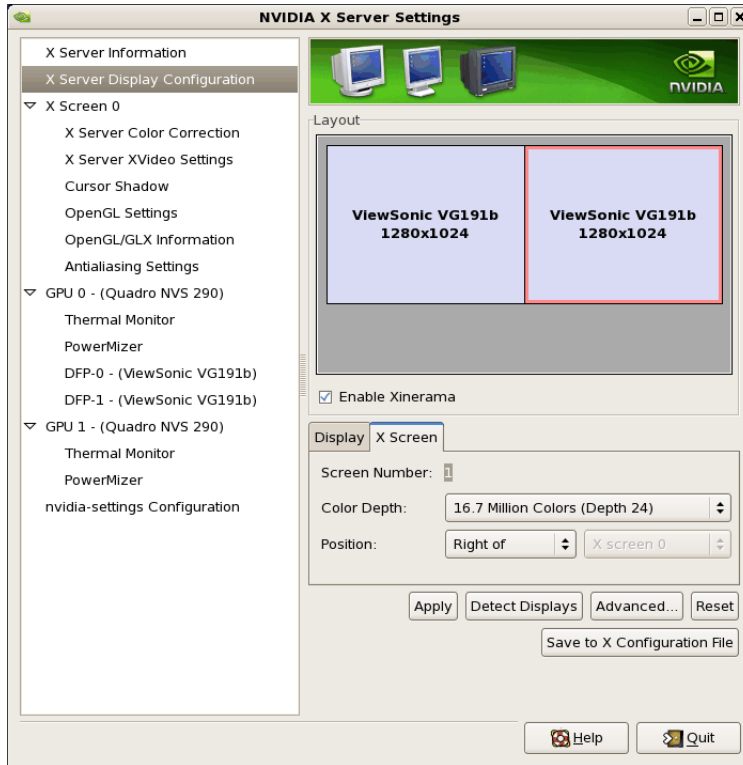
Then click **Save** on the Save X Configuration dialog box as shown above. (Always leave the Merge with existing file box checked).

- After you have saved the new **xorg.conf**, click on **Quit** to exit **nvidia-settings**:

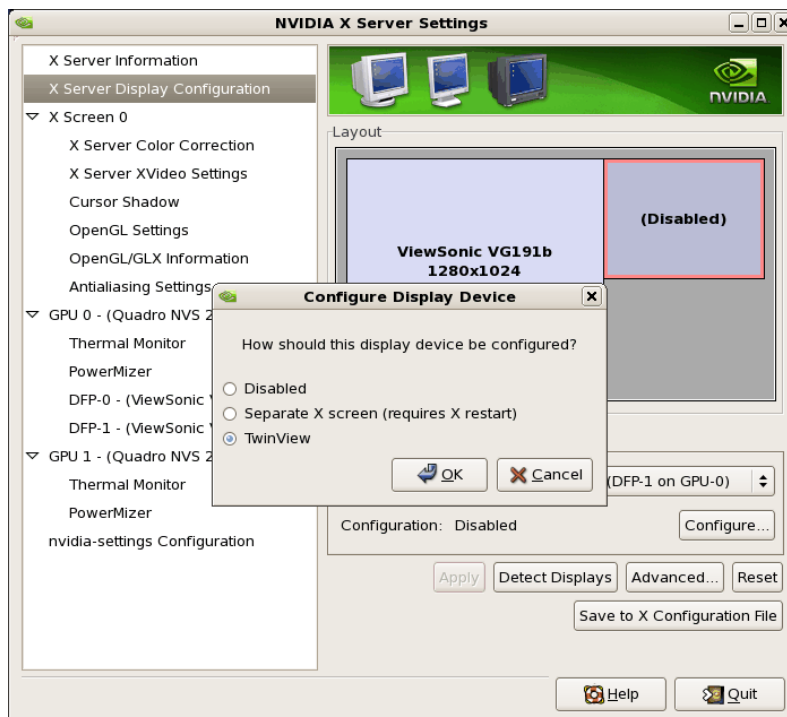


8. Reboot to runlevel 5. You should now be up on two X screens. You can restart `nvidia-settings` to further examine or modify your configuration.



Below shows an example configuring the Xinerama option:

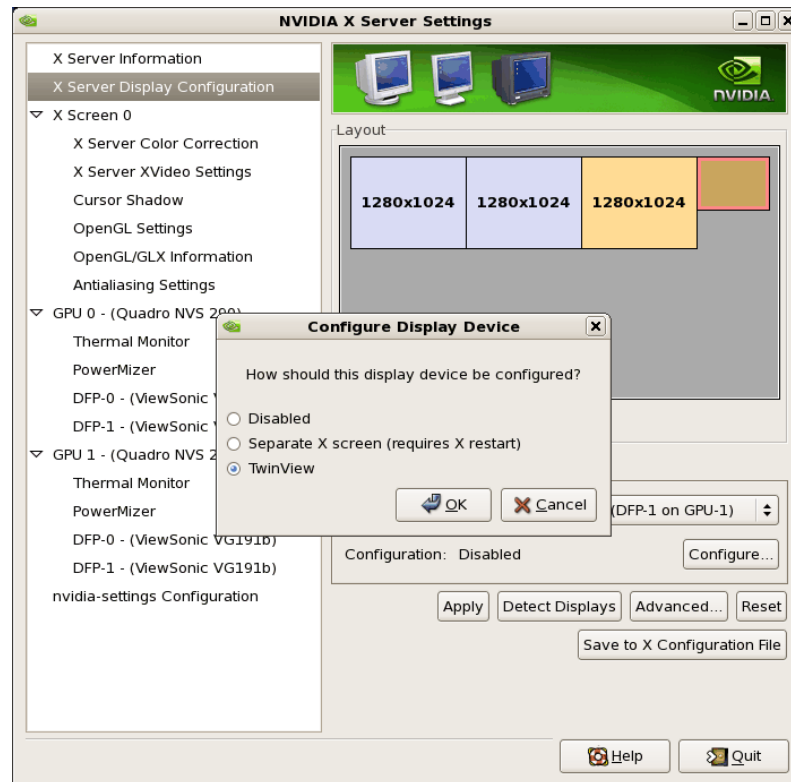


Below shows an example configuring the Twinview option:



Below shows some stages configuring Twinview-Xinerama or Twinview-MultipleDesktop (the difference would be in selecting or not selecting the Xinerama check box):





Note that **nvidia-settings** will allow you to do much more than covered here.

In some configurations that don't require restarting X, you can apply changes to test them before saving the config.

You should examine your `/etc/X11/xorg.conf` files after creating them.

It is easier to hand edit for setting some options, but make sure to make incremental backups of your configurations.

6.6 Modified xorg.conf file generated by nvidia-xconfig

The file shown below was created by the **nvidia-xconfig** utility and then pruned down to remove machine specific entries and some superfluous defaults. Some of the entries in this example are not required, but were left in to provide structure for future hand edits.

There should be no need to modify this template if all you are running is a single display with a typical keyboard and mouse.

You may choose to model your basic **xorg.conf** file after this one.

You can enable SLI options by uncommenting the appropriate one in the "Device" Section.

You can run **nvidia-settings** on this template to create any of the display configurations described in this document.

```
Section "ServerLayout"
    Identifier      "ccur nvidia glx realtime"
    Screen         0      "Screen0"
    InputDevice    "Keyboard0" "CoreKeyboard"
    InputDevice    "Mouse0" "CorePointer"
EndSection
```

```

Section "Module"
    Load      "dbe"
    Load      "extmod"
    Load      "type1"
    Load      "freetype"
    Load      "glx"
EndSection

Section "Extensions"
    Option    "Composite" "Disable"
EndSection

Section "InputDevice"
    Identifier "Mouse0"
    Driver     "mouse"
EndSection

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

Section "Monitor"
    Identifier "Monitor0"
EndSection

Section "Device"
    Identifier "Device0"
    Driver     "nvidia"
    #Option    "SLI" "Auto"
    #Option    "SLI" "AFR"
    #Option    "SLI" "SFR"
    #Option    "SLI" "SLIAA"
EndSection

Section "Screen"
    Identifier "Screen0"
    Device     "Device0"
    Monitor    "Monitor0"
    DefaultDepth 24
    SubSection "Display"
        Depth    24
    EndSubSection
EndSection

```

6.7 Troubleshooting Problems

The X server looks for a configuration in `/etc/X11/xorg.conf` each time it starts up (normally in runlevel 5). You will find that newer versions of X are getting away from depending on configuration files and require only minimum information to create even the largest multidisplay configurations. In fact, X will work in default mode with NO configuration file at all.

It is not uncommon for obsolete options or incorrect settings specified in `xorg.conf` to cause the system to fail to start X correctly. Sometimes X will get so confused it won't be able to do anything but lock up, cause the system to panic or reboot, or just come up in some wacky video state.

Simply stopping and restarting X does not always do the trick. It is best to issue the command(s):

```
init 1; init 5
```

or

```
reboot
```

See the X server actions logged in `/var/log/Xorg.0.log` and `Xorg.0.log.old`.

Normally, X will respond to errors at start up time by telling you there is a problem and offer to show you the log, then create a fresh configuration.

A fresh configuration looks very much like the configuration created by the installer.

If you are unable to create a working `xorg.conf` using the `nvidia-xconfig` utility, it may be because the directions were not followed precisely, or problems may have been created by hand edits after the fact. Take logical steps and verify functionality at the various steps.

Try removing `/etc/X11/xorg.conf` entirely, reboot to runlevel 3, and run `nvidia-xconfig`.

DO NOT BOOT TO RUNLEVEL 1 OR SINGLE USER MODE AND RUN `nvidia-xconfig` – THIS WILL CREATE A BAD `xorg.conf` FILE.

If you have the original `xorg.conf` created by the installer, but it's not working, you probably have the wrong driver defined for the primary video display defined by the BIOS. Try the following config:

```
Section "ServerLayout"
    Identifier      "Default Layout"
    Screen          0  "Screen0"  0 0
    InputDevice     "Keyboard0"  "CoreKeyboard"
EndSection

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

Section "Device"
    Identifier      "Videocard0"
    Driver          "vesa"
EndSection

Section "Screen"
    Identifier      "Screen0"
    Device          "Videocard0"
    DefaultDepth    24
    SubSection "Display"
        Viewport    0 0
        Depth       24
    EndSubSection
EndSection
```

Typical drivers in iHawk systems include:

- `radeon` – Common in many servers with embedded graphics chips
- `nv` – The generic Linux NVIDIA driver
- `nvidia` – Provided by NVIDIA or Concurrent. NVIDIA drivers are modified by Concurrent for real-time performance. Use the correct driver from Concurrent to ensure the best performance for your system.
- `vesa` – The generic fall back low performance driver which should work for all VGA adapters.

NVIDIA has a number of generations of graphics cards which cannot all use the same “`nvidia`” driver version. Not even the generic “`nv`” driver in the Linux kernel works on all NVIDIA cards. The “`nvidia`” driver will actually disable an `nvidia` controller or fail to load if it is the wrong version for that card. You cannot mix cards of different generations if the driver does not support all of them.

There is no need to specify resolutions, refresh rates, and color depths for modern VGA adapters and Displays. It is best to allow the hardware and software to negotiate these settings.

If resolutions, refresh rates, and color depths were generated by one of NVIDIA's utilities (**nvidia-xconfig** or **nvidia-settings**) they can be included, but they are superfluous, as well as most options, unless you want to redefine the compiled-in defaults of the X server or device driver.

WARNING

USING OLD CONFIGURATION FILES IS RISKY. IT IS NOT UNCOMMON FOR OBSOLETE OPTIONS AND ENTRIES TO CAUSE YOUR SYSTEM TO HANG, REBOOT, OR PREVENT NORMAL USAGE. YOU SHOULD START WITH THESE BASIC CONFIGS AND GRADUALLY MODIFY THEM WITH THE OPTIONS YOU THINK YOU NEED IN ORDER TO CREATE CUSTOMIZED CONFIGURATIONS.

7.0. ATI Graphics Configuration

7.1 ATI Graphics Configuration

The ATI FirePro graphics driver is included in the RedHawk 6.0 release in the `ccur-ati-fglrx` RPM. The companion kernel graphics `fglrx` module driver is also included in all RedHawk pre-built kernels.

7.2 Supported ATI Cards

This driver supports several ATI graphics cards, including FirePro, FireStream and FireGL.

7.3 ATI and NVIDIA Configuration Considerations

Please note that the ATI graphics support may not be installed at the same time as the NVIDIA graphics support on the same system.

If the `ccur-nvidia-cuda` and `ccur-nvidia-glx` RPMs are currently installed on your system, you must uninstall these `nvidia` RPMs before installing the `ccur-ati-fglrx` and `ccur-ati-stream` RPMs.

7.4 Display Configurations

Display can be configured as a single monitor or as multiple monitors with one or more controlling keyboards and mice. Screens can be configured to be left, right, above or below each other in any fashion. The differences in how the screens operate and interact are described in the following sections.

The `aticonfig` utility may be used to create various `xorg.conf` ATI graphics card display configurations.

You can either start from scratch and create your own `xorg.conf` file using the appropriate `aticonfig` options, or you can start with the provided `xorg.conf.fglrx` template and use `aticonfig` to finish the creation of the `xorg.conf` file.

Note that the provided template file contains additional mouse and keyboard sections that will not be present if you start from scratch when creating your `xorg.conf` file. However, you can add these and other missing sections to your `xorg.conf` file if necessary, after you have executed the `aticonfig` command.

The `aticonfig` utility contains many options, but only the most commonly options are discussed below. To view a list of all available `aticonfig` options, as root issue the following command:

```
# aticonfig | less
```

7.4.1 Single Head

If you have a single monitor connected to your ATI graphics card, then you can create an `xorg.conf` file with the following steps (as root):

```
# cd /etc/X11
# cp xorg.conf.fglrx xorg.conf
# aticonfig --initial
```

7.4.2 Dual Head

In this mode, the X server is configured to have two screens on two different channels of the same adapter. In the configuration file there will be two Monitor sections and two Screen sections. This will create two independent desktops managed by the window manager.

To create a dual head configuration `xorg.conf` file, execute the following steps as root:

```
# cd /etc/X11
# cp xorg.conf.fglrx xorg.conf
# aticonfig --initial=dual-head
```

You may make additional changes to the screen layout and other attributes of your X session by using the AMD Catalyst Control Center utility. See "Using Catalyst Control Center" on page 52 for more information.

7.5 MultiView

MultiView allows for the merging of independent desktop heads into a unified workspace allowing windows to freely cross X Screen boundaries.

To create a MultiView `xorg.conf` file (as root):

```
# cd /etc/X11
# cp xorg.conf.fglrx xorg.conf
# aticonfig --initial=dual-head --xinerama=on
```

You may make additional changes to the screen layout and other attributes of your X session by using the AMD Catalyst Control Center utility. See "Using Catalyst Control Center" on page 52 for more information.

7.6 Using Catalyst Control Center

In addition to using the `aticonfig` utility, you may also use the AMD Catalyst Control Center Linux Edition utility, `amdcccle`, from an active X session to change the layout and attributes of your display configuration.

To use the Catalyst Control Center, execute `amdcccle` as root from a window in the active X session, or navigate to **System -> Preferences -> ATI Catalyst Control Center** in the Gnome main menu.

The `amdcccle` utility provides various informational and configuration capabilities such as color adjustments, `xorg.conf` configuration, refresh and screen resolutions, 3D settings, etc.

Note that when switching back and forth between various X session layout configurations with `amdcccle`, it may be necessary to execute the following steps as root in order to get the newest screen configuration layout working:

```
# cd /etc/ati
# cp amdpcsdb.default amdpcsdb
```


After this point, you may need to re-enter various settings in an `amdcctl` invocation that were previously saved in the `amdpcsdb` configuration file.

7.7 Known Issues

The AMD ATI Catalyst version release notes are located in the `/usr/share/doc/ati` directory. The release notes contain a section listing the various known issues.

In addition, the following issues have also been observed:

- As was previously mentioned in the configuration sections, when switching back and forth between `xorg.conf` files, it may be necessary to copy the `/etc/ati/amdpcsdb.default` settings back into `/etc/ati/amdpcsdb` in order for the newly installed `xorg.conf` configuration to function properly.
- When running certain X applications such as SPECviewperf it has been observed that shielded CPUs may receive a fair number of page cache LRU drain softirq request cross processor interrupts (CPIs) due to `mlock(2)` calls issued by the X application. To eliminate these types of CPIs on the shielded CPU, you may issue the following command as root so that these LRU drain requests are executed as local events rather than triggered softirqs on the shielded CPU:

```
/bin/echo 0 > /proc/sys/kernel/schedule_on_each_cpu
```

7.8 Supporting Other ATI Graphics Controllers

If your system contains an ATI graphics card other than the ones listed under "Supported ATI Cards" on page 51, you may still make use of most other ATI cards by building a custom kernel that utilizes the ATI Radeon kernel driver module.

In order to configure this support into your system, take the following steps as root:

1. Build a custom kernel with the ATI `fglrx` graphics driver disabled, the DRM Radeon driver enabled, and the preallocated graphics pages feature disabled, since it is not supported by the Radeon driver.

Under the kernel configuration menu modify the following items:

```
Device Drivers ->
```

```
Graphics support ->
```

```
Disable: fglrx ATI kernel support
```

```
Device Drivers ->
```

```
Graphics support ->
```

```
Set to a value of 0: Preallocation of pages for graphics use
```

```
Device Drivers ->
```

```
Graphics support ->
```

```
Direct Rendering Manager (XFree86 4.1.0 and higher DRI support)
```

```
Enable as a module: ATI Radeon
```

```
Enable: Enable modesetting on radeon by default
```

Build the custom kernel and install it.

2. Exit any X sessions that are currently active.
3. As root, remove the **ccur-ati-fglrx** and **ccur-ati-stream** RPMs if they are currently installed:

```
# rpm -e ccur-ati-stream  
# rpm -e ccur-ati-fglrx
```

4. Remove the `/etc/X11/xorg.conf` file. The X startup processing should be able to automatically configure your X session properly using with the Radeon driver.

If you wish, you may optionally create a new `xorg.conf` file using "radeon" as the Driver value in the "Device" Section if there are attributes that are not set up properly by default without an `xorg.conf` file.

5. Reboot the system using the newly built custom kernel.

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

Table 7-1 Files Recommended for Backup

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

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

8.1.2 File System Swap Size Guidelines

Table 7-2 supplies the recommended swap partition sizes for various sizes of main memory. The swap space allocations listed below should be adequate for all real-time systems.

Table 7-2 Swap Partition Size Guidelines

Main Memory Size (GB)	Recommended Swap Partition Size (GB)
0 - 3	2
4 - 15	4
16 - 63	8
64 - 255	16
256 - 511	32
512 - 1024	64

8.1.3 RCIM Connection Modes

Before installing an RCIM on an iHawk system, determine the connection mode. For the RCIM I 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.

8.1.4 How to Use Linux Rescue

An installation or upgrade may not boot if the following conditions exist:

- the disk LABEL scheme is not correct
- there are mismatching LABEL or partition entries in `/etc/fstab` and/or `/etc/grub.conf -> /boot/grub/grub.conf`.

This can happen when dealing with multiple disks and can be extremely confusing.

It is possible to recover from these situations by using “linux rescue”. “Linux rescue” is a boot option on the RHEL 6.0 installation media.

If you cannot boot after the installation or upgrade, perform the following steps.

1. Boot either the RHEL 5.1, 5.2 or 5.4 installation disc.
2. At the prompt, type `linux rescue` or press F5 (F2 for more options).
3. Answer the questions about the keyboard and language set up, skipping the network setup (it can be configured later). When you get to the prompt that says “continue to boot”, choose “continue”.

You will boot into text mode with all partitions mounted.

4. Execute the following commands and use the procedures below to correct errors that prevented you from booting your installed or upgraded RHEL 6.0 installation.

```
chroot /mnt/sysimage
Export MANPATH=/usr/share/man
service gpm start
service network start
```

Use the following procedures to make corrections:

- To scroll up and down, use the `Shift+Page Up` and `Shift+Page Down` key combinations.
- To gain insight into how your files are configured and for naming the various partitions, devices, mount points and filesystem labels, you can use the following commands:

```
mount
cat /proc/partitions
cat /etc/fstab
cat /proc/swaps
cat /etc/grub.conf
```

- To find all these files on every partition, you can use the following commands:

```
find / -name fstab
find / -name grub.conf
```

- To see detailed information on all the filesystems, use `tune2fs` for ext2 and ext3 filesystems. See the `tune2fs(8)` man page for details.

To see the filesystem volume label on a device partition, issue the following command (substituting the proper device name):

```
tune2fs /dev/sda2 -l | grep volume
```

- You may have to change LABEL names and entries in the above mentioned files so they match when booting off the actual hard disk(s). To change the filesystem volume label on a device partition, issue the following command (substituting the proper device name):

```
tune2fs /dev/sda2 -L newlabelname
```

- It is possible to configure grub menus to boot correctly off multiple disk(s).

Normally, you must ensure that the grub entries contain the correct root LABEL or root partition information along with the correct disk order. For example, note the differences in “hd0” and “hd1” and “sda2” and “sdb2” below:

```
title RedHawk Linux 5.1.2-20080916 (Trace=Yes, Debug=No)
    root (hd0,0)
    kernel /vmlinuz-2.6.23.17-RedHawk-5.1.2-trace ro
root=/dev/sda2

title RedHawk Linux 6.0-20091211 (Trace=Yes, Debug=No)
    root (hd1,0)
    kernel /vmlinuz-2.6.31.6-RedHawk-5.4-trace ro
root=/dev/sdb2
```

In this case you would most likely be referencing sda in `/etc/fstab` on `/dev/sda2` and sdb in `/etc/fstab` on `/dev/sdb2`.

8.1.5 How to Install RHEL Remotely Over VNC

Virtual Network Computing (VNC) is a graphical desktop sharing system which uses the RFB protocol to remotely control another computer. It transmits the keyboard and mouse events from one computer to another, relaying the graphical screen updates in reverse direction, over a network.

VNC by default uses TCP ports 5900 through 5906, each port corresponding to a separate screen (:0 to :6).

If you are having difficulties installing in graphical mode and do not want to use text mode, VNC is an option. You need to be able to establish a network connection between the system you are installing (the VNC server) and the system that will remotely display the installation GUI (the VNC client).

Any system with a VNC viewer application can be used as a VNC client. If you are currently running a previously installed RedHawk system, the VNC client application is accessed under the gnome desktop cascading menu selection: **Applications -> Accessories -> VNC Viewer**, or by using the command `/usr/bin/vncviewer`.

Note that VNC installs have only been tested using another RedHawk system as the VNC client.

Perform the following steps to install using VNC.

On the system you wish to install (VNC server):

1. Follow the procedure outlined in section **5.0. Installation Procedures** on page 12 until you reach **5.1.1** step 4.

Select “Install or update an existing system” from the grub menu.

2. Hit the **Tab** key. Add “vnc” (and other options) to the end of the line that appears:

```
> vmlinux initrd=initrd.img vnc [options]
```

Hit the **Enter** key.

3. Continue with step 5 in section **5.1.1**.
4. After step 5, the VNC installer *may* prompt you with a text based GUI for extra setup. Use the **Tab** key to jump between fields, the up/down arrows within the fields, and the space bar to select your settings.

5. At this point you should see output similar to the following:

```
Starting VNC ...
```

```
WARNING!!! VNC server running with NO PASSWORD
You can use the vncpassword=<password> boot option
if you would like to secure the server.
```

```
Please manually connect your vnc client to
129.134.30.100:1 to begin the install.
```


```
Starting graphical installation.
```

6. If you wish to verify your network connection to the VNC client, press **Enter** and use **ifconfig** and **ping**. If you cannot ping your intended VNC client, you will have to troubleshoot your network configuration.
7. On the system that you will remotely display the installation (VNC client), start the the VNC Viewer.

On the gnome menu: Applications -> Accessories -> VNC Viewer

From the command line: **/usr/bin/vncviewer**

8. Establish the VNC connection to your VNC server by entering the *IPaddress:screen* in the VNC Viewer: Connection Details dialog box.



9. Select **OK** and press **Enter** to continue. At this point, you are able to perform a graphical installation remotely.

10. Continue with the graphical installation process as described in section **5.1.1** steps 6 - 21, beginning on page 15. After completing step 21 on page 24, perform step 10 below.

11. The reboot puts the system in runlevel 3 and starts the **firstboot** program that displays a text based menu called **Set Up Agent**.

Use the arrow keys to select **Firewall Configuration** and use the **Tab** key to select **Run Tool**, then press **Enter**.

NOTE

The **Set Up Agent** will time out if no action is taken within approximately one minute. If **Set Up Agent** times out, you will need to execute the following commands to restart it:

```
rm /var/lock/sybsys/firstboot
rm /etc/sysconfig/firstboot
service firstboot start
```

12. Use the space bar to select the following settings:

```
Set Security Level: Disabled
Set SELinux: Disabled
```

Tab to **[OK]** and press **Enter**.

13. From the **Set Up Agent** main menu, tab to **[Exit]** and press **Enter**.

After exiting the **Set Up Agent**, you should see **[OK]** returned by the **firstboot** service.

14. Edit **/etc/inittab** and change the first line:

```
from: id:3:initdefault:
to:   id:5:initdefault:
```

15. Reboot the system and continue with the installation steps.

8.1.6 If the Installation GUI Fails to Start

If during installation the GUI fails to start, it may be caused by graphics driver problems with certain VGA cards and/or flat panel monitors not being probed properly.

Specifying the following parameters at the boot prompt may solve your problem.

- Failure to recognize your VGA card or incompatible driver:

Specify the option **driver=xxx**, where **xxx** is the name of the driver. For example:

```
NVIDIA:      > vmlinux initrd=initrd.img driver=?v
ATI:         > vmlinux initrd=initrd.img driver=?adeon
Generic:     > vmlinux initrd=initrd.img driver=?esa
```

You can also specify resolution. For example, add: **resolution=1280x1024**

- The flat panel monitor may not be probed properly by the VGA card. Specify the native resolution, or if not known, try the following suggestions:

```

20" - 22" Flat Panels:  add: resolution=1600x1200
17" - 19" Flat Panels:  add: resolution=1280x1024
15" Flat Panels:       add: resolution=1024x768
24" Wide Flat Panels:  add: resolution=1920x1080
22" Wide Flat Panels:  add: resolution=1680x1050

```

You might have to specify more than one option. For example:

```
> vmlinux ... driver=vesa resolution=1280x1024
```

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-4.4.4** or later.

8.2.2 Hyper-thread Activation

Hyper-threading is configured by default in each of the standard RedHawk Linux kernels. 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.

9.0. Known Issues

Special consideration should be given to the following areas.

BIOS “Console Redirection”

The “Console Redirection” BIOS feature, when enabled, has been observed to interfere with the proper operation of the integrated VGA video and the Xorg 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_c?pability(8)** feature with **ssh**, the **/etc/ssh/sshd_config** file may need to 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.

irqbalance

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

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

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

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

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

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

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

Compatibility Issues for Ada programs

RedHawk 6.0 is based on RHEL Enterprise Linux 6.0. Binary and source compatibility issues exist for MAXAda programs between RHEL 3.0 and RHEL 6.0.

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.

Excessive ksoftirqd Runs Affecting Determinism

In more recent kernel.org kernels 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 determinism is an 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 more recent kernel.org kernels the default setting for the **sysctl** flag that allows ping to a multicast address has changed. Previously it was set to enable broadcast and multicast ICMP echo (pings) and timestamp requests. Now this functionality is disabled.

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., trace and debug kernels add “crashkernel=64M@48M,” 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.

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
```

Issues with Tigon3 Network Adapter

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

Crash Dump Support

On 32-bit systems with Physical Address Extension (PAE) enabled in the RedHawk kernels, the `vm -p` command will be unable to translate some virtual page addresses into their corresponding physical page addresses.

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 Software Portal. There are three ways of installing product updates:

- Using the Network Update Utility (NUU) installed on your RedHawk system
- Manual installation after browsing and downloading individual RPMs from Concurrent's software repositories
- Building a customized Update disc 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.

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

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

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

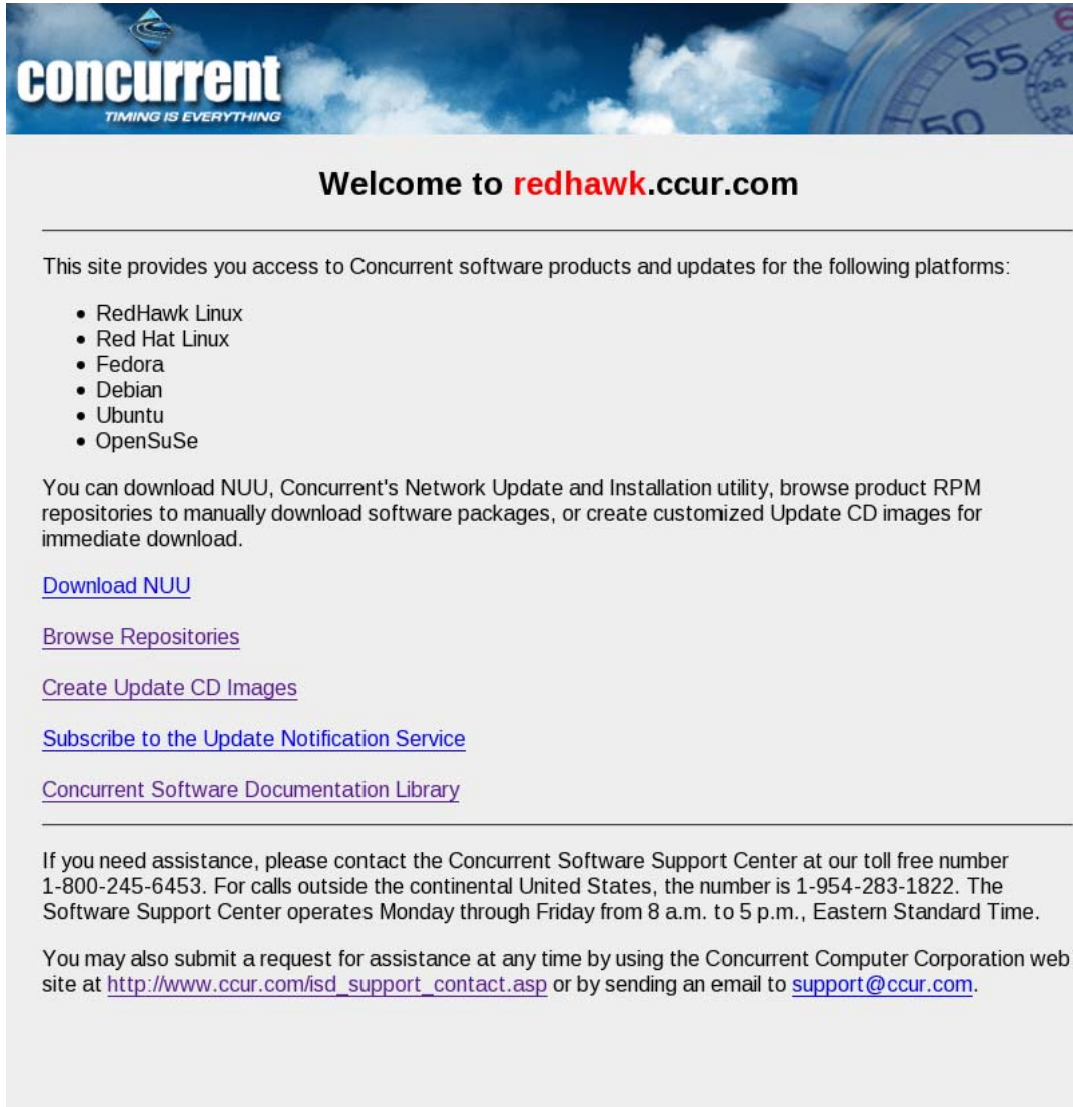
The first time you invoke NUU you will need to specify your redhawk.ccur.com Login ID and Password that were provided in the shipping documents accompanying your system. If you require assistance, refer to "Direct Software Support" on page 65.

Before using NUU to install any updated software modules, check for NUU updates separately. Apply any NUU updates and then restart NUU before applying any other updates.

10.2.2 Installing Manually Downloaded RPMs

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

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



concurrent
TIMING IS EVERYTHING

Welcome to **redhawk.ccur.com**

This site provides you access to Concurrent software products and updates for the following platforms:

- RedHawk Linux
- Red Hat Linux
- Fedora
- Debian
- Ubuntu
- OpenSuSe

You can download NUU, Concurrent's Network Update and Installation utility, browse product RPM repositories to manually download software packages, or create customized Update CD images for immediate download.

[Download NUU](#)

[Browse Repositories](#)

[Create Update CD Images](#)

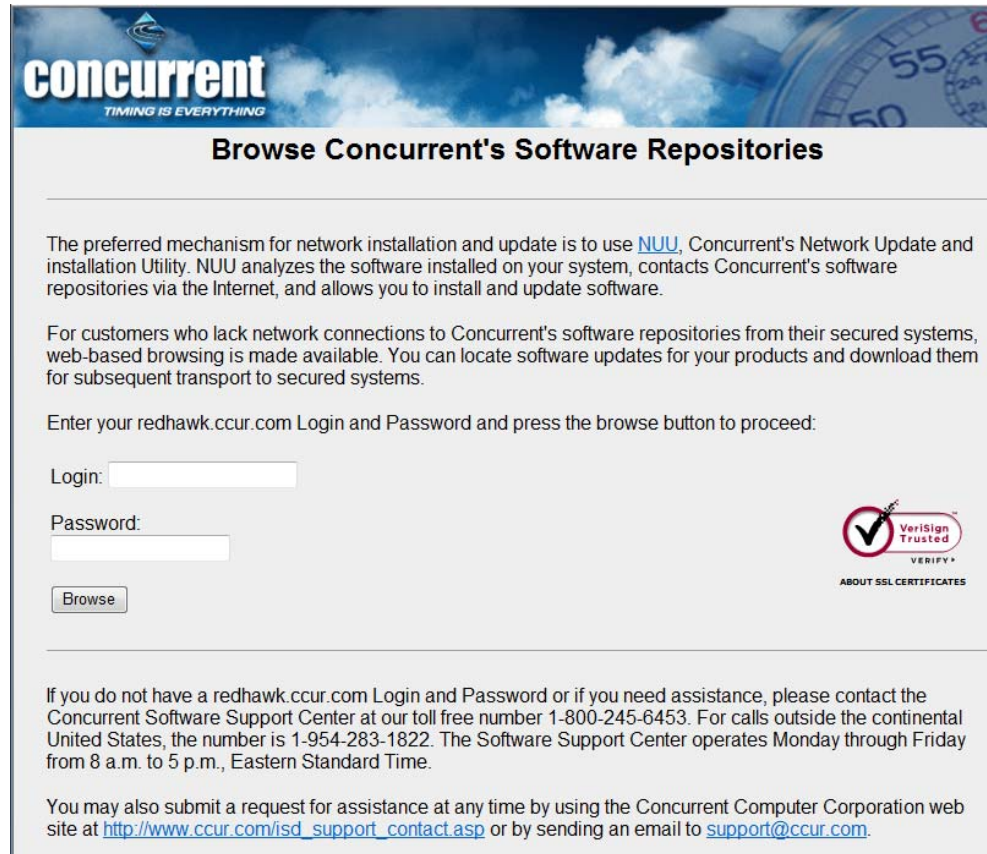
[Subscribe to the Update Notification Service](#)

[Concurrent Software Documentation Library](#)

If you need assistance, 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.

You may also submit a request for assistance at any time by using the Concurrent Computer Corporation web site at http://www.ccur.com/isd_support_contact.asp or by sending an email to support@ccur.com.

Clicking on the [Browse Repositories](#) link takes you to an authentication page.



concurrent
TIMING IS EVERYTHING

Browse Concurrent's Software Repositories


The preferred mechanism for network installation and update is to use [NUU](#), Concurrent's Network Update and installation Utility. NUU analyzes the software installed on your system, contacts Concurrent's software repositories via the Internet, and allows you to install and update software.

For customers who lack network connections to Concurrent's software repositories from their secured systems, web-based browsing is made available. You can locate software updates for your products and download them for subsequent transport to secured systems.

Enter your redhawk.ccur.com Login and Password and press the browse button to proceed:

Login:

Password:


VeriSign
Trusted
VERIFY
ABOUT SSL CERTIFICATES




If you do not have a redhawk.ccur.com Login and Password or if you need assistance, 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.

You may also submit a request for assistance at any time by using the Concurrent Computer Corporation web site at http://www.ccur.com/isd_support_contact.asp or by sending an email to support@ccur.com.

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.

Index of /home/redhawk/buffet/RedHawk/6.0

Name	Last modified	Size	Description
 Parent Directory/		-	
 i386/	10-May-2011 14:52	-	
 x86_64/	10-May-2011 14:52	-	

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 Discs

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

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

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

You will need to enter your redhawk.ccur.com Login ID and Password and then you can select the products to update. A disc 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>.