



# Version 5.4 Release Notes

December 2009

0898003-5.4



---

## 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 5.4	8
4.1.1	RedHawk Based on RHEL 5.4	8
4.1.2	Updated NVIDIA Support	8
4.1.3	New Pagemap Utility	8
4.1.4	Support For More Than 32 CPUs	8
4.1.5	Performance Counter Support	8
4.1.6	Support for KVM Virtualization	8
4.1.7	Pre-Installed NVIDIA CUDA Support	9
4.1.8	NUMA Support For 32-bit Systems	9
4.1.9	Kexec/Kdump Enhancements	9
4.1.10	NVIDIA Graphics Cards Now Use MSI Interrupts	9
4.1.11	Crash Dump Kernel Pre-Configured on Trace Kernels	9
4.1.12	New Per-CPU Migration Statistics	9
4.2	Modifications in Version 5.4	9
4.2.1	Critical Pthread Library Change	9
4.2.2	Run Command Output Change	10
4.2.3	TRACE_EVENT Macro Renamed	10
4.2.4	FBS Library Now Returns EINVAL On NULL	10
5.0	Installation Procedures	12
5.1	Installing Red Hat Software	12
5.1.1	Installing Red Hat	13
5.1.2	Installing Red Hat Updates	33
5.2	Installing RedHawk Linux	35
5.3	Installing the RCIM	37
5.3.1	Quick Hardware Installation Checklist	37
5.3.2	Verifying RCIM Operation	38
5.3.3	Installing NTP Updates for GPS Module	38
5.4	Installing Frequency-Based Scheduler (FBS) Software	39
5.5	Installing the PCI-to-VME Bridge	39
5.5.1	Installing the Hardware	39
5.5.2	Installing the Software	40
5.6	Installing Additional RedHawk Products	41
5.7	Installing RedHawk Updates from Disc	41

5.8	Installing RedHawk Updates from the Updates Web Site	42
6.0	Upgrade Procedures	43
6.1	Pre-Upgrade Procedure	43
6.2	Upgrading to RHEL5.4	44
6.3	Updating Red Hat Software	48
6.4	Upgrading to RedHawk Version 5.4	49
6.5	Upgrading to Frequency-Based Scheduler (FBS) Version 5.4	50
6.6	Upgrading to PCI-to-VME Bridge Version 5.4	51
6.7	Installing RedHawk Updates from Disc	51
6.8	Installing RedHawk Updates from the Updates Web Site	52
6.9	Additional Requirements	52
7.0	NVIDIA Graphics Configuration	54
7.1	Supported NVIDIA GPUs	54
7.2	Multiple NVIDIA Video Card Considerations	55
7.3	Display Configurations	55
7.3.1	Single	55
7.3.2	Xinerama	55
7.3.3	MultiDesktop	56
7.3.4	Twinview	56
7.3.5	Twinview-Xinerama	57
7.3.6	Twinview-MultiDesktop	57
7.4	Using nvidia-xconfig	57
7.5	Using nvidia-settings	58
7.6	Modified xorg.conf file generated by nvidia-xconfig	65
7.7	Troubleshooting Problems	66
8.0	Additional Information	69
8.1	Installation/Configuration Issues	69
8.1.1	File System Backup Recommendations	69
8.1.2	File System Swap Size Guidelines	70
8.1.3	RCIM Connection Modes	70
8.1.4	How to Use Linux Rescue	70
8.1.5	How to Install Using VNC	72
8.1.6	If the Installation GUI Fails to Start	75
8.2	Considerations	75
8.2.1	Compiler Requirements	75
8.2.2	Hyper-thread Activation	75
9.0	Known Issues	77
10.0	Software Updates and Support	83
10.1	Direct Software Support	83
10.2	Software Updates	83
10.2.1	Updating via NUU	83
10.2.2	Installing Manually Downloaded RPMs	84
10.2.3	Customized Update Discs	86
10.3	Documentation Updates	86

---

## 1.0. Introduction

This document provides release information and installation instructions for Concurrent Computer Corporation's RedHawk™ Linux® operating system Version 5.4

### 1.1 Product Description

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

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

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

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

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

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

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

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

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

## 1.2 Product Media

The tables below list the optical media discs that are delivered with RedHawk Linux Version 5.4 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

<b>Processor Type:</b>	<b>Intel® Xeon™ i386 and EM64T 32-bit</b>
<b>Model:</b>	<b>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]</b> <b>iHawk 880 Series [ HQHS2 ]</b> <b>ImaGen [ HQ0G1, HR0G1 ]</b>
<b>Operating System Discs:</b>	
<b>Red Hat Enterprise Linux 5.4 (32-bit) (1 DVD)</b> <b>Red Hat Enterprise Linux 5.4 Source (32-bit) (1 DVD)</b> <b>Red Hat Enterprise Linux 5.4 Updates (32-bit) (1 CD) [LXUPDATE WA-IA32 Version 5.4]</b> <b>RedHawk Linux OS (32-bit) (1 CD) [WA-IA32 Version 5.4]</b>	
<b>Optional Product Discs:</b>	
<b>RedHawk Linux Frequency-Based Scheduler (32-bit) [WU1210-JL Version 5.4]</b> <b>NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version x.x (x.x=latest version)]</b> <b>RedHawk Linux PCI-to-VME Bridge Software Library for i386 &amp; EM64T Systems [WU-HS630-LIB (WCS-SBS-620)] Version 5.4</b> <b>RedHawk Linux Cluster Manager (32-bit) [ WA9017-L] Version 5.4</b> <b>RedHawk Global File System for i386 &amp; EM64T Systems [ WA9018-L] Version 5.4</b> <b>RedHawk High Availability NFS for i386 &amp; EM64T Systems [ WA9019-L] Version 5.4</b>	

## 64-Bit Product Media

<b>Processor Type:</b>	<b>AMD Opteron™ 64-bit and EM64T x86_64 64-bit</b>
<b>Model:</b>	<b>iHawk 860 Series [ HQ685, 680, 285, 280, 185, 069, 067, 049, 047, T74, T54, R54, R90, 000 ]</b> <b>iHawk 870 Series [ HR844, 430, 221, 210, 00W, 00T ]</b> <b>iHawk 880 Series [ HQHS2 ]</b> <b>ImaGen [ HQ0G1, HR0G1 ]</b>
Operating System Discs:	
<b>Red Hat Enterprise Linux 5.4 (64-bit) (1 DVD)</b> <b>Red Hat Enterprise Linux 5.4 Source (64-bit) (1 DVD)</b> <b>Red Hat Enterprise Linux 5.4 Updates (64-bit) (1 CD) [LXUPDATE WA-AMD64 Version 5.4]</b> <b>RedHawk Linux OS (64-bit) (1 CD) [WA-EM64T-AMD64 Version 5.4]</b>	
Optional Discs:	
<b>RedHawk Linux Frequency-Based Scheduler for x86_64 Systems [WU1210-JA Version 5.4]</b> <b>NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version x.x (x.x=latest version)]</b> <b>RedHawk Linux PCI-to-VME Bridge Software Library for x86_64 Systems [WU-HS630-LIB (WCS-SBS-620)] Version 5.4</b> <b>RedHawk Linux Cluster Manager (64-bit) [ WA9017-L] Version 5.4</b> <b>RedHawk Global File System for x86_64 Systems [ WA9018-L] Version 5.4</b> <b>RedHawk High Availability NFS for x86_64 Systems [ WA9019-L] Version 5.4</b>	

### 1.3 Version Numbers

All RedHawk Linux version numbers take the following form:

$$major.minor[.update]$$

where:

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

For example, 5.1 is the initial release of major version 5, and 5.2 is a subsequent release to version 5.1. Both 5.1 and 5.2 constitute complete product releases and do not require a previous release to be currently installed on the system; both are available on 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 5.1.1 kernel update can only be installed on a system currently running RedHawk Linux version 5.1.

## 1.4 System Updates

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

### NOTE

Concurrent does not recommend downloading Red Hat updates.

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

Instructions for downloading updates from the website are provided in the section "Software Updates" on page 83.

## 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> .
<b>list bold</b>	User input appears in <b>list bold</b> type and must be entered exactly as shown. Names of directories, files, commands, options and man page references also appear in <b>list bold</b> 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 5.4 as PDF files. After installing RedHawk Linux documentation, the files are available for viewing:

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

**Table 2-1 RedHawk Linux Documentation**

<b>RedHawk Linux Operating System Documentation</b>	<b>Document Number</b>
<i>RedHawk Linux Online Documentation Roadmap</i>	0898002-540
<i>RedHawk Linux Version 5.4 Release Notes</i>	0898003-5.4
<i>RedHawk Linux User’s Guide</i>	0898004-620
<i>Real-Time Clock &amp; Interrupt Module (RCIM) User’s Guide</i>	0898007-400
<i>iHawk Optimization Guide</i>	0898011-110
<i>RedHawk Linux FAQ</i>	N/A
<b>Optional RedHawk Product Guides</b>	
<i>RedHawk Linux Frequency-Based Scheduler (FBS) User’s Guide</i>	0898005-310
<i>RedHawk High Availability NFS (HA-NFS) Installation Guide</i>	0898018-100
<i>RedHawk Global File System (GFS) Installation Guide</i>	0898020-100
<b>Optional RedHawk Product Documentation Sets</b>	
<i>RedHawk Linux Cluster Manager</i>	
<i>PCI-to-VME Bridge</i>	

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

The latest documentation for all Concurrent software products, including updated Release Notes and User Guides, can be accessed from Concurrent’s documentation web site at <http://redhawk.ccur.com/docs>. An updated RedHawk FAQ can be accessed at <http://www.ccur.com/isdfaq/RedHawk-FAQ.pdf>.

Printed copies can be ordered by contacting the Concurrent Software Support Center. Refer to “Direct Software Support” on page 83 for details.

## 2.2 Vendor Documentation

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

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

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

---

---

## 3.0. Prerequisites

### 3.1 Software

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

- Red Hat Enterprise Linux 5.4 with updates
- RedHawk Linux operating system
- NightStar RT for RedHawk tools (optional)
- RedHawk Linux Frequency-Based Scheduler (FBS) (optional)
- RedHawk Linux PCI-to-VME bridge software (optional)
- RedHawk Linux Cluster Manager (optional)
- RedHawk High Availability NFS (HA-NFS) (optional)
- RedHawk Global File System (GFS) (optional)

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

#### NOTE

Real-time extensions and features are *not* available to 32-bit applications running under a 64-bit operating system (i.e. x86\_64). In order to use real-time features, migrate 32-bit applications to 64-bit (described in Appendix E of the *RedHawk Linux User's Guide*) or boot a 32-bit operating system instead.

### 3.2 Hardware

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

## 4.0. Changes in this Release

This section describes enhancements and other changes introduced in Version 5.4.

### 4.1 Enhancements in Version 5.4

#### 4.1.1 RedHawk Based on RHEL 5.4

RedHawk Version 5.4 uses Red Hat Enterprise Linux 5.4 as its base with a merge of kernel.org version 2.6.31.

#### 4.1.2 Updated NVIDIA Support

RedHawk Version 5.4 includes the NVIDIA Linux Display Driver version 190-42. Legacy driver support is no longer included but is available for download from the NVIDIA website. Refer to "NVIDIA Graphics Configuration" on page 54 for details.

#### 4.1.3 New Pagemap Utility

RedHawk 5.4 now includes a new utility called **pagemap(1)** that may be used to view page mapping or swap information for all or a subset of user space translations within a specified process's address space. For pages that are currently resident it displays the page frame number, NUMA node ID, number of mappings to the page and page flag information. Refer to the **pagemap(1)** man page for more information.

#### 4.1.4 Support For More Than 32 CPUs

RedHawk 5.4 is the first release of the new FBS support that allows more than 32 CPUs to be specified for CPU affinity. Support for more than 32 CPUs includes a revised Frequency Based Scheduler User's Guide document, with more than 20 new services documented and described.

#### 4.1.5 Performance Counter Support

RedHawk 5.4 provides kernel support for the new "Performance Counter" subsystem. This support provides a generic mechanism for counting key hardware and software events that occur during a specified range in time. Tools and documentation for Performance Counters are included with the RedHawk 5.4 product. See the **perf(1)** man page for more information.

#### 4.1.6 Support for KVM Virtualization

RedHawk 5.4 is the first release of the KVM virtualization solution for RedHawk. KVM (for Kernel-based Virtual Machine) is a full virtualization solution for Linux on x86 hardware containing virtualization extensions. Using KVM, one can run multiple virtual machines running unmodified Linux or Windows images. Each virtual machine has private virtualized hardware: a network card, disk, graphics adapter, etc. For more information, see the README file in `/usr/share/doc/ccur/kvm` and visit <http://www.linux-kvm.org> for much more documentation.

### 4.1.7 Pre-Installed NVIDIA CUDA Support

RedHawk 5.4 includes pre-installed NVIDIA CUDA libraries and header files. In addition, example code is provided. See `/usr/share/doc/ccur/examples/cuda` for more information.

### 4.1.8 NUMA Support For 32-bit Systems

Concurrent's NUMA extensions and enhancements are now available on 32-bit systems, in addition to 64-bit systems. This allows 32-bit systems to take advantage of any hardware NUMA support that is available, and this can be combined with memory shielding to guarantee that a process's pages are all local to a specific NUMA node.

### 4.1.9 Kexec/Kdump Enhancements

The kexec/kdump support used for creating system crash dump files has been enhanced to provide better reliability and flexibility for a wider variety of systems.

### 4.1.10 NVIDIA Graphics Cards Now Use MSI Interrupts

The `ccur-vidia-glx` rpm installation now updates the `/etc/modprobe.conf` file with the following line:

```
options nvidia NVreg_EnableMSI=1
```

This will make the NVIDIA driver use an MSI (*message signalled interrupt*) instead of the PCI INTA interrupt, which may be shared with other devices in the system. Setting this option to a value of zero (0), or removing this line will make the NVIDIA driver revert back to using a PCI INTA interrupt. The driver must be unloaded and reloaded, or the system must be rebooted, in order for a change in interrupt use to occur.

### 4.1.11 Crash Dump Kernel Pre-Configured on Trace Kernels

In previous releases of RedHawk, only the debug kernel was pre-configured with a crash dump kernel, however in RedHawk 5.4 both the trace and debug kernels are pre-configured with a crash dump kernel. Note that the following grub boot option can be added to any kernel to enable a crash dump kernel:

```
crashkernel=64M@48M
```

### 4.1.12 New Per-CPU Migration Statistics

RedHawk 5.4 now keeps track of the number of process migrations to each CPU since the machine was booted. This statistic is available by examining the contents of the `/proc/ccur/switches` file.

## 4.2 Modifications in Version 5.4

### 4.2.1 Critical Pthread Library Change

Starting with RedHawk 5.4, the application writer must use `pthread_mutexattr_setshared(3)` on their shared POSIX mutexes, and `pthread_condattr_setshared(3)` on their shared POSIX condition variables.

For private variables, the services mentioned are optional. Note that a shared variable is one which different processes access, and a private variable is one which only multiple threads of a single process access.

Failure to correctly mark shared variables as shared will cause silent corruption in glibc behavior, resulting in an eventual user-space stall, in glibc itself, in one or more of the applications referencing the shared mutex or condition variable.

#### 4.2.2 Run Command Output Change

The `run` command `--info` output now contains a CPU field, which displays the CPU that the process or thread is running on (or, if it is not running, it shows the last CPU that the process or thread ran on) at the time of the `run` command invocation.

#### 4.2.3 TRACE\_EVENT Macro Renamed

The `TRACE_EVENT` macro has been renamed to `TRACE_LOG_EVENT` in RedHawk 5.4. This was necessary because of a macro name conflict that the 2.6.31 kernel introduced.

#### 4.2.4 FBS Library Now Returns EINVAL On NULL

Many FBS library routines accept a *name* or a *name\_ptr* argument. In previous releases of RedHawk, passing a NULL pointer into the *name* or *name\_ptr* argument resulted in a program SEGFAULT exception occurring.

In RedHawk 5.4, these FBS library routines have been changed to properly detect a NULL pointer and simply return `EINVAL` instead of producing a SEGFAULT exception.

The following FBS library routines have been modified with this new behavior:

- `fbsinfo(3)`
- `fbsinfo_big(3)`
- `pgmremove(3)`
- `pgmremove_big(3)`
- `pmclrpgm_big(3)`
- `pmprogram(3)`
- `pmprogram_big(3)`
- `pmqrypgm_big(3)`
- `sched_pgm_deadline_query(3)`
- `sched_pgm_deadline_query_big(3)`
- `sched_pgm_deadline_test(3)`
- `sched_pgm_deadline_test_big(3)`
- `sched_pgm_set_deadline(3)`
- `sched_pgm_set_deadline_big(3)`
- `sched_pgm_set_soft_overrun_limit_big(3)`
- `sched_pgm_soft_overrun_query(3)`
- `sched_pgm_soft_overrun_query_big(3)`
- `sched_pgmadd(3)`
- `sched_pgmadd_big(3)`
- `sched_pgmqry(3)`
- `sched_pgmqry_big(3)`
- `sched_pgmresched_big(3)`

Note that this change applies only to the FBS C library. The FBS Fortran library has not been modified for this release.

---

---

## 5.0. Installation Procedures

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

**If upgrading from Version 5.1 or 5.2 to Version 5.4, follow the steps in “Upgrade Procedures” on page 43.**

### Review the following points before installing the software:

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

### Before you start:

- **Main Memory Size** — You need to know the system’s main memory size in order to properly set up the file system swap size. If you are not sure of the main memory size, access the system’s BIOS to ascertain this information.
- **Network Addresses** — If you do not plan to configure networking DHCP, you will need to have the following information for your system before beginning the installation:
  - IP address
  - 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 anaconda installer instead of from disc.

## 5.1 Installing Red Hat Software

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

- Installing Red Hat
- Installing Red Hat Updates



### 5.1.1 Installing Red Hat

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

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



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

#### NOTES

##### If 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 75 for some boot options that may cause the GUI to launch. Some NVIDIA cards require specifying the display resolution at the boot prompt before the install GUI will start. If successful, continue to follow the installation directions below.

In addition to GUI mode, Red Hat Enterprise Linux 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 72 for full instructions.

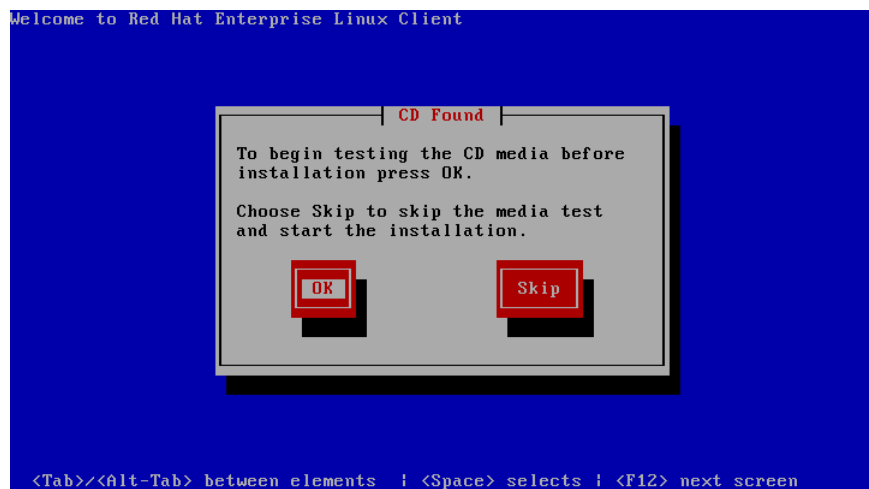
**If your system contains disk drives or devices that are not supported by the Red Hat installation program**, you will need to obtain the necessary driver(s) in order to install Red Hat Enterprise Linux. A diskette containing the driver may be included with your system or it may be available on the Red Hat installation media. If not available, contact the driver manufacturer for an updated driver or call Concurrent Support (see "Direct Software Support" on page 83).

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 Red Hat media, press **Enter**, otherwise, navigate to the **Skip** selection using the **Tab** key and press **Enter** to immediately start the installation.

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



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



- Select your desired language and keyboard layout on the next two screens, pressing Next at each to proceed.
- After selecting your keyboard layout, a dialog will appear which requests an installation number. Select the Installation Number radio button, enter `6d70-bd33-6d3c-fcbc` in the text field, and click the OK button.



- The dialog as shown below will appear.

**If you are installing Red Hat**, select the Install Red Hat Enterprise Linux Client radio button and click Next.

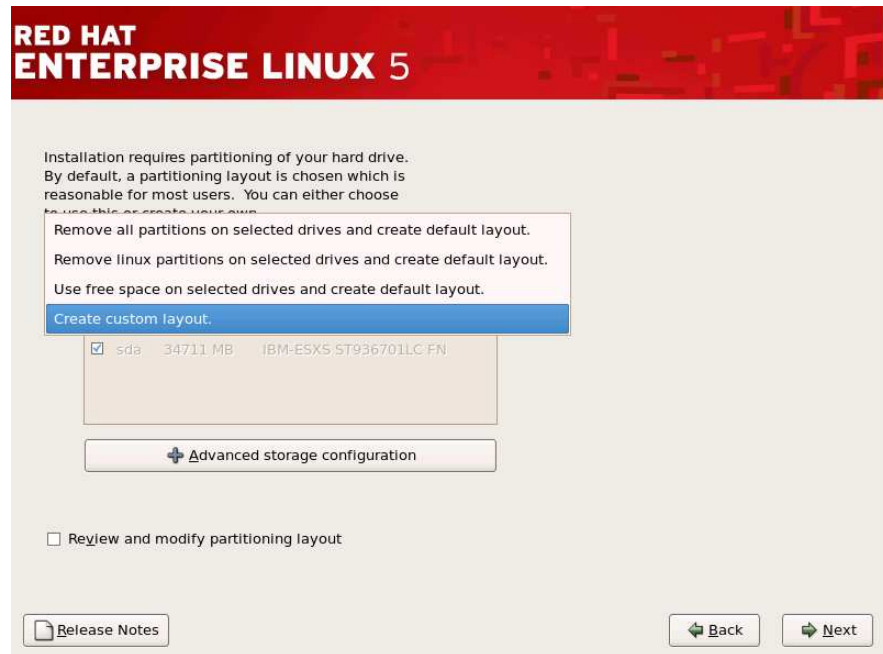
**If you are upgrading from RedHawk Version 5.1 or 5.2**, select “Upgrade an existing Installation” and go to page 44 for further instructions.



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

### WARNING

Do *not* use any Red Hat **default** partitioning scheme. If used, you will not be able to install RedHawk Linux. In particular RedHawk does not support booting from a Logical Volume Manager (LVM) partition.

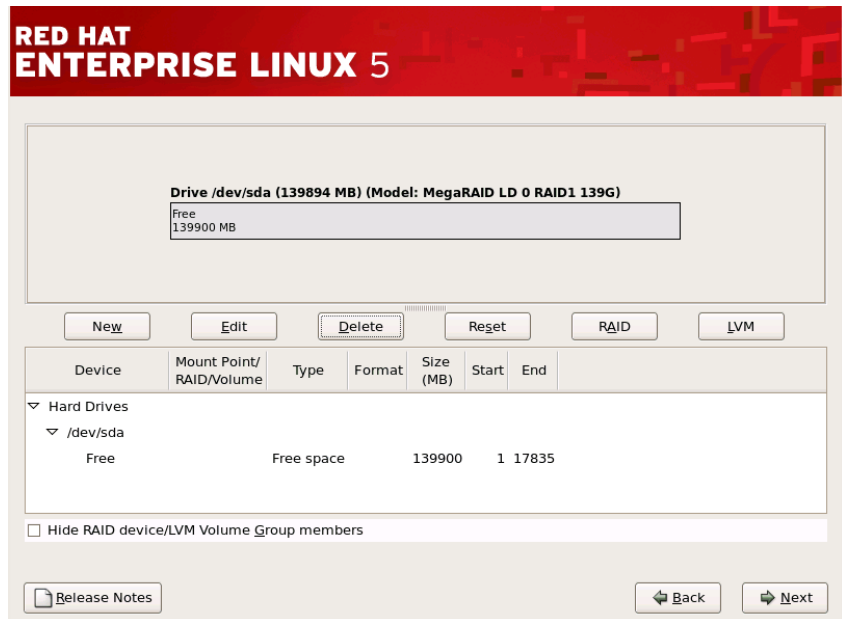


### NOTE

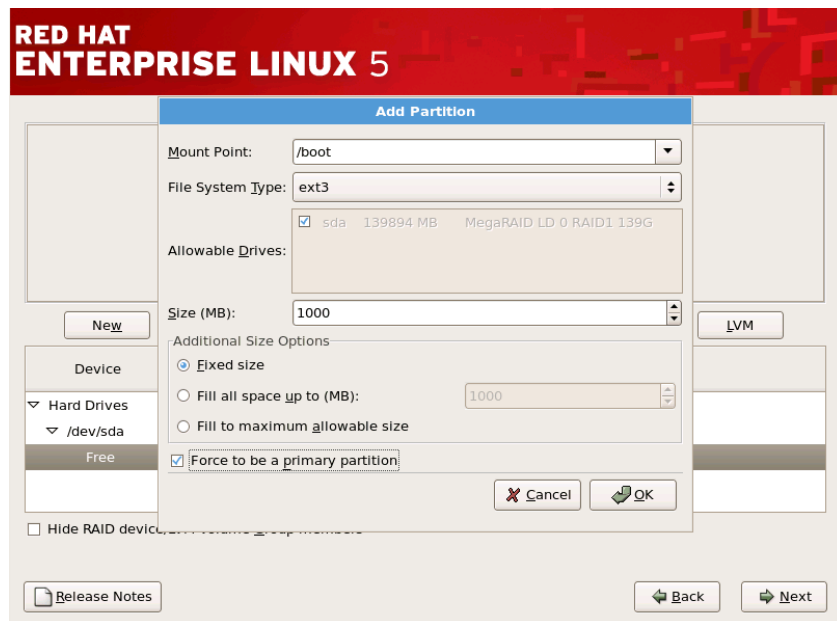
If you encounter a problem because the installer cannot read the partition table on your disk, select **Create custom layout**, click on **Review and modify partitioning layout**, then click the **Next** button. Click **OK** on any warning dialogs that may appear.

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

**WARNING:** This action will erase all data on the disk when the disk is formatted in the next step.



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

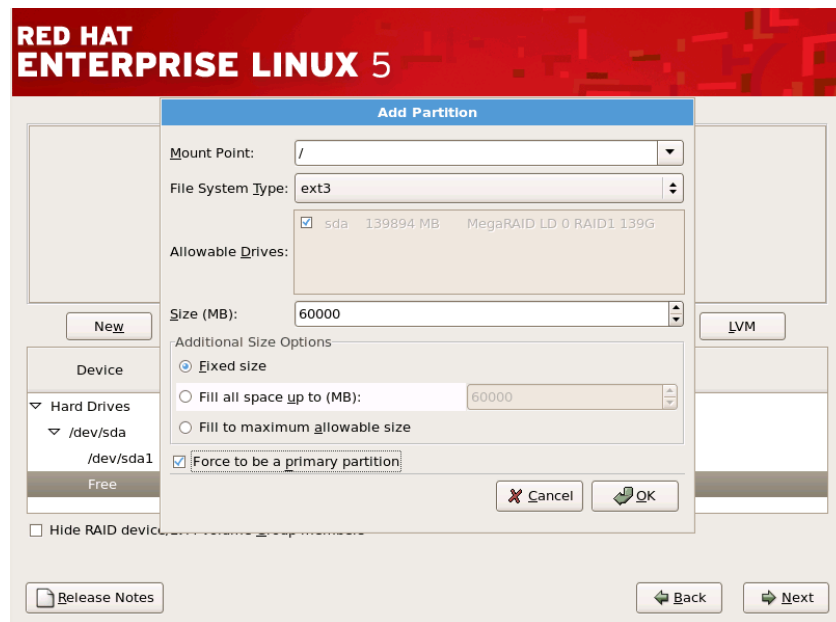


- c. Click the **New** button to add the root partition. Type `/` into the Mount Point text field, select `ext3` as the File System Type.

Select the Fixed size radio button and check the Force to be a primary partition checkbox.

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

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

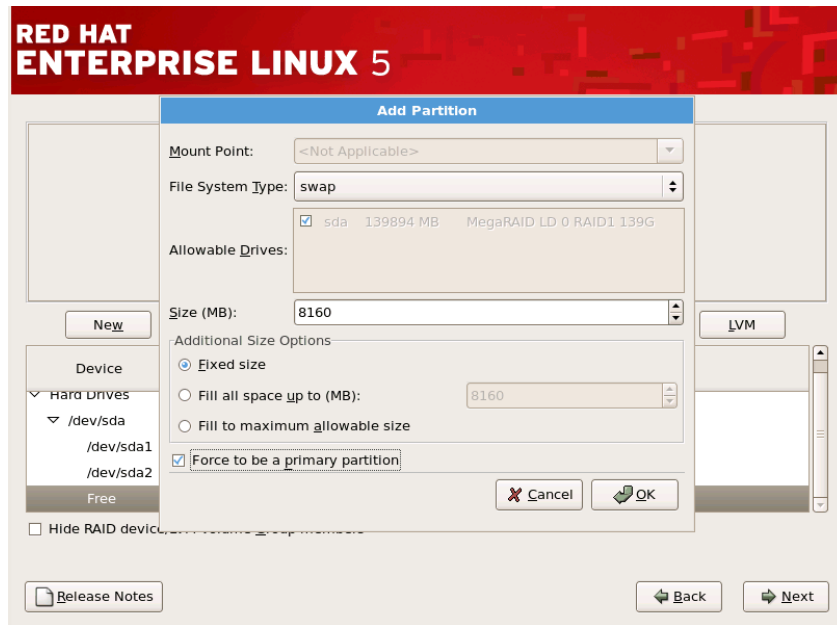


- d. Click the **New** button to add the swap partition. Select **swap** as the **File System Type**. Select the **Fixed size** radio button and check the **Force to be a primary partition** checkbox.


Select an appropriate size for the swap partition -- you should specify a size which matches the amount of physical memory you have in your system, plus 5% (refer to "File System Swap Size Guidelines" on page 70).

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

Click the OK button.



- e. Unless you used the entire contents of the disk in step c above, configure a `/home` partition for general use which uses the remaining free space on the disk.
- f. The disk partitioning dialog should look similar to the following, except for device names and sizes as selected in the previous steps.



Click Next to continue.



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

**RED HAT ENTERPRISE LINUX 5**

The GRUB boot loader will be installed on /dev/sda.  
 No boot loader will be installed.

You can configure the boot loader to boot other operating systems. It will allow you to select an operating system to boot from the list. To add additional operating systems, which are not automatically detected, click 'Add.' To change the operating system booted by default, select 'Default' by the desired operating system.

Default	Label	Device
<input checked="" type="checkbox"/>	Red Hat Enterprise Linux Client	/dev/sda2

A boot loader password prevents users from changing options passed to the kernel. For greater system security, it is recommended that you set a password.

Use a boot loader password

Configure advanced boot loader options

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

**RED HAT ENTERPRISE LINUX 5**

**Edit Interface**

**Broadcom Corporation NetXtreme BCM57045 Gigabit Ethernet**  
**Hardware address: 00:14:5E:3D:BB:C6**

Enable IPv4 support  
 Dynamic IP configuration (DHCP)  
 Manual configuration  
 IP Address: 129.134.30.200 / Prefix (Netmask): 255.255.255.0

Enable IPv6 support  
 Automatic neighbor discovery  
 Dynamic IP configuration (DHCPv6)  
 Manual configuration  
 IP Address: / Prefix:

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

- b. For Manual configuration, enter your IP Address and Net-mask in the text fields.
- c. Click OK to close the Edit dialog.
- d. The Networking screen now shows your configuration. If you choose Manual configuration, you also need to fill in the host name next to the manually text field, and the Gateway, Primary DNS and Secondary DNS fields in the Miscellaneous Settings section of the Networking dialog. Make any final required changes.
- e. Follow the same procedure for any other network devices that are shown in the list and which you wish to be enabled.
- f. Click Next to continue.



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

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



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



16. In addition to selecting the major sub-sections in each category, you can add additional optional packages to the installation. It is recommended that you select every package available in **every** sub-section, **except** in the **Language** category (explained in the next step).

**In GUI mode:**

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

To add multiple selections in a sub-section, select the first sub-section item in the list by clicking on it. Then go to the bottom of the list (you may need to scroll down), and click on the last sub-section while holding the **Shift** key down. This selects all sub-sections. Right-click and choose **Select all optional packages** from the context menu.

**In text mode:**

Use the instructions at the bottom of the screen to navigate. Press **F2** to open each sub-section and press the space bar to add packages to the packages menu.

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

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

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

If you do require additional optional packages, you may attempt to install them manually after all Red Hat updates have been installed, but it is likely that the Red Hat Installation media does not have the required (updated) versions. Contact Concurrent Software Support, as described in "Direct Software Support" on page 83, for aid in obtaining and installing such packages.

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



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



20. After package dependencies have been completed, installation begins. No action is required from you during this stage.

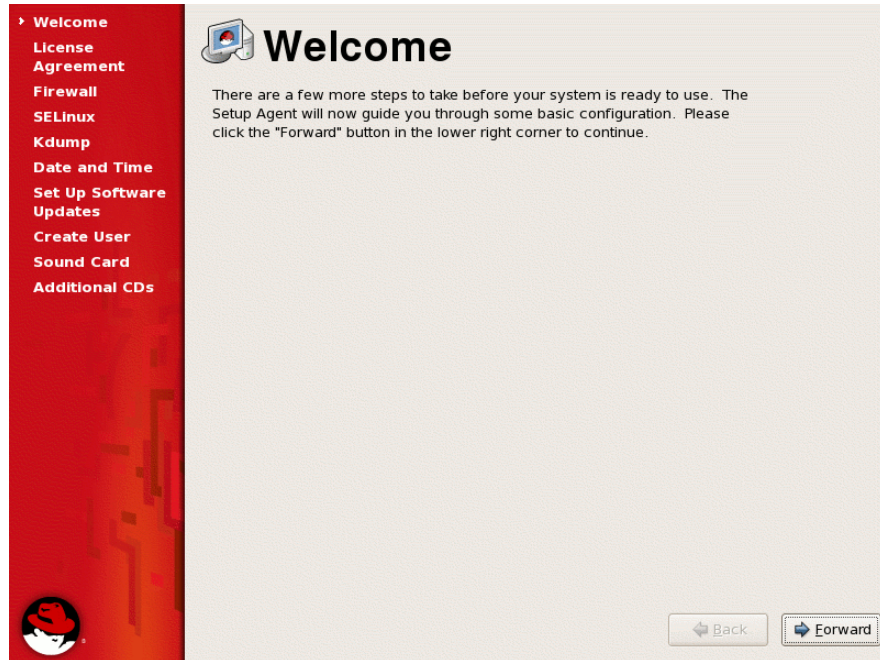


21. When installation is complete, the following dialog appears. Remove the media from the DVD device and click **Reboot** to continue.



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

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




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

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



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






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



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

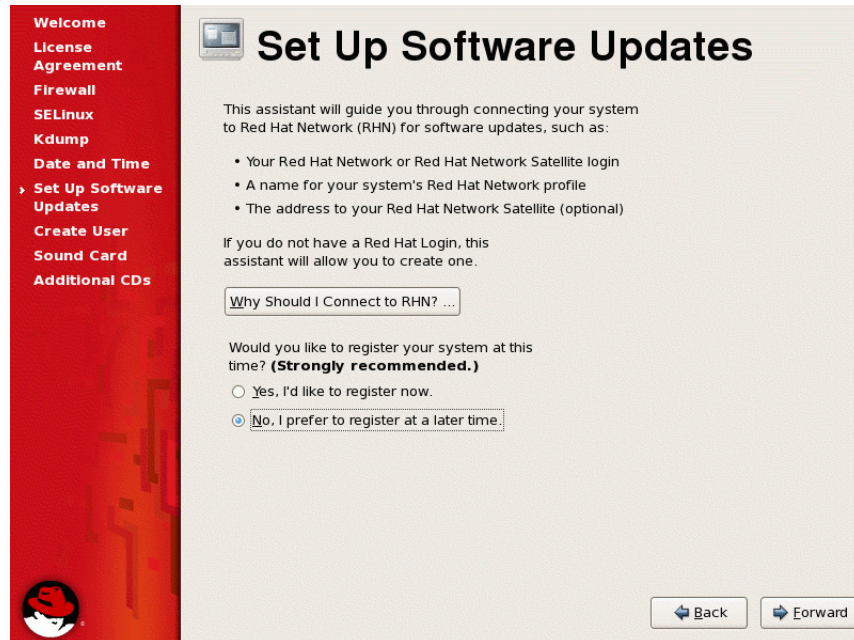


28. The next screen controls the configuration of Kdump. Ensure that the **Enable Kdump** checkbox is not set. The installation of RedHawk in a subsequent step installs an improved system dump utility which will be automatically configured. Click **Forward**.

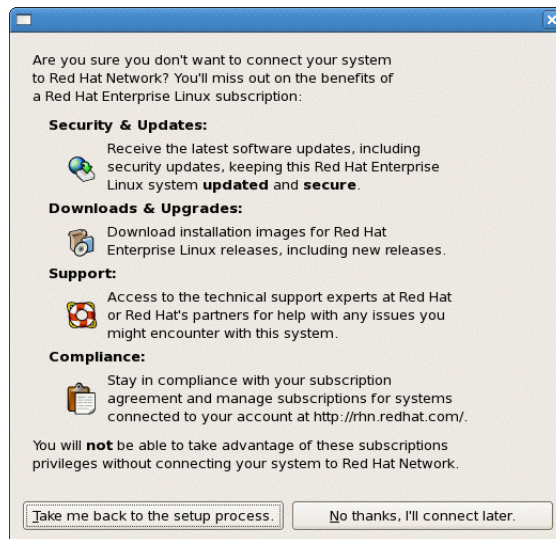


29. On the next screen, set the date and time and configure the Network Time Protocol daemon if desired.
30. The next screen asks you to register with Red Hat for software updates. Select the **No, I prefer to register at a later time** radio button. Software updates to the Red Hat Installation software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


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



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



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



- Optionally configure a user, NIS, and your sound card on the next two screens.
- The final screen will prompt you for additional CDs. Do not click the Install... button -- the remaining software will be installed in a subsequent step. Click Finish, and then subsequently OK in the dialog which indicates a reboot is required.



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

## 5.1.2 Installing Red Hat Updates

Install the disc labeled RHEL 5.4 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
```

3. 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
```

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

**NOTE**

If you cannot execute the **install-updates** script (it fails with Permission Denied), or if you attempt to execute it by launching it from a file browser and it immediately returns, the 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:

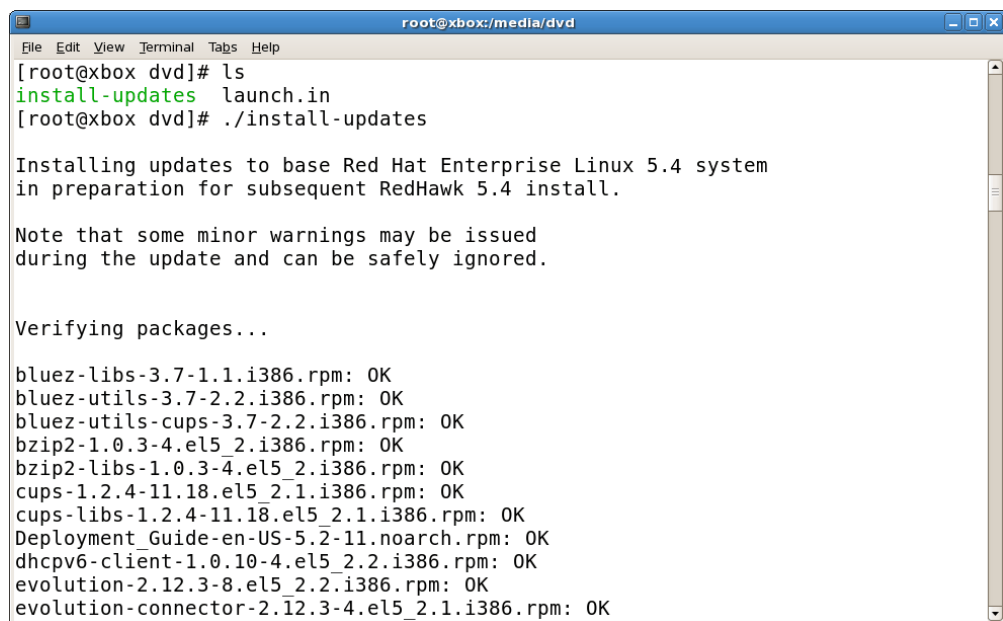
```
mount -o remount,exec mount-point
```

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/RHEL5.4-Updates type iso9660 (ro,noexec)

> mount -o remount,exec /media/RHEL5.4-Updates
```

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



```
root@xbox:/media/dvd
File Edit View Terminal Tabs Help
[root@xbox dvd]# ls
install-updates launch.in
[root@xbox dvd]# ./install-updates

Installing updates to base Red Hat Enterprise Linux 5.4 system
in preparation for subsequent RedHawk 5.4 install.

Note that some minor warnings may be issued
during the update and can be safely ignored.

Verifying packages...

bluez-libs-3.7-1.1.i386.rpm: OK
bluez-utils-3.7-2.2.i386.rpm: OK
bluez-utils-cups-3.7-2.2.i386.rpm: OK
bzip2-1.0.3-4.el5_2.i386.rpm: OK
bzip2-libs-1.0.3-4.el5_2.i386.rpm: OK
cups-1.2.4-11.18.el5_2.1.i386.rpm: OK
cups-libs-1.2.4-11.18.el5_2.1.i386.rpm: OK
Deployment_Guide-en-US-5.2-11.noarch.rpm: OK
dhcpcv6-client-1.0.10-4.el5_2.2.i386.rpm: OK
evolution-2.12.3-8.el5_2.2.i386.rpm: OK
evolution-connector-2.12.3-4.el5_2.1.i386.rpm: OK
```

Dependency checking is performed, and finally, updating of packages is executed. No action is required of you during this step.

```

root@xbox:/media/dvd
File Edit View Terminal Tabs Help
Updating yum-skip-broken

Installing 78 RPMs...

Preparing... ##### [ 100%]
 1:libxml2 ##### [ 1%]
 2:initscripts ##### [ 3%]
 3:perl ##### [ 4%]
 4:net-snmp-libs ##### [ 5%]
 5:cups-libs ##### [ 6%]
 6:libxslt ##### [ 8%]
 7:openoffice.org-core ##### [ 9%]
 8:bluez-libs ##### [ 10%]
 9:hpijs ##### [ 12%]
10:libxml2-python ##### [ 13%]
11:evolution ##### [ 14%]
12:openssh ##### [ 15%]
13:libpurple ##### [ 17%]
14:yum ##### [ 18%]
15:selinux-policy ##### [ 19%]
16:samba-common ##### [ 21%]
17:bzip2-libs ##### [ 22%]
18:openoffice.org-pyuno ##### [ 23%]
19:cups ##### [ 24%]

```

4. 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**.

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

## 5.2 Installing RedHawk Linux

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

1. Log in as root and take the system down to single-user mode:
  - a. If in GUI mode, right click on the desktop and select Open Terminal.
  - b. At the system prompt, type **init 1**.
2. Locate the disc labeled “RedHawk Linux Version 5.4” 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.

- 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@wario:/media/cdrom
File Edit View Terminal Tabs Help
[root@wario cdrom]#
[root@wario cdrom]# ls
documentation  install          launch.in  nvidia      uninstall
extras          install-redhawk launch.un  RedHawk-eula  uninstall-redhawk
[root@wario cdrom]# ./install-redhawk

Installing RedHawk Linux 5.4 (Lima)

This script will install the RedHawk 5.4 release on
the current system (wario.ccur.com).

Do you want to install RedHawk Linux 5.4? [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 /etc/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: 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.
- 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.
  - At the system prompt, type `init 5`. Proceed through the windows that display to configure screen resolution and color depth for your system. When you exit this session, the RedHawk Linux desktop should appear.



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

## 5.3 Installing the RCIM

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

### 5.3.1 Quick Hardware Installation Checklist

1. Before installing the RCIM, determine if you will be using the RCIM to accept or deliver external interrupts and the mode in which the RCIM will run (see "RCIM Connection Modes" on page 70).

#### NOTE

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

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

#### NOTE

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

4. Open the case of your system and identify the PCIe slot (RCIM III) or PCI slot (RCIM II or RCIM I) where you want the RCIM to reside. In general, it is best for the RCIM to be configured in a slot where minimal or no contention with other devices occurs and at the highest IRQ priority possible. For more information, refer to the *iHawk Optimization Guide*, publication number 0898011. The PDF file for this manual is available for viewing by clicking on the "Documents" icon on the desktop.
5. Insert the RCIM into the selected PCI slot and lock it down.
6. If this is to be part of an RCIM chain, attach the synchronization cable as required (see the *Real-Time Clock and Interrupt Module (RCIM) User's Guide* manual for complete details).

7. If you have the optional GPS module, attach the antenna lead and mount the antenna. The antenna should be mounted on the rooftop or in an open area.
8. Close the case and reconnect all power cords.
9. Power up the system and verify the system boots up properly. Refer to the RCIM guide for configuration options.

### 5.3.2 Verifying RCIM Operation

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

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

You should see an output similar to that shown below:

```
RCIM-III board 0 is at revision 6 eeprom 2.0.  
This is a standalone (isolated) rcim.  
Has IRQ 19 and major number 253  
GPS present, RCIM firmware version 1
```

#### NOTE

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

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

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

You should see output similar to the following:

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

### 5.3.3 Installing NTP Updates for GPS Module

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

1. The **ccur-ntp** RPM is available via Concurrent's Software Repository system. Execute the following steps to install this RPM:
  - a. Remove (erase) the existing **ntp**. Because there are some rpms which are normally installed which depend on **ntp**, use the **--nodeps** option to bypass the dependency test. This will save the existing **/etc/ntp.conf** file as **/etc/ntp.conf.rpmsave**.

```
# rpm --nodeps -e ntp
```
  - b. Download the **ccur-ntp** RPM from Concurrent's web site as described in section 10.2.2 on page 84.

- c. Install the RPM via the following command:

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

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

## 5.4 Installing Frequency-Based Scheduler (FBS) Software

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

1. With RedHawk Linux Version 5.4 running, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk Linux Frequency-Based Scheduler Version 5.4" 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 VMEbus adapter and a fiber optic cable.

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

1. To configure the jumpers on the VMEbus adapter card correctly, refer to Chapter 10 of the *SBS Technologies Model 618-3 Hardware Manual* included with the RedHawk Linux documentation. The PDF

file for this manual is available for viewing by clicking on the “Documents” icon on the desktop.

2. To install the PCI adapter in your iHawk system:
  - a. Power down the system.
  - b. Locate a vacant PCI card slot in the chassis that supports a bus master.
  - c. Remove the metal plate that covers the cable exit at the rear of the chassis.
  - d. Insert the PCI adapter card into the connector.
  - e. Fasten the adapter card in place with the mounting screw.
  - f. Replace the cover.
3. To install the VMEbus adapter card:
  - a. Ensure that the VMEbus chassis is powered down.
  - b. Decide whether the VMEbus adapter card is the system controller. If so, it must be installed in slot 1. If not, locate an unoccupied 6U slot in the card cage for the adapter.
  - c. Insert the card into the connector of the selected slot.
4. To install the fiber optic cable:
  - a. Ensure that both systems are powered down.
  - b. Remove the rubber boots on the fiber-optic transceivers as well as the ones on the fiber-optic cables. Be sure to replace these boots when cables are not in use.
  - c. Plug one end of the fiber-optic cable into the PCI adapter card’s transceiver.
  - d. Plug the other end of the cable into the VMEbus adapter card’s transceiver.
  - e. Power on both systems.
  - f. Ensure that the READY LEDs on both adapter cards are lit.

## 5.5.2 Installing the Software

1. With RedHawk Linux Version 5.4 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 5.4” 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
```

- To install, execute the following commands:

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

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

- When the installation completes, execute the following commands:

```
cd ..
eject
```

- Remove the disc from the optical media drive and store.

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

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

## 5.6 Installing Additional RedHawk Products

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

## 5.7 Installing RedHawk Updates from Disc

Discs labeled “Software Updates for RedHawk” include updates to RedHawk and its products. If a RedHawk Updates disc is provided, install it now.

- Ensure that the system is in multi-user mode.
- Locate the disc labeled “Software Updates for RedHawk” appropriate to your system’s architecture, insert it into the drive.

- If the disc device does not mount automatically, execute the **mount** command; for example:

```
mount /media/dvd
```

- 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.

## 5.8 Installing RedHawk Updates from the Updates Web Site

1. Additional updates may have been issued. They should be downloaded and installed now. Refer to “Software Updates” on page 83 for details.
2. When complete, or if no updates are available, reboot the system choosing the new RedHawk kernel and verify that the system boots properly.

## 6.0. Upgrade Procedures

**Use the upgrade procedures provided here only if you are upgrading from Version 5.1 or 5.2 to Version 5.4.**

**If generating or regenerating a RedHawk 5.4 system, migrating a system from Version 4.2 or earlier, or converting a system from 32-bit to 64-bit, follow the steps in “Installation Procedures” on page 12.**

### Review the following points before upgrading the software:

- It is recommended that the installer first read through the entire procedure once before actually performing the upgrade.
- It is **strongly recommended** that installation and updates of software be performed in single-user mode to guard against the possibility of corrupting a running application.
- Install all available upgrades. No quality claims are made by Concurrent unless all appropriate version levels of all RedHawk Linux software are installed.
- The media in use is accessed at various times and must be kept in the drive during the installation until you are instructed to install different media.
- Prior to upgrade, remove any floppy media that may have been left in the machine.
- The Red Hat kernel should be running when upgrading to RedHawk Linux Version 5.4.

## 6.1 Pre-Upgrade Procedure

- Before beginning the upgrade procedure, care should be taken to avoid overwriting and thereby destroying files or filesystems that you cannot restore. Identify and back up all important files to a safe location, preferably on a system other than the one on which you are performing the upgrade. For maximum safety, it is recommended that you disconnect all disks that you do not want the upgrade process to accidentally overwrite.
- On systems with multiple Red Hat installations, identify all the appropriate mount points and partitions so that you can correctly answer questions from the installer during the upgrade process.

To identify the partition that contains the root file system to be upgraded, boot the installation to be upgraded into run level 1 and issue the `mount` command. For example:

```
[root@ihawk ~]$ mount | grep 'on / type'
/dev/sda2 on / type ext3 (rw)
```

In this example, the root partition on which this RHEL5.4 was installed is `/dev/sda2`.

Resolve any discrepancies before continuing with the upgrade.

- Avoid adding or removing drives to a system before or after an install or upgrade to prevent having difficulties with grub, the BIOS, and the operating system correctly identifying various drives and partitions.
- If you cannot reboot your system after the upgrade, refer to the section “How to Use Linux Rescue” on page 70.

## 6.2 Upgrading to RHEL5.4

### NOTE

If you are upgrading to RedHawk 5.4 from a RedHawk 5.1 system, please note that you might see problems with 3 packages during the Red Hat upgrade procedure: openmpi, rhnsd and lam-libs. These problems are due to Red Hat, not Concurrent. For the current recommended resolution of these issues, see these links:

<http://kbase.redhat.com/faq/docs/DOC-19442>

<http://rhn.redhat.com/errata/RHBA-2008-0481.html>

The system must first be upgraded to Red Hat Enterprise Linux 5.4. Perform the following steps.

1. Insert the RHEL 5.4 installation disc #1 into the bootable optical media drive.
2. Follow steps 1-8 of the installation instructions in section **5.1.1** on page 13, up to the window which prompts you to either install or upgrade, shown below. Select “Upgrade an existing Installation”.

Make sure that the correct root partition you wish to upgrade is selected in the combo box.

If you are satisfied the upgrade is correctly configured, click the Next button.



3. Select the default option to configure grub: “Update boot loader configuration”. This updates the grub identified as the primary grub used by the system.



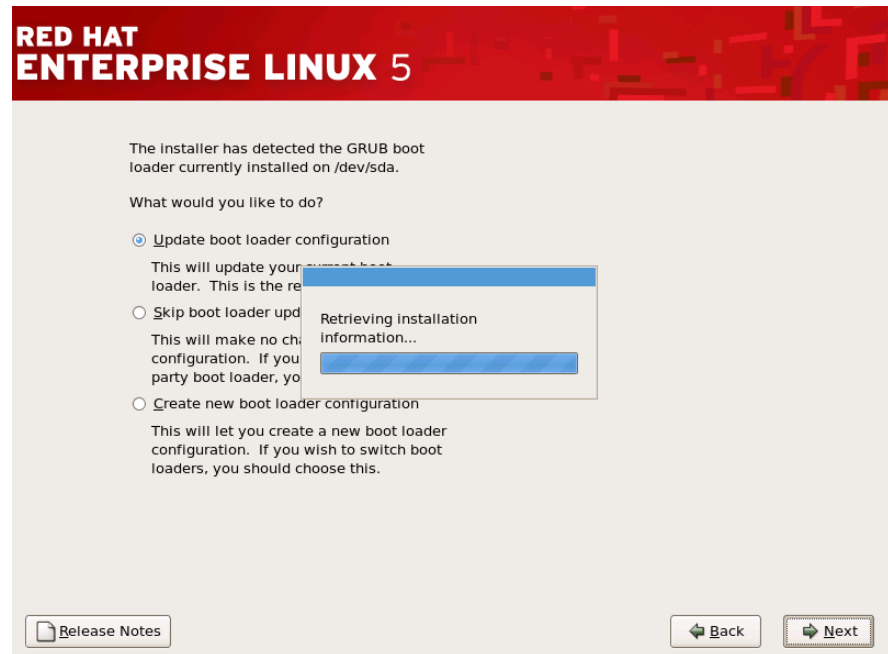
**NOTE**

Unless you are an advanced user knowledgeable about this type of procedure, do not select other options; otherwise, you may not be able to boot the system after the upgrade process completes.



Click Next to continue.

You will see the “progress bar” showing “Retrieving Installation Information” followed by “Checking dependencies in packages selected for installation”.




4. After the dependency check completes, you are prompted to begin the upgrade process. This is your last to chance to abort.

Click Next when you are ready to begin the upgrade.

No action is required from you during this stage.



5. Monitor the progress of the upgrade: “Transferring the install image to hard drive”, “Starting the install process, this may take several minutes”.



After the install process begins, another progress bar displays with messages that provide details as packages are updated, including the remaining time estimated until completion (the actual time is usually faster than indicated) along with the current percentage of packages that have completed installing.



6. After 100% of the packages are installed, remove the installation media and reboot.



## 6.3 Updating Red Hat Software

The disc labeled RHEL 5.4 Updates contains the latest updates released by Red Hat and validated against RedHawk Linux. These updates are important to the proper operation of RedHawk Linux.

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

1. If not already running **Red Hat Enterprise Linux**, boot the system with the Red Hat kernel.
2. Log in as root and take the system down to single-user mode:
  - a. Right click on the desktop and select **Open Terminal**.
  - b. At the system prompt, type **init 1**.
3. Insert the RHEL 5.4 updates disc into the drive.
4. If the disc device does not mount automatically, execute the **mount** command with the mount point that is designated in **/etc/fstab**.

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

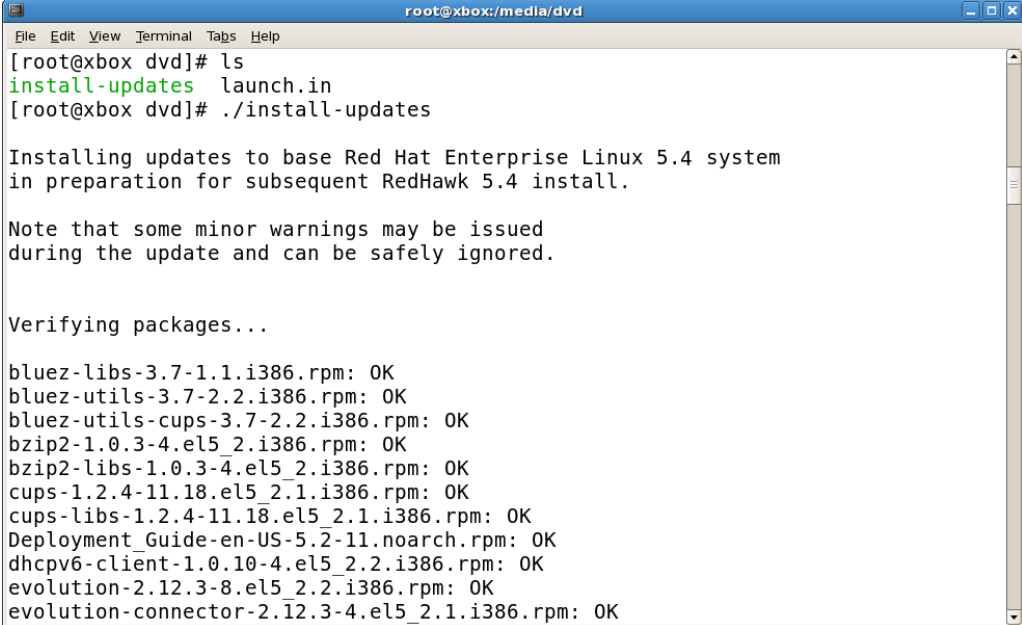
```
mount /media/dvd
```

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

```
mount /dev/dvd /media/dvd
```

5. Execute the following commands to install the update software:

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



```
root@xbox:/media/dvd
File Edit View Terminal Tabs Help
[root@xbox dvd]# ls
install-updates launch.in
[root@xbox dvd]# ./install-updates

Installing updates to base Red Hat Enterprise Linux 5.4 system
in preparation for subsequent RedHawk 5.4 install.

Note that some minor warnings may be issued
during the update and can be safely ignored.

Verifying packages...

bluez-libs-3.7-1.1.i386.rpm: OK
bluez-utils-3.7-2.2.i386.rpm: OK
bluez-utils-cups-3.7-2.2.i386.rpm: OK
bzip2-1.0.3-4.el5_2.i386.rpm: OK
bzip2-libs-1.0.3-4.el5_2.i386.rpm: OK
cups-1.2.4-11.18.el5_2.1.i386.rpm: OK
cups-libs-1.2.4-11.18.el5_2.1.i386.rpm: OK
Deployment_Guide-en-US-5.2-11.noarch.rpm: OK
dhcpv6-client-1.0.10-4.el5_2.2.i386.rpm: OK
evolution-2.12.3-8.el5_2.2.i386.rpm: OK
evolution-connector-2.12.3-4.el5_2.1.i386.rpm: OK
```

- When complete, issue the following commands to remove the disc:

```
cd ..
eject
```

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

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

## 6.4 Upgrading to RedHawk Version 5.4

To upgrade to RedHawk Linux Version 5.4, perform the following steps.

- With the Red Hat Enterprise Linux kernel updated in the previous step running on the system, log in as root and take the system down to single-user mode:
  - Right click on the desktop and select Open Terminal.
  - At the system prompt, type **init 1**.
- Insert the RedHawk 5.4 OS disc into the drive and issue the following commands:

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

```
root@wario:/media/cdrom
File Edit View Terminal Tabs Help
[root@wario cdrom]#
[root@wario cdrom]# ls
documentation  install          launch.in  nvidia      uninstall
extras         install-redhawk launch.un  RedHawk-eula  uninstall-redhawk
[root@wario cdrom]# ./install-redhawk

Installing RedHawk Linux 5.4 (Lima)

This script will install the RedHawk 5.4 release on
the current system (wario.ccur.com).

Do you want to install RedHawk Linux 5.4? [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 /etc/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: 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.
- When the installation completes, issue the following commands:
 

```
cd ..
eject
```
  - Remove the disc from the drive and store.
  - Reboot the system choosing the new RedHawk kernel and verify that the system boots properly.
  - Continue to the next step to upgrade FBS.

## 6.5 Upgrading to Frequency-Based Scheduler (FBS) Version 5.4

If you are using FBS, upgrade to FBS Version 5.4 at this time by performing the following steps. If you are not using FBS, proceed to “Upgrading to PCI-to-VME Bridge Version 5.4” on page 45.

- Insert the updates disc into the drive and issue the following commands:

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

2. When the installation completes, issue the following commands:

```
cd ..
eject
```

3. Remove the optical media disc from the drive and store.

## 6.6 Upgrading to PCI-to-VME Bridge Version 5.4

If you are using the PCI-to-VME bridge supplied by Concurrent to connect your iHawk system with a VME system, upgrade the software to Version 5.4 at this time by performing the following steps. If you are not using this bridge, proceed to “Installing RedHawk Updates from Disc” on page 51.

1. With RedHawk Linux Version 5.4 running, log in as root and take the system down to single-user mode.
2. Locate the disc labeled “RedHawk Linux PCI-to-VME Bridge Library Version 5.4” appropriate to your system’s architecture and insert it into the drive.
3. If the device does not mount automatically, execute the **mount** command; for example:

```
mount /media/dvd /media/dvd
```

4. Execute the following commands to install the PCI-to-VME software:

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

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

5. Upon completion, execute the following commands:

```
cd ..
eject
```

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

## 6.7 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 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.

## 6.8 Installing RedHawk Updates from the Updates Web Site

1. Additional updates may have been issued that are not included on the "Software Updates for RedHawk" disc(s). They should be downloaded and installed now. Refer to "Software Updates" on page 83 for details.
2. When complete, or if no updates are available, reboot the system choosing the updated RedHawk kernel and verify that the system boots properly.

## 6.9 Additional Requirements

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

1. Change directory (**cd**) to **/usr/src/redhawk-release-directory**.
2. Run the **ccur-config(8)** tool as follows:

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



Wait for on-screen messages to complete.

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

1. Change directory (**cd**) to **/usr/src/redhawk-release-directory**.
2. Run the **ccur-config(8)** tool as follows:

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

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

---

---

## 7.0. NVIDIA Graphics Configuration

The Unified NVIDIA Linux Display Driver (version 190.42) 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 **7.1 Supported NVIDIA GPUs** and **7.2 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 5.4.

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 **7.4 Using nvidia-xconfig** to generate the new `xorg.conf` file.

After you have generated the new `xorg.conf`, follow the process described in section **7.5 Using nvidia-settings** to configure your displays, whether single or multiple displays. The types of display configurations are described in **7.3 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.

### 7.1 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-190/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.

## 7.2 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.

## 7.3 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 **7.5 Using nvidia-settings**.

### 7.3.1 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 **7.4 Using nvidia-xconfig**.

### 7.3.2 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.

### 7.3.3 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.

### 7.3.4 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-190/README.txt` for a detailed discussion on Twinview.

### 7.3.5 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.

### 7.3.6 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.

## 7.4 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 **7.6** for reference.

Follow this procedure exactly to generate the file:

1. Read sections **7.0** - **7.3** 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`.

## 7.5 Using nvidia-settings

The `nvidia-settings` utility should be used to configure any of the multidisplay configurations described in section 7.3, and is the preferred method for tuning single display settings.


Before using `nvidia-settings`, complete the `nvidia-xconfig` procedure outlined in section 7.4. 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:

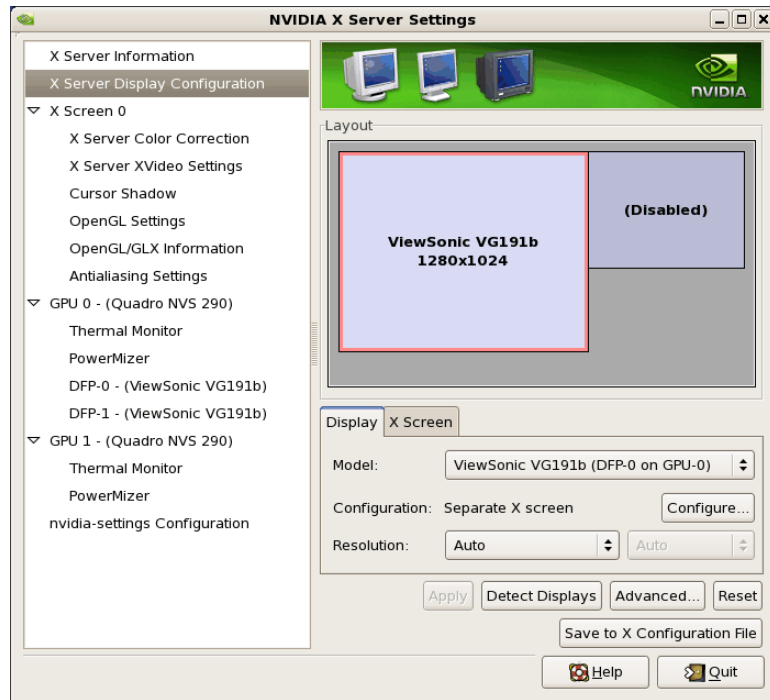
```
nvidia-settings
```

The GUI will start:



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

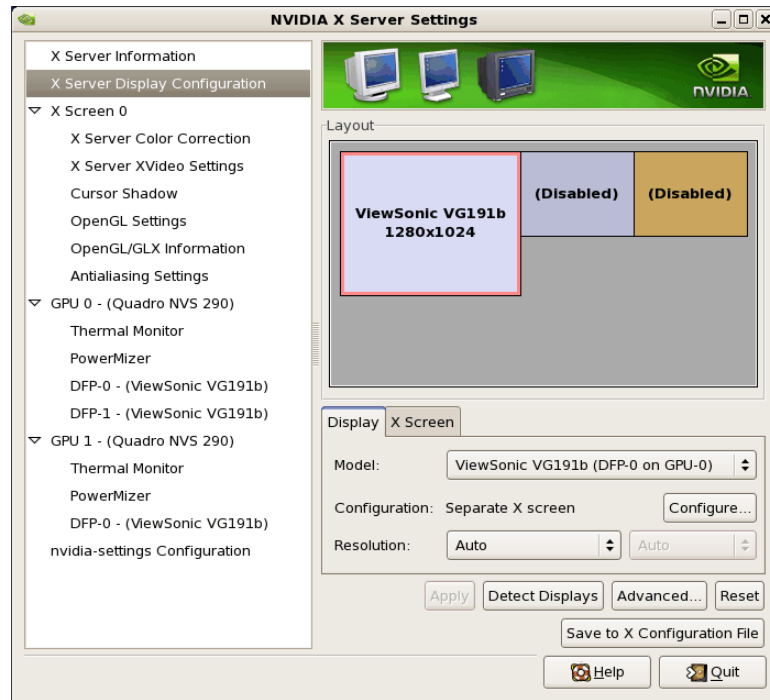
- To create a new “X Server Display Configuration” select that option from the list on the left as shown below:



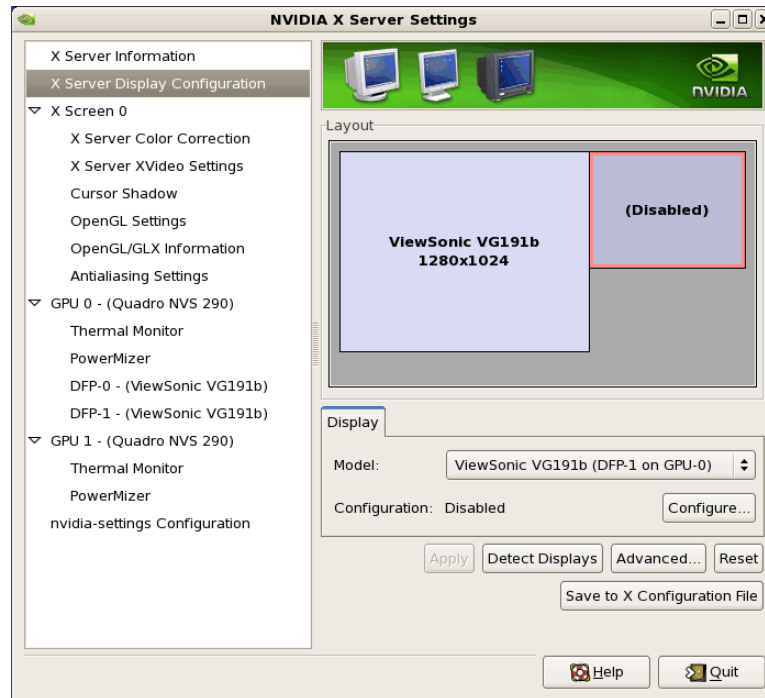
In the Layout window, we see a graphical depiction of the number of displays detected and configured. Notice that the large box is the primary monitor already configured by **nvidia-xconfig**. The smaller box to the right (Disabled) is a display that has not yet been configured.

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:






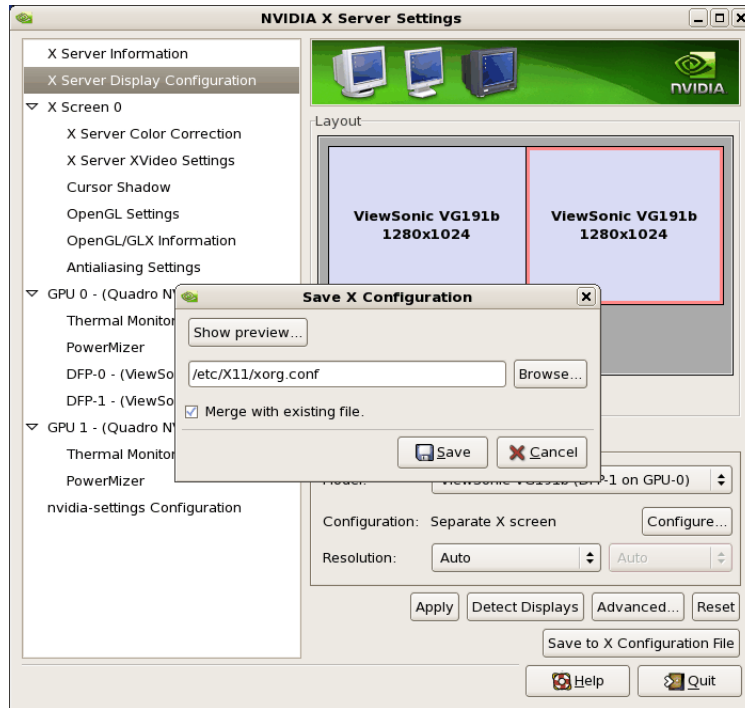
- Click on the Configure button below the Layout window and select the “Separate X screen” option, then click OK:



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

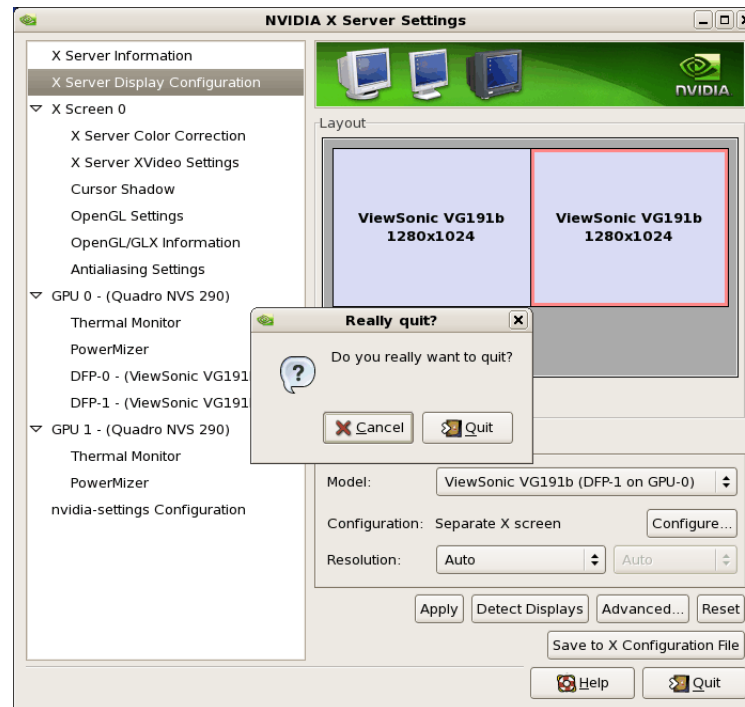


- To save the configuration to **xorg.conf**, click on the **Save to X Configuration File** button near the bottom of the GUI:




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.

## 7.6 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.

The "Module" Section of this file is important. If the Module Section is not present, X will load more modules than required including "dri". The "dri" module has been known to cause race conditions that can cause X to lock up. You should not load the "dri" module or add it to the "Module" Section.

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

```

## 7.7 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 realtime 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.



## 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 <b>/bin/rpm -qal --last   more</b> 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 partition size is based on the size of the system's main memory plus an additional 5 percent. This swap space allocation should be adequate to accommodate system crash dumps.

**Table 7-2 Swap Partition Size Guidelines**

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

### 8.1.3 RCIM Connection Modes

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

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

### 8.1.4 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 every RHEL 5.1, 5.2 and 5.4 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 5.4 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 5.4-20091211 (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 Using 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.
2. At step 4, enter “linux vnc” then press **Enter** to start the VNC installer:

```
boot: linux [options] vnc
```

3. Continue with step 5 in section **5.1.1**.
4. After step 5, the VNC installer prompts 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.
  - a. Select the language. The default is “English”. Select **[OK]** and press **Enter** to continue.

- b. Select the keyboard layout. The default is “US”. Select [OK] and press **Enter** to continue.
- c. Configure the network device used for the VNC connection to the client. A list of all the network devices seen by the installation kernel displays.

The network settings you choose at this point are for establishing the VNC connection. You will be able to change the network configuration from the VNC client at a later point in the installation process.

Select [OK] and press **Enter** to continue.

- d. Configure TCP/IP settings. The IPV4 Manual Configuration is recommended.

Enter the IP address of the system you are installing (VNC server), the IPV4 Netmask, gateway, and name server, as prompted.

In this example, the IP address is “129.134.30.200”. You will see this IP referenced by the VNC server when it starts; for example:

```
129.134.30.200:1
```

“:1” represents the TCP port screen number. Do not enter the TCP port screen number as part of the IP address. The VNC server will decide which TCP port to use.

- e. Select [OK] and press **Enter** to continue. The VNC server program will now start up.

```
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.
```

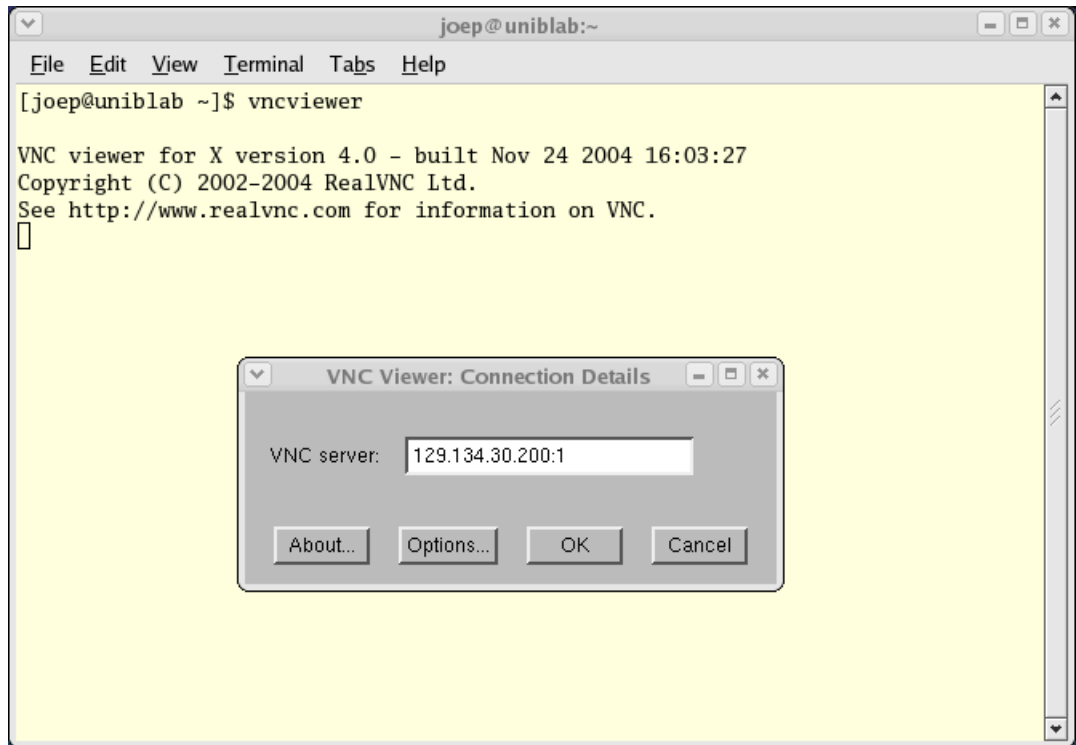
```
The VNC server is now running.
Please connect to 129.134.30.200:1 to begin the install ...
Press <enter> for a shell Starting graphical installation ...
XKB extension not present on :1
```

5. 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.
6. 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**

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



8. Select **OK** and press **Enter** to continue. At this point, you are able to perform a graphical installation remotely.
9. 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 26, perform step 10 below.
10. 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
```

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

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

Tab to [OK] and press Enter.

12. 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.

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

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

14. Reboot the system and continue with section **5.1.2** on page 33.

### 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:           boot: linux driver=nv
ATI:              boot: linux driver=radeon
Generic (incompatible driver): boot: linux driver=vesa
```

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

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

```
boot: linux 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-4.1.2-14** or later.

### 8.2.2 Hyper-thread Activation

Hyper-threading is configured by default in each of the standard RedHawk Linux kernels running on iHawk Series 860 systems. It can be disabled on a per-CPU basis using the **cpu(1)** command. However, hyper-threading is also defined in

the system BIOS. The setting in the BIOS will supersede any configuration of this feature in the kernel. Therefore, if hyper-threading is turned off in the BIOS, it will not be available on your system even if it is configured in the kernel.

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



---

## 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 XFree86 X server with some iHawk platforms, such as the Dell PowerEdge™ 6650 (iHawk Model HQ665).

### NMI Button

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

### kdb and USB Keyboards

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

### PAM Capabilities

- Note that in order to use the **pam\_capability(8)** feature with **ssh**, the **/etc/ssh/sshd\_config** file must have the following option set:

```
UsePrivilegeSeparation no
```

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

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

The krb5-telnet xinetd service should be disabled:

```
chkconfig krb5-telnet off
```

### sadc(8) Issue

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

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

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

## irqbalance

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

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

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

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

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

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

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

## Compatibility Issues for Ada programs

RedHawk 5.4 is based on Red Hat Enterprise Linux 5.4. Binary and source compatibility issues exist for MAXAda programs between RHEL 3.0 and RHEL 5.4.

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

## Inactive Virtual Terminal State After Changing Run Levels

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

## Instability with NVIDIA 8174 Driver and NVS400 Graphics Card

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

## Excessive ksoftirqd Runs Affecting Determinism

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

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

## Pings to Multicast Addresses Disabled by Default

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

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

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

```
# sysctl -w net.ipv4.icmp_echo_ignore_broadcasts=0
```
- To initialize the parameter to the desired value on every reboot, add the following command and `sysctl` parameter to `/etc/sysctl.conf`:
 

```
# Controls broadcast and multicast ICMP echo and timestamp requests
net.ipv4.icmp_echo_ignore_broadcasts = 0
```

## Grub Option Holdovers from Previous Kernel Installations

Be aware that when kernel packages are installed, they add a grub entry with options associated with that kernel (e.g., trace and debug kernels add “crashkernel=64@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.

## vdso and glibc 2.3.2

Changes to the location of the vdso page in the user address space in RedHawk Version 4.2 may result in a segfault if using `glibc` version 2.3.2 (RedHawk 5.1, 5.2 and 5.4 use `glibc` 2.5).

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

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

## USB Ports Non-functional on Supermicro Boards

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

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

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

### Configure Quad Port Ethernet

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

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

### Start haldaemon to burn files onto optical media discs

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

```
$ service haldaemon start
```

After the copying process is complete, stop the service:

```
$ service haldaemon stop
```

### Issues with Tigon3 Network Adapter

Significant reliability and operability problems have been detected with the Tigon3 (BCM95700A6) network adapter. As yet, there are no solutions for these problems under RedHawk 5.1, 5.2 or 5.4.

### Warning Issued When Using sudo

Occasionally, when using `sudo`, the following warning may be issued:

```
audit_log_user_command(): Connection refused
```

The warning can be safely ignored and there is no functionality loss.

### GNOME Related Problems

A number of problems with GNOME have been discovered.

- A conflict exists between GNOME and X that can result in displays blacking out and not returning to normal. The solution is to remove the GNOME screensaver and Power Management software from the system. To check if these rpms are present on your system, issue the following:

```
rpm -q gnome-screensaver
rpm -q gnome-power-manager
```

To remove them, issue the following:

```
rpm -e gnome-screensaver
rpm -e gnome-power-manager
```

Removing these rpms does not affect applications or the stability of the system. This issue has been reported to GNOME.

- GNOME settings sometimes fail to initialize when booting to init 5. This results in the GNOME Desktop failing to function properly and can be recognized when the usual

Desktop wallpaper and or GNOME panels do not appear or when one of the following errors displays:

An error occurred while loading or saving configuration information for gnome-settings-daemon.

Some of your configuration settings may not work properly.

or

GConf error: Failed to contact configuration server; some possible causes are that you need to enable TCP/IP networking for ORBit, or you have stale NFS locks due to a system crash.

See <http://www.gnome.org/projects/gconf/> for information. (Details - 1: IOR file '/tmp/gconfd-root/lock/ior' not opened successfully, no gconfd located: No such file or directory 2:IOR file '/tmp/gconfd-root/lock/ior' not opened successfully, no gconfd located: No such file or directory)

To workaround this problem, perform the following:

- a. Remove the `/tmp/gconfd-root` directory:

```
rm -rf /tmp/gconfd-root
```

- b. Drop to runlevel 3 then back to runlevel 5:

```
init 3; init 5
```

- The GNOME System Tray keeps “utility” programs such as the calendar and clock running and out of the way. You may see the following error dialog after X starts, and it may occur frequently on some systems.



Click on Don't Delete, drop to runlevel 3 then back to runlevel 5.

If you select Delete, you will have to add the applet back to the panel. The applet continues running after issuing the error, so there's no need to delete it.

- If you notice that no panels are running on the desktop, usually accompanied by the following error, a reboot is required. This sometimes occurs after changing runlevels. Click OK then reboot the system.



### Crash Dump Support

On RedHawk systems, the grub kernel option `percpu_alloc=4k` should not be used in combination with the `crash(1)/kdump(8)` support.

Also, as was the case in RedHawk 5.2, the `crash(1)` `mach-c` command only displays the first CPU information properly.

### Boot Time lockd Error Message

During the RedHawk boot process, you may see the following error appear on the console:

```
...
...
Starting kernel logger: [ OK ]
Starting portmap: [ OK ]
FATAL: Module lockd not found.
Starting NFS statd: [ OK ]
Starting system message bus: [ OK ]
Mounting other filesystems: [ OK ]
...
...
```

This error message does *not* indicate a true problem. It is simply issued by `modprobe` because it cannot find the `lockd` kernel module. There is no `lockd` kernel module on RedHawk; the `lockd` support is built statically into the kernel, so no module is required.

### System Hang Due To Incorrect Usage of rcim\_clocksync(1) Utility

If the RCIM is configured as the clocksource device for a system *and* that system has a standalone RCIM II or RCIM III, then the `rcim_clocksync(1)` utility should *not* be used to zero or synchronize the tick clock.

In addition, if the RCIM is configured as the clocksource device for a system *and* that system has a cable-connected master RCIM II, then the `rcim_clocksync(1)` utility should *not* be used to stop the tick or POSIX clocks, or to disable the cable clock signal.

In these situations, it is recommended to first change the system clocksource to some other device, such as the TSC clock or HPET timer, before using the `rcim_clocksync(1)` utility.

For more information about configuring other devices as the system clocksource, see the *System Clocks and Timers* chapter in the *RedHawk Linux User's Guide*.

---

---

## 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](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>.

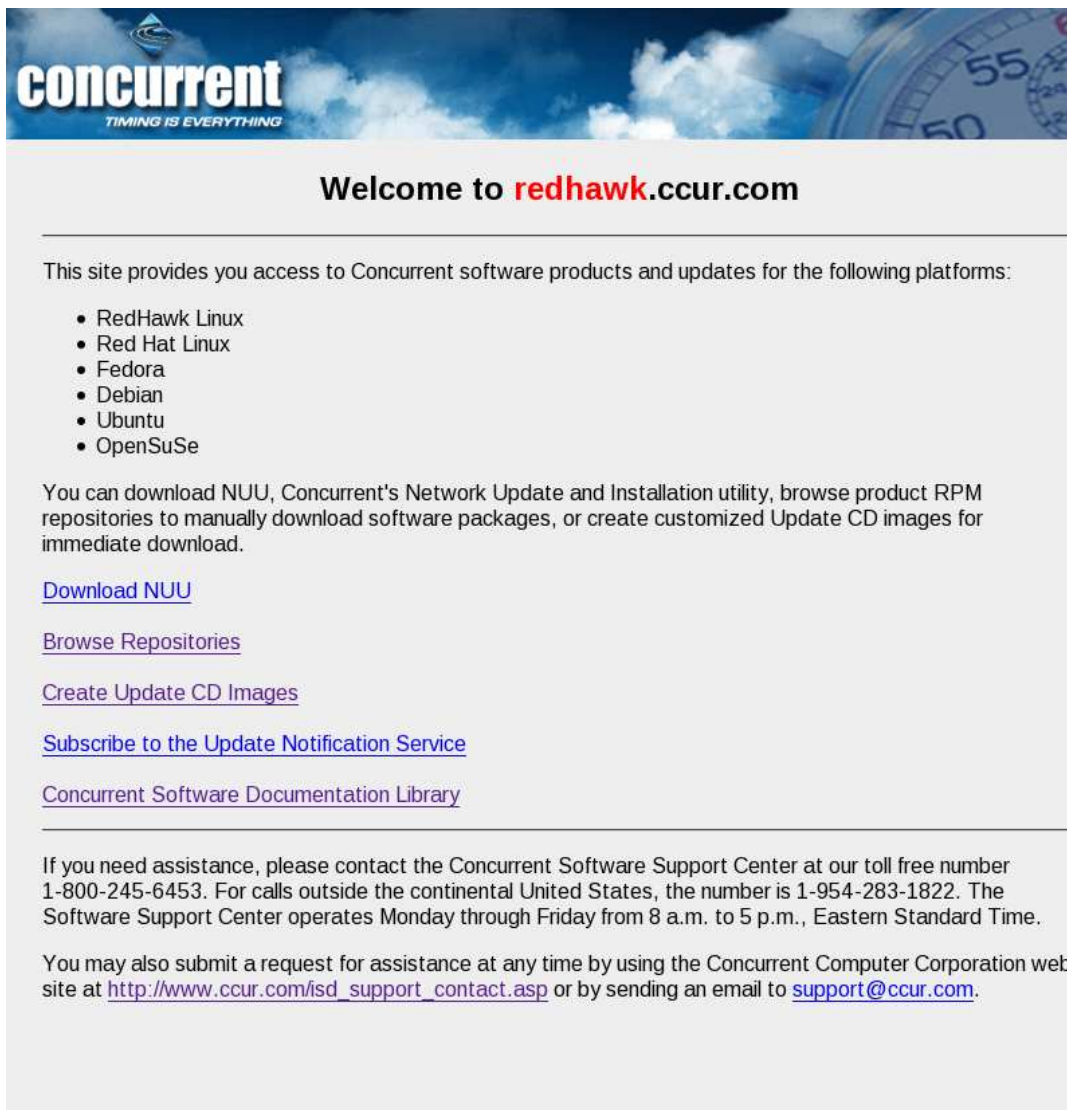
The first time you invoke NUU you will need to specify your redhawk.ccur.com Login ID and Password that were provided in the shipping documents accompanying your system. If you require assistance, refer to "Direct Software Support" on page 83.

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](#)

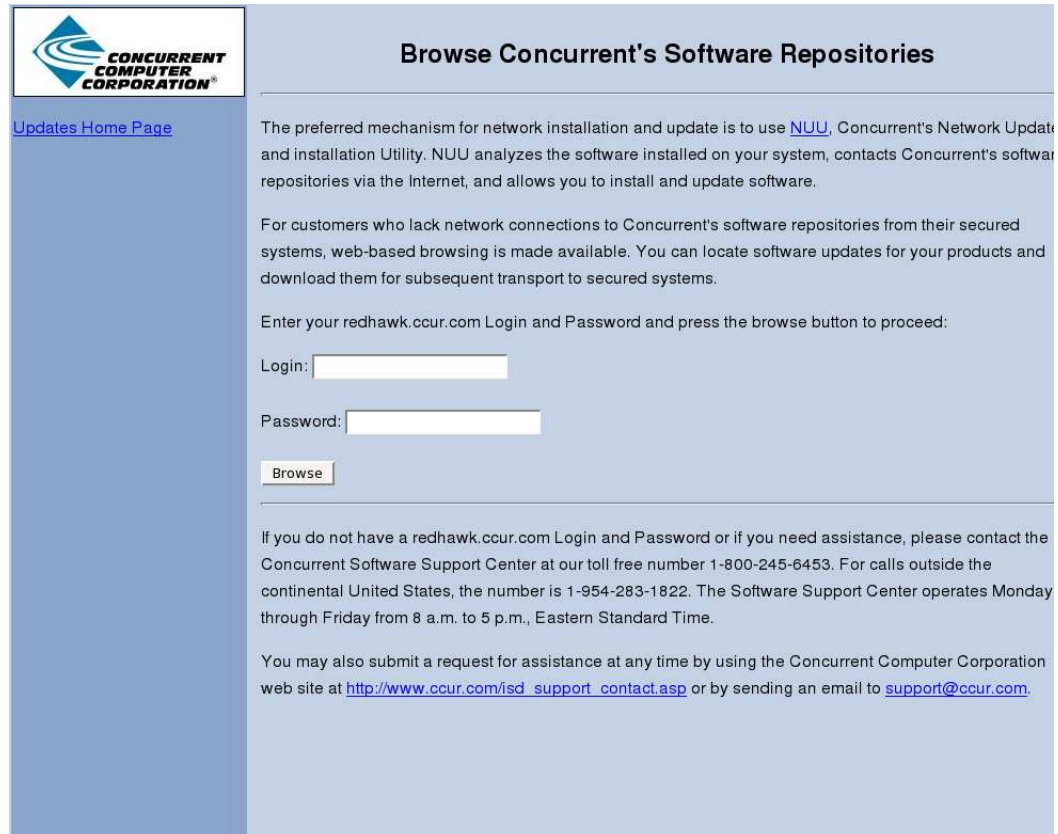
---

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](http://www.ccur.com/isd_support_contact.asp) or by sending an email to [support@ccur.com](mailto:support@ccur.com).



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



**Browse Concurrent's Software Repositories**

[Updates Home Page](#)

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:




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](http://www.ccur.com/isd_support_contact.asp) or by sending an email to [support@ccur.com](mailto: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/5.4

<a href="#">Name</a>	<a href="#">Last modified</a>	<a href="#">Size</a>	<a href="#">Description</a>
 <a href="#">Parent Directory/</a>		-	
 <a href="#">i386/</a>	10-Dec-2009 20:58	-	
 <a href="#">x86_64/</a>	10-Dec-2009 21:00	-	

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>.