RED HAWK Linux®
Version 7.0
Release Notes

May 2015

0898003-7.0a

concurrent

Disclaimer

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

License

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.

© 2015 Concurrent Computer Corporation – All Rights Reserved

Concurrent Computer Corporation
2881 Gateway Drive
Pompano Beach, FL 33069

Note: Information subject to change without notice.

Contents

1.0 Introduction 1
1.1 Product Description 1
1.2 Product Media 2
1.3 Version Numbers 2
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 7.0 8
4.1.1 Based on CentOS 7.0 8
4.1.2 Based on kernel.org 3.16 8
4.1.3 RedHawk 7.0 now 64-bit Only 8
4.1.4 Full 32-bit Real-Time API Compatibility 8
4.1.5 Updated NVIDIA Driver 8
4.1.6 Updated CUDA SDK 8
4.1.7 NUMA Balancing 9
4.1.8 Updated PTP Daemon 9
4.1.9 Optional OpenOnload® Network Stack 9
4.1.10 Additional Grub Utility 9
4.1.11 Compressed RAM Disk (zram) Memory Shielding . . 9
4.1.12 Alternate Deterministic Math Functions Available . . 9
4.2 Modifications in Version 7.0 10
4.2.1 Now Using Grub 2 10
4.2.2 UEFI Fully Supported 10
4.2.3 /proc/sysvipc/shm Changes 11
4.2.4 New /usr/ccur Directories 11
4.2.5 RedHawk crash(1) Changes 11
4.2.6 kexec-tools Changes 11
4.2.7 RedHawk Linux strace(1) Changes 11
4.2.8 Deprecated clock and timer services in libccur_rt . 11
4.2.9 Moving to 7.0 from 6.x 12
4.2.10 Gnome 3 Graphical Desktop 12
4.2.11 New Network Interface Device Names 12
4.2.12 Docker not yet supported in 7.0 12
5.0 Installation Procedures 13
5.1 Installing CentOS Software 13
5.1.1 Installing CentOS 14
5.1.2 Installing CentOS Updates 28
5.2 Installing RedHawk Linux 32
5.3 Installing the RCIM 34
5.3.1 Quick Hardware Installation Checklist 35
5.3.2 Verifying RCIM Operation 35
5.3.3 Installing NTP Updates for GPS Module 36
5.4 Installing Frequency-Based Scheduler (FBS) Software 36

5.5 Installing the PCI-to-VME Bridge 37
5.5.1 Installing the Hardware 37
5.5.2 Installing the Software 38
5.6 Installing the PCIe-to-VME Bridge 38
5.6.1 Installing the Hardware 39
5.6.2 Installing the Software 39
5.6.3 Configuring the PCIe-to-VME Windows 40
5.7 Installing Additional RedHawk Products 41
5.8 Installing RedHawk Updates from Disc 41
5.9 Installing RedHawk Updates from the Updates Web Site 42
6.0 NVIDIA Graphics Configuration 43
6.1 Supported NVIDIA GPUs 43
6.2 Multiple NVIDIA Video Card Considerations 44
6.3 Display Configurations 44
6.3.1 Single 44
6.3.2 Xinerama 44
6.3.3 MultiDesktop 45
6.3.4 Twinview 45
6.3.5 Twinview-Xinerama 46
6.3.6 Twinview-MultiDesktop 46
6.4 Using nvidia-xconfig 46
6.5 Using nvidia-settings 47
6.6 Modified xorg.conf file generated by nvidia-xconfig 54
6.7 Troubleshooting Problems 55
7.0 Additional Information 58
7.1 Installation/Configuration Issues 58
7.1.1 File System Backup Recommendations 58
7.1.2 File System Swap Size Guidelines 59
7.1.3 RCIM Connection Modes 59
7.1.4 How to Use Linux Rescue 59
7.1.5 How to Install CentOS Remotely Over VNC 61
7.1.6 If the Installation GUI Fails to Start 64
7.2 Considerations 64
7.2.1 Compiler Requirements 64
7.2.2 Hyper-thread Activation 64
8.0 Known Issues 65
9.0 Software Updates and Support 68
9.1 Direct Software Support 68
9.2 Software Updates 68
9.2.1 Updating via NUU 68
9.2.2 Installing Manually Downloaded RPMs 69
9.2.3 Customized Update Discs 71
9.3 Documentation Updates 71

1.0. Introduction

This document provides release information and installation instructions for Concurrent Computer Corporation's RedHawk™ Linux® operating system Version 7.0.

1.1 Product Description

RedHawk Linux is a real-time version of the open source Linux operating system. Modifications are made to standard Linux version 3.16 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 CentOS Community ENTerprise Operating System® version 7.0. However, note that RedHawk can also be installed on *any* CentOS-compatible version 7.0 distribution.

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

The RedHawk kernel integrates both open source patches and Concurrent developed features derived from the real-time UNIX® implementations that Concurrent has supported in its over 44 years experience developing real-time operating systems.

RedHawk is included with each Concurrent iHawk™ system. iHawks are symmetric multi-processor (SMP) systems available in a variety of architectures and configurations, though only 64-bit capable iHawks can run RedHawk 7.0 as it is only available in a 64-bit version.

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

RedHawk Linux matches the level of POSIX® conformance of other Linux distributions based on the 3.16 series of kernels. Linux on the Intel® x86 architecture has defined a defacto binary standard of its own which allows shrink-wrapped applications that are designed to run on the Linux/Intel x86 platform to run on Concurrent's iHawk platform.

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

- NightView™ source-level debugger – allows multi-language, multi-processor, multi-program and multi-thread monitoring and debugging from a single graphical interface.
- NightTrace™ run-time analyzer – analyzes the dynamic behavior of a running application.

- NightSim™ periodic scheduler – allows the user to easily schedule applications that require periodic execution.
- NightProbe™ data monitor – used to sample, record or modify program data in multiple running programs.
- NightTune™ performance tuner – used for analyzing system and application performance.

1.2 Product Media

The tables below list the optical media discs that are delivered with RedHawk Linux Version 7.0 for each Concurrent system.

WARNING

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

Product Media

Processor Type:	**Intel and AMD x86_64 64-bit**
Model:	**iHawk Intel-based systems [HQxxx]** **iHawk AMD-based systems [HRxxx]** **ImaGen Intel-based systems [HQ0G1]** **ImaGen AMD-based systems [HR0G1]**
Operating System Discs:	
CentOS 7.0 (64-bit) (1 DVD)	
CentOS 7.0 Source (64-bit) (2 DVDs)	
CentOS 7.0 Updates (64-bit) (1 DVD) [LXUPDATE WA-AMD64 Version 7.0]	
RedHawk Linux OS (64-bit) (1 DVD) [WA-EM64T-AMD64 Version 7.0]	
Software Updates for RedHawk (1 DVD or 1 CD, if present)	
Optional Discs:	
RedHawk Linux Frequency-Based Scheduler for x86_64 Systems [WU1210-JA Version 7.0]	
NightStar RT for RedHawk [WU1220-LDV, WU1220-LDVF Version *x.x* (*x.x*=latest version)]	
RedHawk Linux PCI-to-VME Bridge Software Library for x86_64 Systems [WU-HS630-LIB (WCS-SBS-620)] Version 7.0	
RedHawk Linux PCIe-to-VME Bridge Software Library for x86_64 Systems [WU-HS630-C2243] Version 7.0	
RedHawk Architect (64-bit) [W-RHLE-E64] Version 7.0	

1.3 Version Numbers

All RedHawk Linux version numbers take the following form:

major.*minor*[.*update*]

where:

major is the product's major version number

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

For example, 7.0 is the initial release of major version 7, and 7.3 is a subsequent release to version 7.0. Both 7.0 and 7.3 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 7.0.1 kernel update can only be installed on a system currently running RedHawk Linux version 7.0.

1.4 System Updates

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

NOTE

Concurrent does not recommend downloading CentOS updates.

The RedHawk Linux kernel replaces the standard CentOS kernel and is likely to work with any version of the CentOS 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 68.

1.5 Syntax Notation

The following notation is used throughout this document:

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

2.0. Documentation

2.1 RedHawk Linux Documentation

Table 2-1 lists the documentation provided with RedHawk Version 7.0 as PDF files. After installing RedHawk Linux documentation, the files are available for viewing:

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

Table 2-1 RedHawk Linux Documentation

RedHawk Linux Operating System Documentation	**Document Number**
RedHawk Linux Version 7.0 Release Notes	0898003-7.0
RedHawk Linux User’s Guide	0898004-700
Real-Time Clock & Interrupt Module (RCIM) User’s Guide	0898007-620
RedHawk Linux FAQ	N/A
Optional RedHawk Product Guides	
RedHawk Linux Frequency-Based Scheduler (FBS) User’s Guide	0898005-350
Optional RedHawk Product Documentation Sets	
RedHawk Architect Release Notes	0898600-7.0
RedHawk Architect User’s Guide	0898601-7.0
SBS PCI-to-VME Bridge Series 600 Hardware Manual	85851150
SBS PCI-to-VME Bridge Series 800 Hardware Manual	85913163
SBS PCI-to-VME Bridge Software Manual	85221990

The Optional RedHawk Product Documentation Sets include Concurrent’s user guides as PDF files along with additional vendor documentation PDFs that facilitate using these RedHawk products.

The latest documentation for all Concurrent software products, including updated Release Notes and User Guides, can be accessed from Concurrent’s documentation web site at **http://redhawk.ccur.com/docs**. An updated RedHawk FAQ can be accessed at **http://redhawk.ccur.com/docs/root/1redhawk/FAQ/RedHawk-FAQ.pdf**.

Printed copies can be ordered by contacting the Concurrent Software Support Center. Refer to “Direct Software Support” on page 68 for details.

2.2 Vendor Documentation

CentOS operating system documentation is included as PDF files on the CentOS 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:

- CentOS 7.0 with updates
- RedHawk Linux operating system (with optional updates)
- NightStar RT for RedHawk tools (optional)
- RedHawk Linux Frequency-Based Scheduler (FBS) (optional)
- RedHawk Linux PCI-to-VME bridge software (optional)
- RedHawk Linux PCIe-to-VME bridge software (optional)
- RedHawk Architect (optional)

These packages operate only in a 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

In RedHawk 7.0, real-time extensions and features are available to 32-bit compatibility-mode applications running under a 64-bit operating system (i.e. x86_64).

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 7.0 (and any subsequent kernel updates if applicable).

4.1 Enhancements in Version 7.0

4.1.1 Based on CentOS 7.0

RedHawk 7.0 uses CentOS Community ENTerprise Operating System 7.0 as its base. CentOS 7.0 is the world-wide standard enterprise-class Linux distribution that provides a rock-solid foundation.

4.1.2 Based on kernel.org 3.16

RedHawk 7.0 uses the `kernel.org` Linux 3.16 kernel as the basis for its core kernel. This version includes many exciting enhancements. For more information examine the kernel change summaries at the **http://kernelnewbies.org** website.

4.1.3 RedHawk 7.0 now 64-bit Only

RedHawk 7.0 is the first release of RedHawk Linux to only release 64-bit kernels and 64-bit user-level software. Thus, RedHawk7.0 will only install and run on systems that are capable of running in a true 64-bit mode (x86_64).

4.1.4 Full 32-bit Real-Time API Compatibility

RedHawk 7.0 now provides full 32-bit real-time API compatibility. Thus, existing 32-bit real-time applications will continue to run unmodified under RedHawk 7.0 using the 32-bit compatibiltiy mode. New 32-bit real-time applications can also still be created, for example, by using the `gcc -m32` option. See the **gcc**(1) man page for more information.

4.1.5 Updated NVIDIA Driver

RedHawk 7.0 includes the NVIDIA Linux Display Driver version 346.35. Legacy driver support for older NVIDIA GPUs is available for download from the NVIDIA website. Refer to "NVIDIA Graphics Configuration" on page 43 for details.

4.1.6 Updated CUDA SDK

RedHawk 7.0 includes pre-installed NVIDIA CUDA 6.5 libraries and header files. CUDA 6.5 has significant improvements over previous versions, including Unified Memory and Multi-GPU Scaling. See the many resources at **http://www.nvidia.com** fore more information.

Example CUDA code is also provided with RedHawk. See the files in the **/usr/share/doc/nvidia/cuda/samples** directory for more information.

4.1.7 NUMA Balancing

RedHawk 7.0 includes a new kernel.org NUMA Balancing feature that performs NUMA balancing on non-shielded CPUs. Over time, both pages and tasks will be optimized for *locality* such that there will be fewer and fewer foreign-node accesses occurring on the system.This feature may be enabled at boot time with a grub option, or enabled dynamically via a **sysctl**(8) parameter.

For more details, refer to the section on NUMA Balancing in Chapter 10 “Non Uniform Memory Access (NUMA)” in the RedHawk Linux User’s Guide.

4.1.8 Updated PTP Daemon

RedHawk 7.0 now supports both hardware PTP and software PTP using the software developed by The Linux PTP Project. This new PTP implementation moves timekeeping into the kernel for both software PTP and hardware PTP. See the **http://linuxptp.sourceforge.net** for more information.

4.1.9 Optional OpenOnload® Network Stack

OpenOnload is a high performance network stack that moves the packets of each data stream up and down the network stack at the priority of the process creating or consuming the packets. This results in dramatic improvements to both latency and jitter for those data streams owned by high priority tasks.

RedHawk 7.0 comes with OpenOnload version 201405-u2, however OpenOnload is initially disabled. To enable OpenOnload, execute the following commands as the root user:

```
systemctl enable openonload
systemctl start openonload
```

OpenOnload works only with network cards that have the vNIC interface. These cards are available from SolarFlare Communications, Inc. and also from SMC Corporation. See **http://www.openonload.org** for more information about OpenOnload software and supported cards.

4.1.10 Additional Grub Utility

RedHawk 7.0 includes a new grub utility, **ccur-grub2**, which may be used to quickly view a list of the available grub kernel entries. It may also be used to change the default grub kernel to boot. See the **ccur-grub2**(1) man page for more information.

4.1.11 Compressed RAM Disk (zram) Memory Shielding

RedHawk 7.0.1 provides a mechanism for ensuring that all memory associated with a compressed RAM disk is contained within a specified set of NUMA nodes. For more information, refer to the following file on an installed RedHawk 7.0.1 system:

```
/usr/src/linux-3.16.*RedHawk7.0*/Documentation/blockdev/zram-numa.txt
```

4.1.12 Alternate Deterministic Math Functions Available

RedHawk 7.0.1 includes the ccur-math RPM, which installs alternative deterministic implementations of the sin, cos, asin, acos and apow math functions. See the **sin**(3ccur) man page for more information.

4.2 Modifications in Version 7.0

4.2.1 Now Using Grub 2

RedHawk 7.0 now uses the Grub 2 bootloader. While Grub 2 offers some new advantages, it has also significantly complicated grub configuration and extra care must be taken to ensure that the grub configuration files remain correct.

In particular, ignore the comments at the top of **grub2.cfg** which tell you not to edit the file directly; editing the file directly or preferably using the **ccur-grub2** or **grubby** commands to make modifications is the recommended procedure for making Grub 2 configuration changes on RedHawk systems.

NOTE

Do not use **grub2-mkconfig** on RedHawk systems, and instead you should use **ccur-grub2** or **grubby** or hand edit the **grub2.cfg** file directly. See the **ccur-grub2**(1) and **grubby**(8) man pages for more information.

NOTE

Be careful when using dollar sign characters (e.g. "$") in the `grub2.cfg` file; dollar sign characters must be escaped with a backslash (e.g. "\$") or they will be interpred as Grub 2 variables. For example, a memory reservation kernel grub option could be specified as follows: `memmap=32M\$1G`.

4.2.2 UEFI Fully Supported

CentOS 7.0 and RedHawk 7.0 have full support for hardware that uses the Unified Extensible Firmware Interface (UEFI). More information about UEFI can be found via this web link:

http://en.wikipedia.org/wiki/Unified_Extensible_Firmware_Interface

Unfortunately, there is a known issue with UEFI in the CentOS 7.0 boot loader: it ignores the usual **grub2** logic for setting the default kernel. To work around this issue on a UEFI system, simply edit the **/boot/efi/EFI/centos/grub.cfg** file and set the default kernel index manually. For example, add the following line before the first **menuentry** line to boot the third listed **menuentry** kernel by default:

```
set default=2
```

Remember that the kernel index numbers start with 0 to select the first **menuentry** kernel listed.

4.2.3 /proc/sysvipc/shm Changes

The `physaddr` field in the output of **`/proc/sysvipc/shm`** has been removed. A new **`/proc/sysvipc/ccur-shm`** file has been added that contains the same output of **`/proc/sysvipc/shm`** plus the additional `physaddr` field.

4.2.4 New /usr/ccur Directories

RedHawk now installs some executable binaries into a new **`/usr/ccur/bin`** directory. These executables are generally also provided by the base distribution. For example, **`strace`** is provided by the base distribution in **`/usr/bin`**, however the version in the **`/usr/ccur/bin`** directory has been enhanced to know about additional RedHawk-specific system calls.

In addition, new **`/usr/ccur/man`** and **`/usr/ccur/share/man`** directory hierarchies have been created to hold man pages that would otherwise conflict with man pages provided by the base distribution.

4.2.5 RedHawk crash(1) Changes

The RedHawk version of the **`crash`**(1) utility is now located at **`/usr/ccur/bin/crash`**, and its associated man page is located at **`/usr/ccur/man/man8/crash.8`**.

The man page for the RedHawk version can be viewed by issuing the following shell command:

```
# man -M /usr/ccur/man crash
```

4.2.6 kexec-tools Changes

The standard CentOS **`kexec-tools`** utilities are now being used for creating vmcore kdump crash images. See "Chapter 12 Kernel Debugging" in the RedHawk Linux User's Guide for information about using and configuring this support.

4.2.7 RedHawk Linux strace(1) Changes

The RedHawk **`strace`**, **`strace-graph`** and **`strace-log-merge`** utilities are now in the **`/usr/ccur/bin`** directory, and the man page for the RedHawk version of **`strace`** is located in the **`/usr/ccur/share/man/man1`** directory.

4.2.8 Deprecated clock and timer services in libccur_rt

Duplicate `clock_*` and `timer_*` services in the **`libccur_rt`** library have been deprecated and will be removed in a future version of RedHawk. Existing binaries will run unmodified, but newly built programs will require specifying **`-lrt`** to reference these services from the system's **`librt`** library.

Note that the deprecated version of **`clock_nanosleep`**(2) could in some cases return error codes that were not Posix compliant. Users are encouraged to switch to the more recently updated version in the **`librt`** library.

4.2.9 Moving to 7.0 from 6.x

As with all RedHawk major releases, existing RedHawk 6.x systems cannot be upgraded to RedHawk 7.0 and instead a complete system regen is required. See "Installation Procedures" on page 13 for details on installing RedHawk 7.0.

4.2.10 Gnome 3 Graphical Desktop

The graphical desktop interface has been significantly overhauled in RedHawk 7.0 from previous versions of RedHawk. However, users preferring the older Gnome 2 desktop can install MATE by issuing the following shell commands as root (*requires Internet access*):

```
# yum install epel-release
# yum groups install "MATE Desktop"
```

After these commands complete and the system is rebooted, a Gnome 2 compatible `mate` option can be selected from the graphical login screen's gear icon.

4.2.11 New Network Interface Device Names

If you do anything at all with networking, you will quickly notice that the network interface device names in RedHawk 7.0 have dramatically changed. For example, two network interfaces that used to be named **`eth0`** and **`eth1`** may now be named **`enp7s0`** and **`enp8s0`**. For those interested, the rationale for this change can be found at this link:

http://www.freedesktop.org/wiki/Software/systemd/PredictableNetworkInterfaceNames/

While this new system does have some advantages, many people dislike this change, and thankfully you can return to the previous network interface naming scheme simply by adding these two grub options to the grub kernel command line:

```
net.ifnames=0 biosdevname=0
```

4.2.12 Docker not yet supported in 7.0

Docker operating system-level virtualization containers are not fully compatible with the initial release of RedHawk 7.0. Full compatibility support for Docker will be available in a future RedHawk release.

5.0. Installation Procedures

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

NOTE: Upgrades from RedHawk 6.x to RedHawk 7.0 are *not* supported.

NOTE: Resources to assist users supplying their own Red Hat software are available in the Concurrent Software Documentation Library:
http://redhawk.ccur.com/docs/document.cgi?document=1redhawk/1OS/7.0/RedHat

Review the following points before installing the software:

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

Before you start:

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

5.1 Installing CentOS Software

The version of CentOS Community ENTerprise Operating System (CentOS) that is appropriate to your system's architecture must be installed on the system first by *specifically completing the steps below*.

Installation of CentOS is accomplished in two parts:

- Installing CentOS
- Installing CentOS Updates

If you already have a CentOS system installed at the 7.0 version level you can skip directly to "Installing CentOS Updates" on page 28.

5.1.1 Installing CentOS

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

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

4. Initially the option to test the media is highlighted; press the Up Arrow to hightlight Install CentOS 7 and press Enter.

NOTES

If the installation GUI does not launch:

If you wish to use GUI mode but it does not launch, refer to section **7.1.6** on page 64 for some boot options that may cause the GUI to launch. Some NVIDIA cards require specifying the display resolution before the install GUI will start. If successful, continue to follow the installation directions below.

In addition to GUI mode, CentOS 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 **7.1.5** on page 61 for full instructions.

If your system contains disk drives or devices that are not supported by the CentOS installation program, you will need to obtain the necessary driver(s) in order to install CentOS Community ENTerprise Operating System. A diskette containing the driver may be included with your system or it may be available on the CentOS installation media. If not available, contact the driver manufacturer for an updated driver or call Concurrent Support (see "Direct Software Support" on page 68).

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. Once the installation has started you will be presented with the following welcome screen.

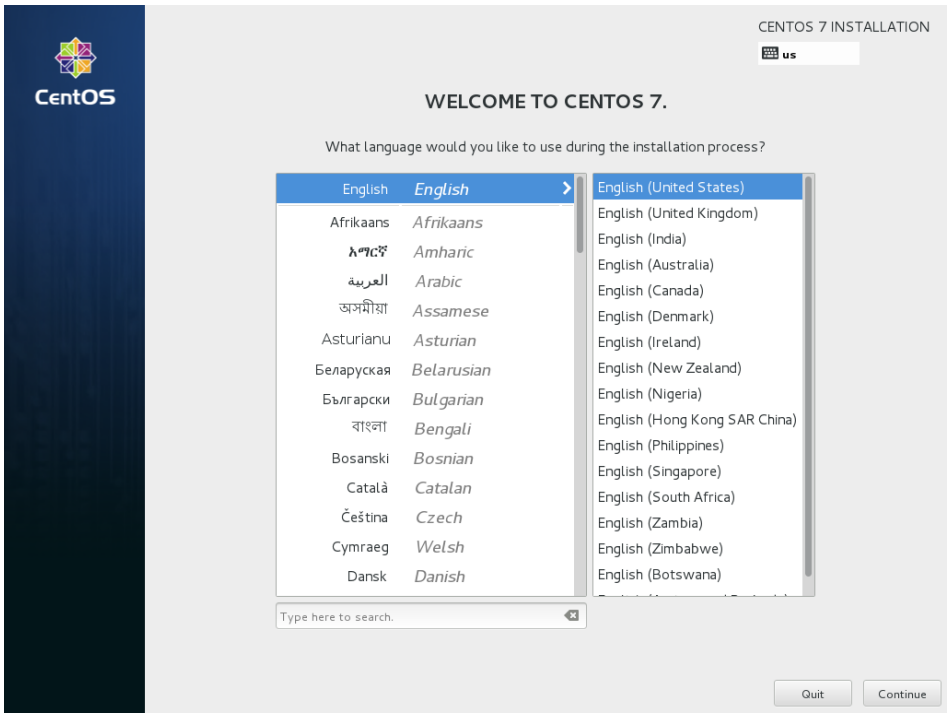

6. Select your desired language and and press Continue proceed.

7. You will now be presented with the Installation Summary page, which is the main launching point for all installation activities.

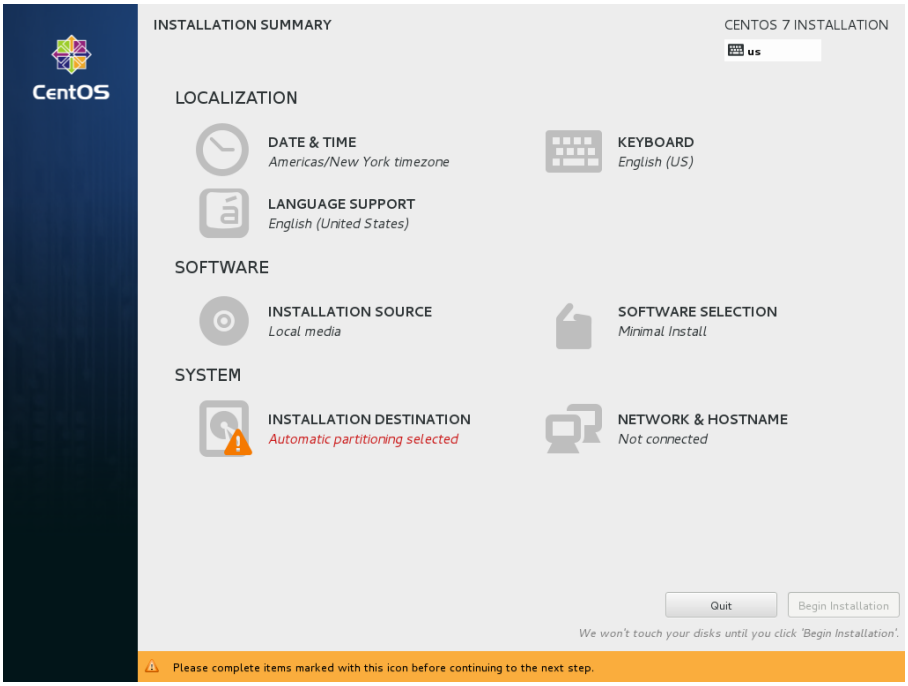

Click on the settings under Localization if you wish to make any site-specific changes to the date, time, language and keyboard settings.

8. Click on Software Selection to select the package installation profile for the installation. The following screen will be displayed.

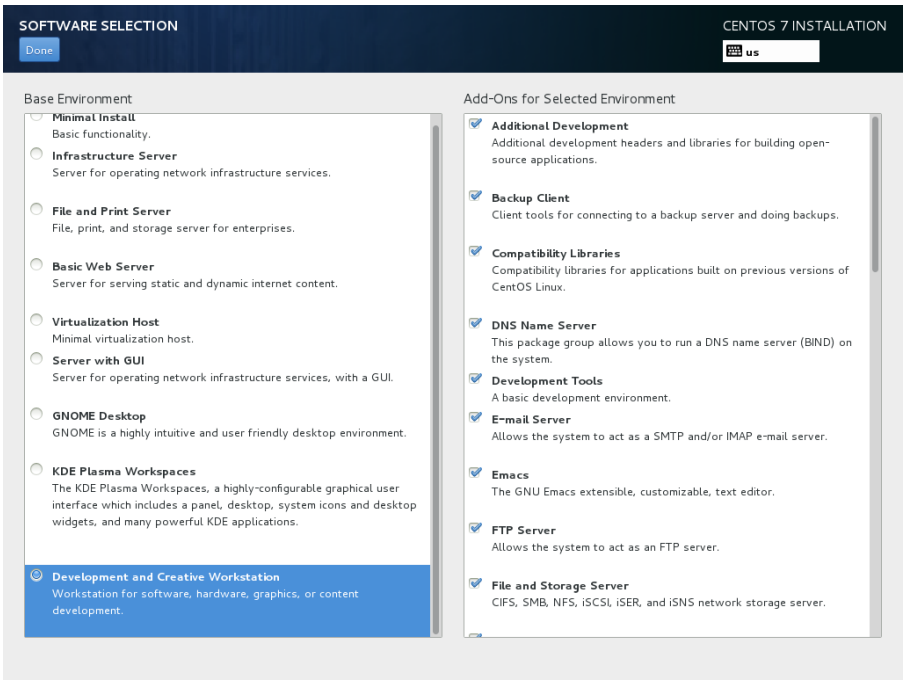

Choose Development and Creative Workstation and select all of the Add-Ons categories displayed on the right side of the screen.

NOTE

Click on the top add-on (Additional Development) and then repeat pressing the down arrow on the keyboard, followed by the space bar, until all add-ons are selected. Note that the list will continue to scroll across multiple pages of add-ons and will stop scrolling only after reaching the last add-on.

Then press Done in the upper left corner to complete software selection and return to the Installation Summary page. The Installation Summary page should eventually appear as shown below.

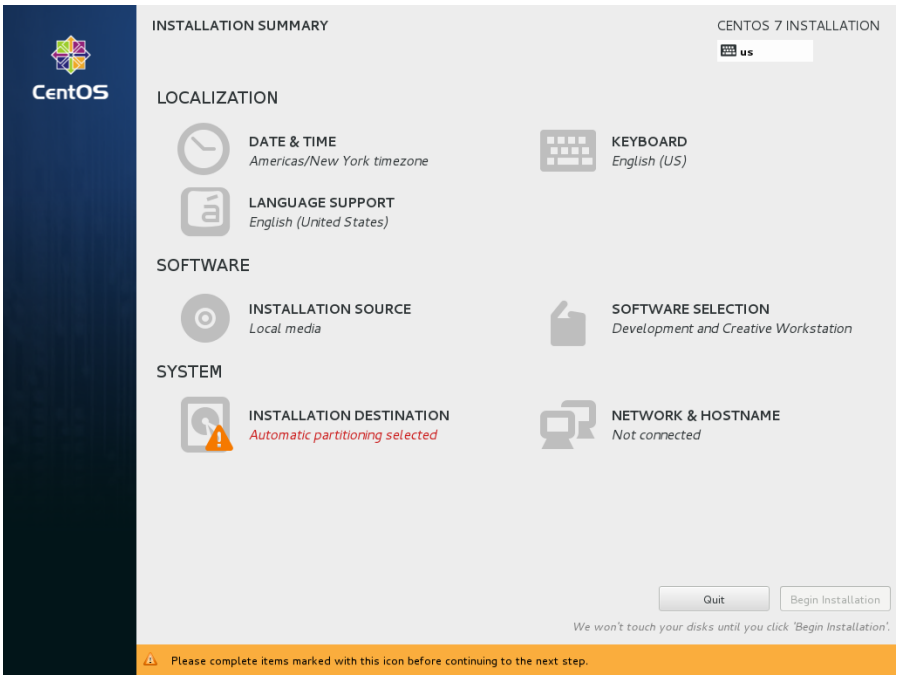

9. Next, click on Installation Destination to manually partition the system's hard drive for the best operation with RedHawk Linux. A page

similar to the following page will be displayed.

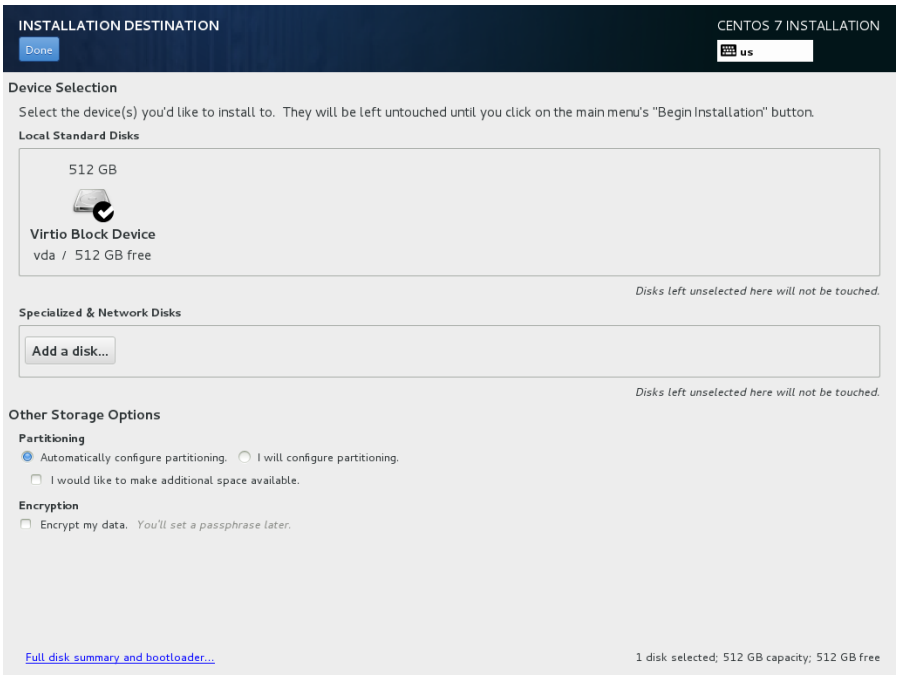

NOTE

If there is more than one hard drive attached to the system, you will see multiple hard drive icons displayed. In this case, select the hard drive where you wish to place the RedHawk installation, as well as any other drives that you wish to configure and/or format during the installation.

10. On this page, select I will configure partitioning under Other Storage Options and then press Done in the upper left to continue.

You will now be presented with the Manual Partitioning page, which should appear similar to the following.

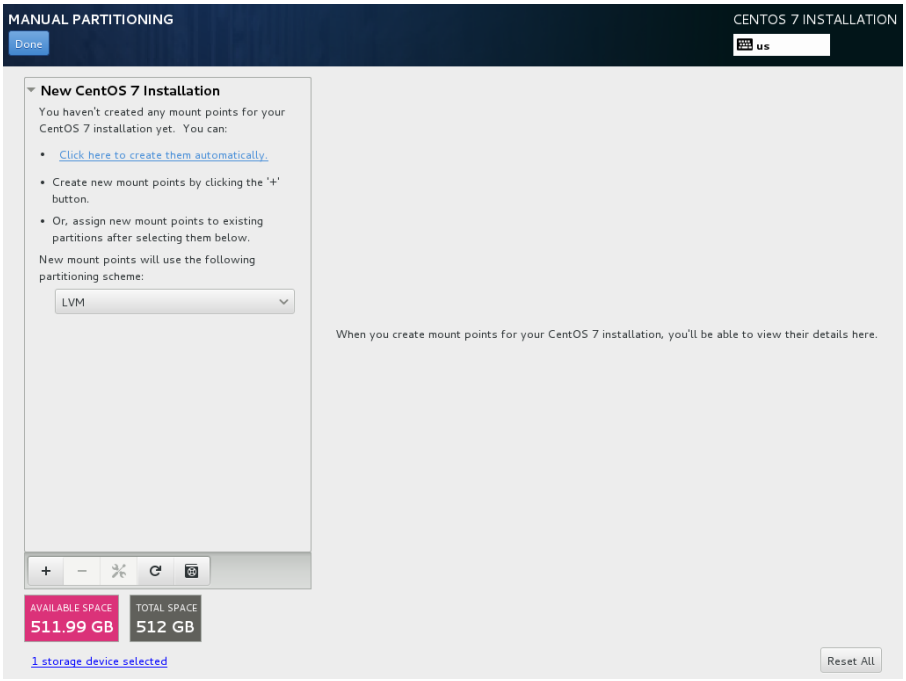

If the system's hard drive has been used previously, you should delete any existing partitions to start from a clean state.

NOTE

Deleting existing partitions is recommended unless you wish to create a multi-boot system. For example, one system with the ability to boot into multiple versions of RedHawk or possibly even Microsoft ® Windows ® operating systems. In these cases, consult with your system administrator as this document does not cover setting up a multi-boot system.

The following page illustrates selecting an existing `/boot` partition, clicking the "–" (minus) icon, checking the Delete all other filesystems checkbox, and finally pressing the Delete It button to deleting an entire exting partitioning scheme, if necessary.

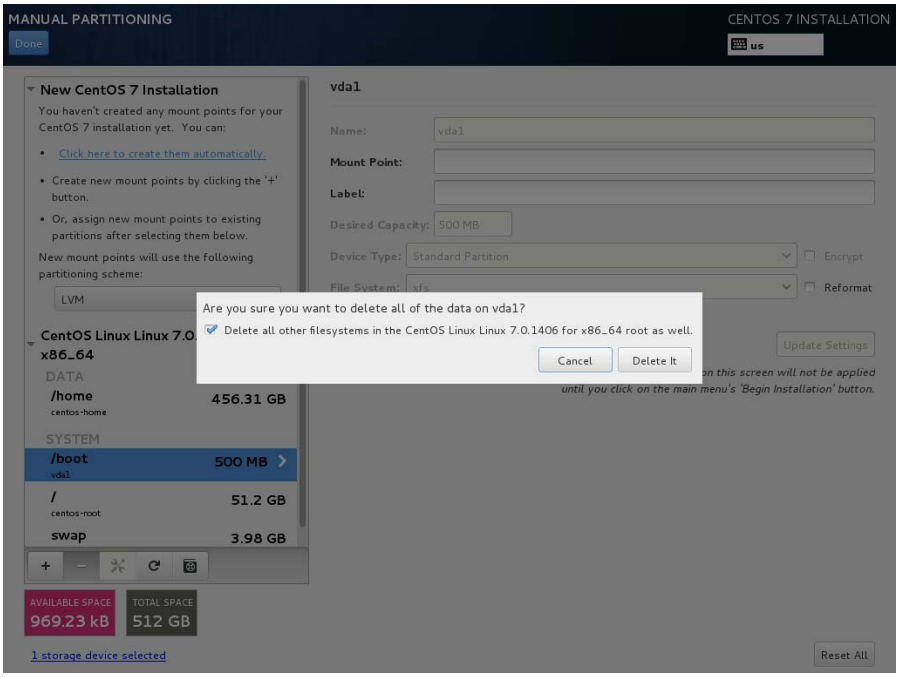

11. On the Manual Partitioning page, press the Click here to create them automatically link to automatically create the necessary partitions required for your installation.

    The following page shows an example of the automatically created partitions.

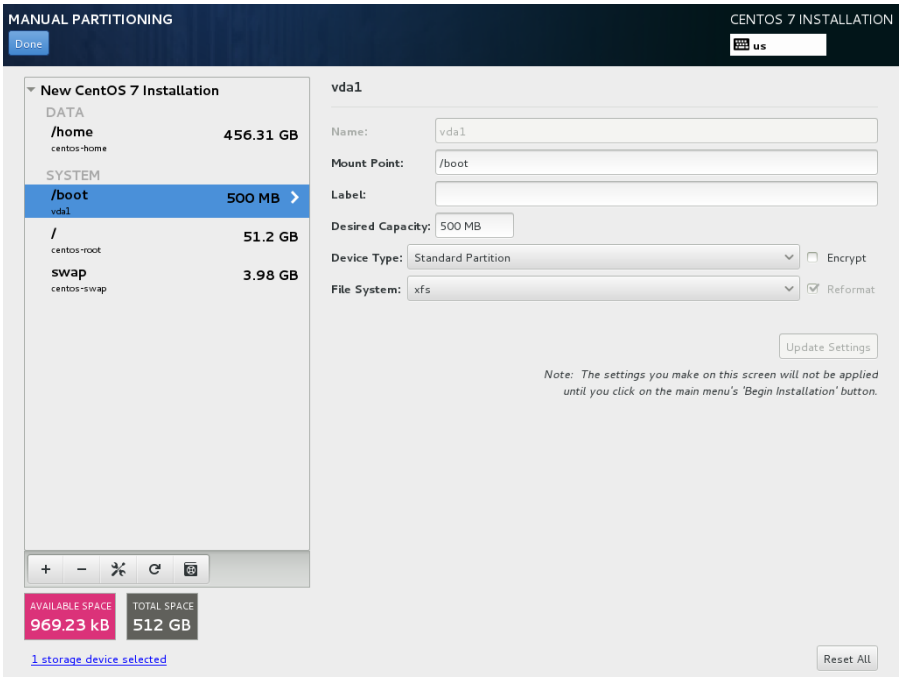

NOTE

The sizes of the automatically created partitions are inadequate; the following steps will adjust the partition sizes to be optimal for a RedHawk installation.

NOTE

Be careful not to accidentally press the Done or the Reset All buttons at this point. You must only press the Done button at the very end of modifying the partitions.

NOTE

On systems configured to use UEFI, an additional /boot/efi partition will be automatically created. This partition is not included in the following examples, however do not modify this partition as it is already configured properly for a RedHawk installation.

12. On the Manual Partitioning page, select the /home partition, change the Desired Capacity to 10 GB, and then press the Update Settings button.

    The following snippet shows an example of what should now be displayed:

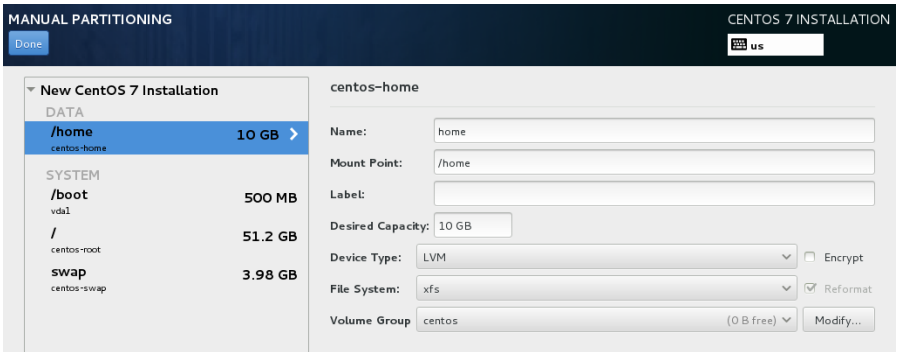

NOTE

Setting the /home partition to 10 GB is only a temporary change that is required in order to allow the other partitions to be resized.

13. Next, select the /boot partition, change the Desired Capacity from 500 MB to 4 GB, and then press the Update Settings button.

The following snippet shows an example of what should now be displayed:

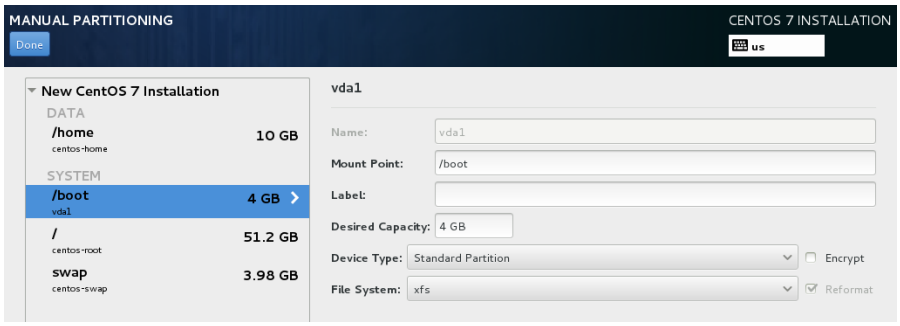

14. Next, select the / partition, change the Desired Capacity to 100 GB, and then press the Update Settings button.

    The following snippet shows an example of what should now be displayed:

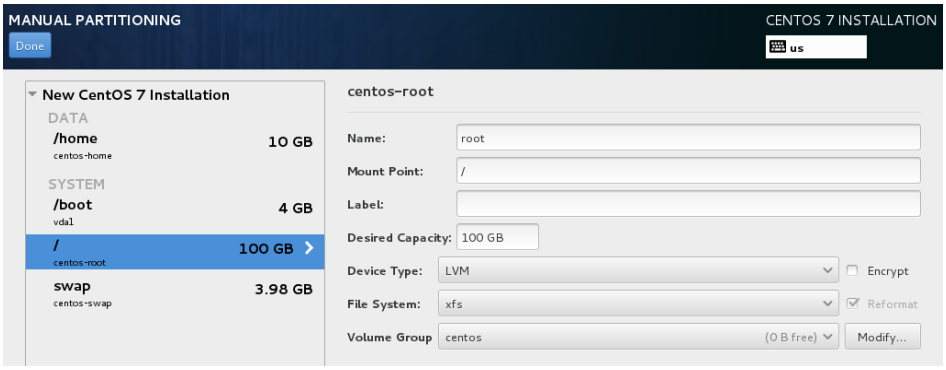

15. Next, select the swap partition, change the Desired Capacity to match the size recommended in the following table, and then press the Update Settings button.

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

The following snippet shows an example of what should now be displayed:

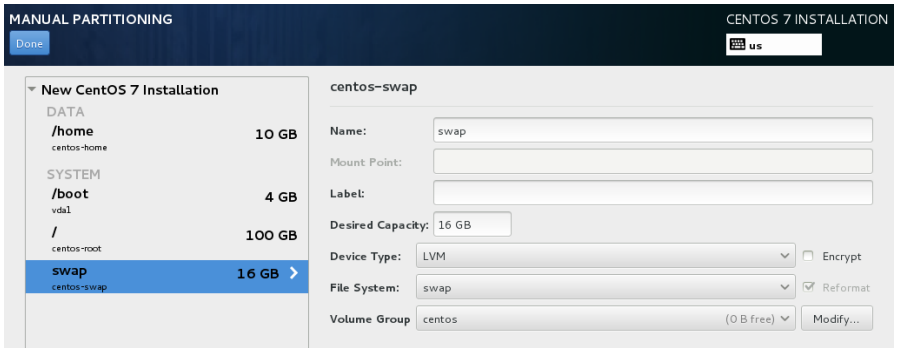

16. Lastly, select the /home partition again, delete all of the text in the Desired Capacity text field, and then press the Update Settings button to set the size to the maximum size available from the disk space remaining.

    The following snippet shows an example of what should now be displayed:

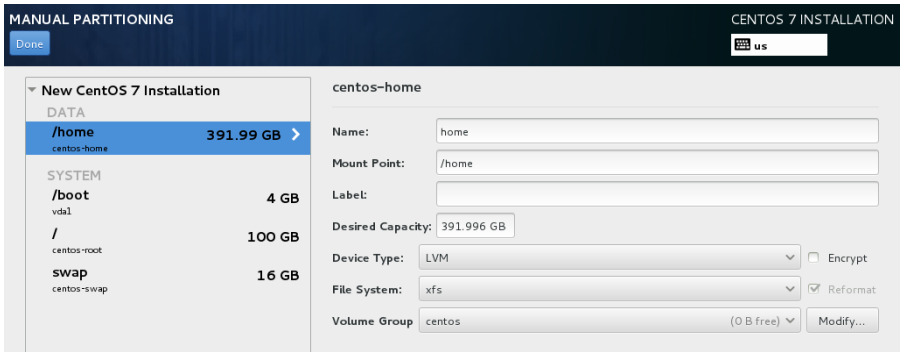

17. At this point, you should now see a Manual Partitioning page similar to the following displayed.

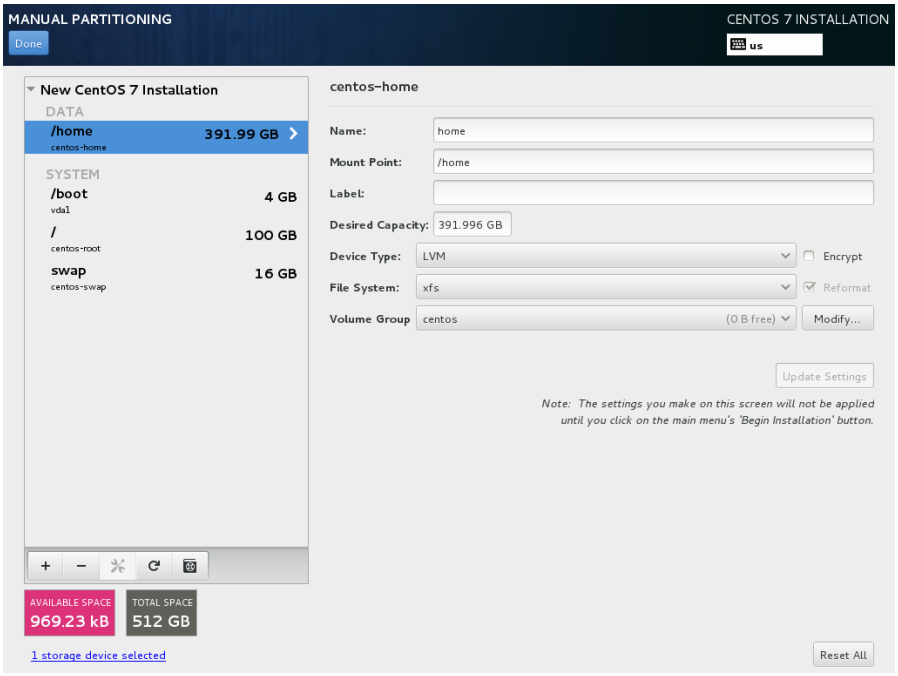

Finally, press the Done button to use this layout for the system's partitioning.

18. You will now be presented with a Summary of Changes confirmation screen similar to the following:

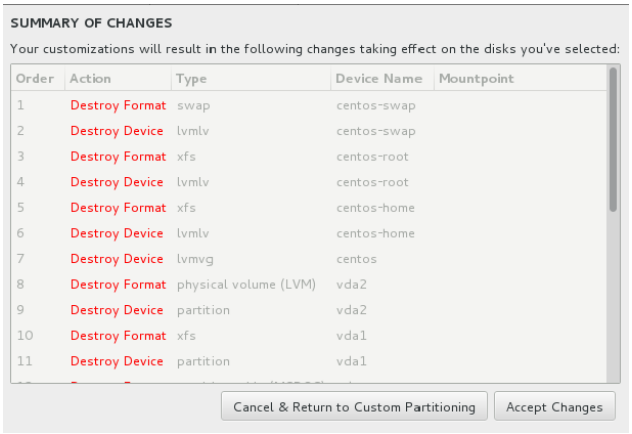

Next, press Accept Changes to accept the custom partitioning.

19. Back on the Installation Summary page, click on Network & Hostname to set the system's hostname and networking configuration, as shown in the following example screen:

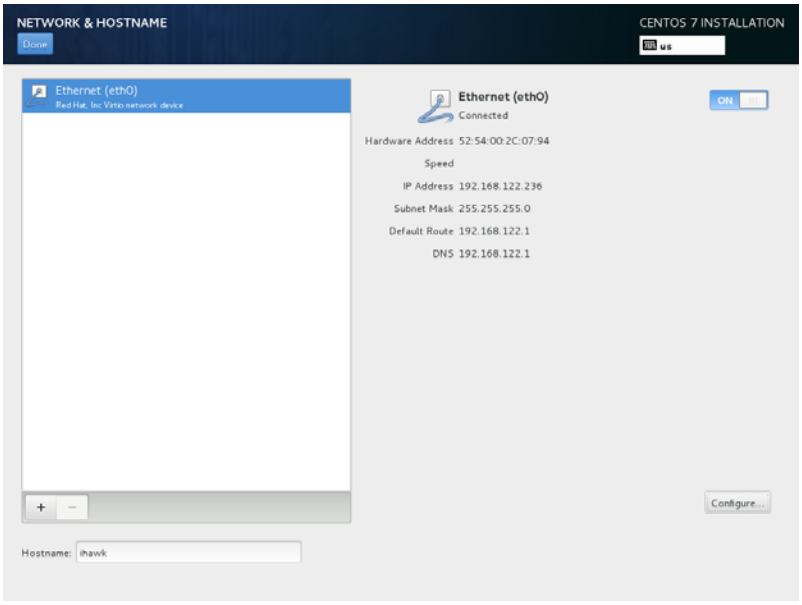

Press Done when you have finished configuring all of the networking details.

20. Back on the Installation Summary page, press the Begin Installation button on the lower-right of the screen to start package installation. During installation, you will be presented with a screen similar to the following.

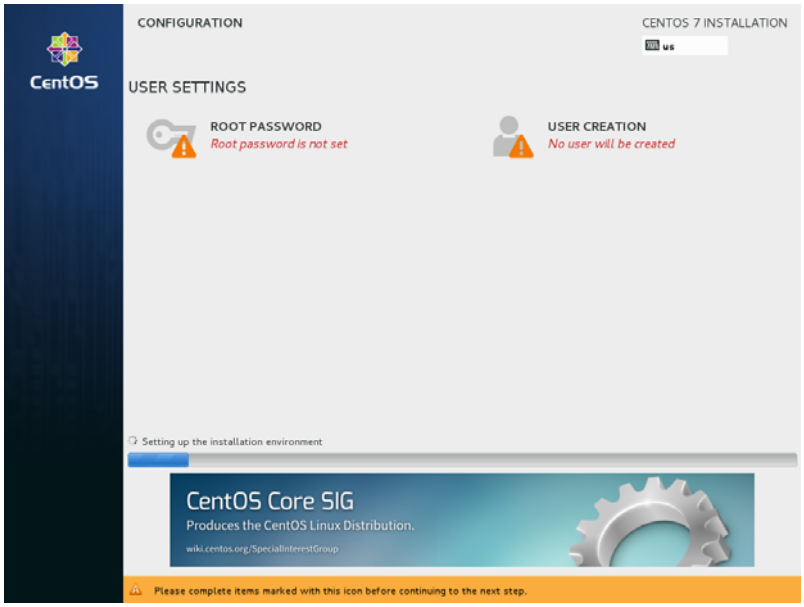

NOTE

During the system installation you must set the system's root password and also create a non-root user.

21. After all packages have been installed you will see the following screen.

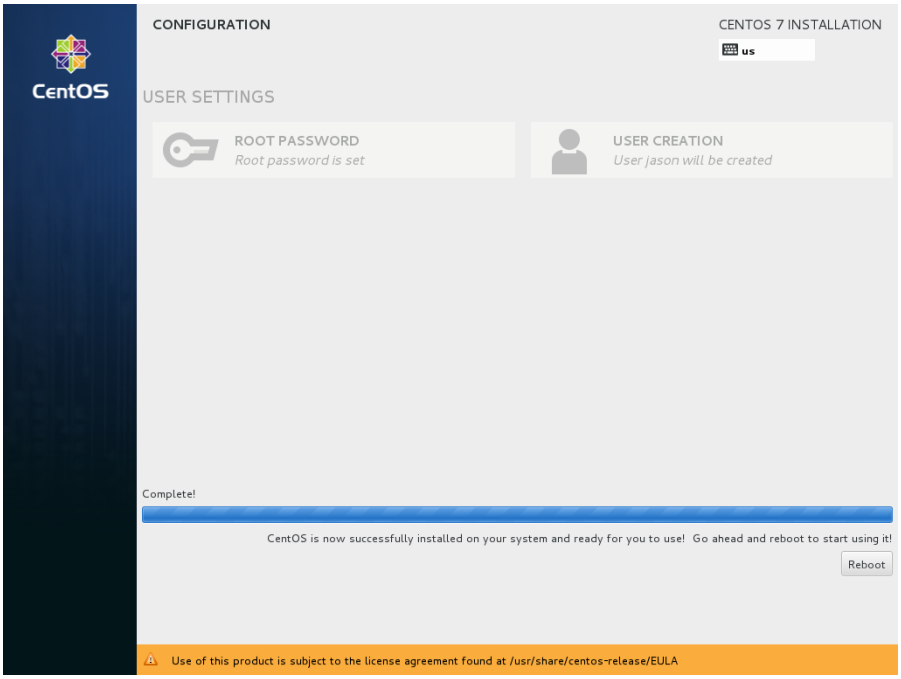

Remove the installation DVD and press Reboot in the lower right to restart the system and continue.

22. You should see the following grub screen displayed after the system restarts.

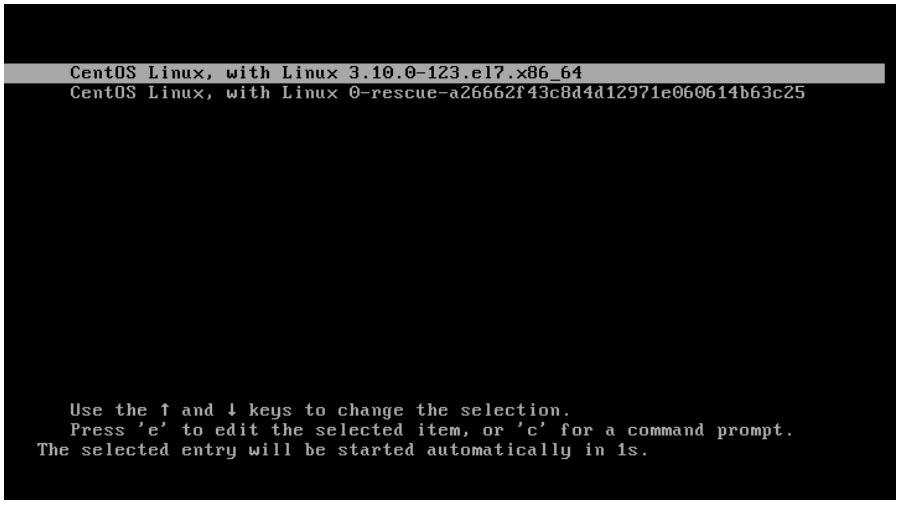

Press the Enter key or allow the grub screen to timeout to boot the newly installed system.

23. Once the system boots, you will be presented with the following setup screen.

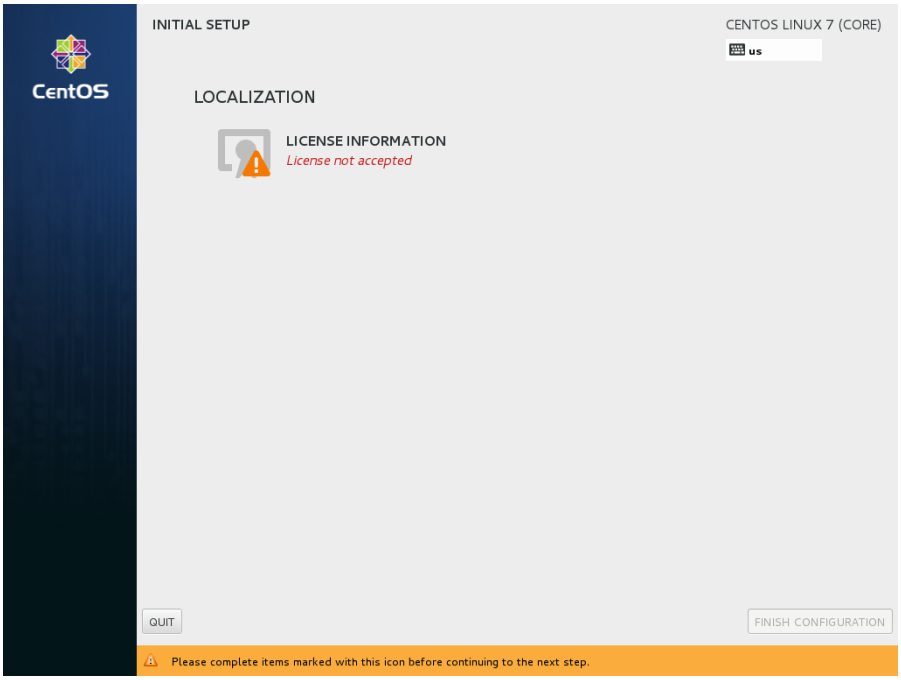

Click on License Information and the License Information screen will be displayed. Check the checkbox at the bottom left to accept the license. Then, click the Done button in the upper left, and back on the Initial Setup screen click Finish Configuration in the lower right.

24. The next screen prompts you to configure Kdump settings.

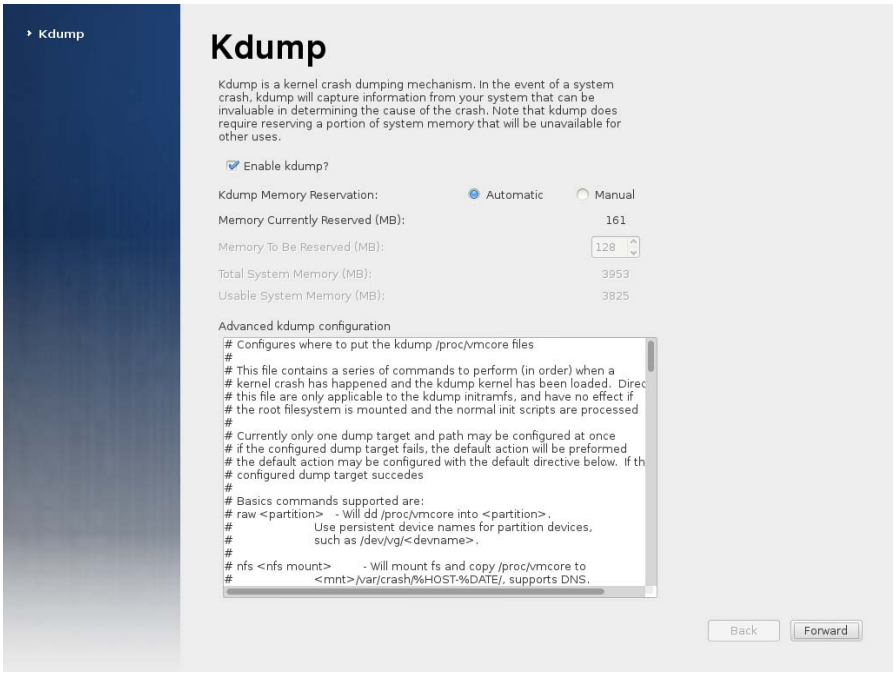

Click Forward and the login screen will appear.

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

5.1.2 Installing CentOS Updates

Install the disc labeled CentOS 7.0 Updates. This contains the latest updates released by CentOS 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`**.
   c. Enter the root password at the rescue prompt to access the shell.
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 **`/run`**, **`/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: Do *not* mount the device on the **`/mnt`** or **`/media`** directories; you must use a subdirectory of one of these directories. Failure to follow this advice may cause the installation to abort.

NOTE 3: 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
   ```

   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/CentOS7.0-Updates type iso9660 (ro,noexec)

   > mount -o remount,exec /media/CentOS7.0-Updates
   ```

4. The installation script will first ask you to confirm that you want to proceed with updating CentOS 7.0 packages. Type Y and press ENTER to continue.

   The installation script will then ask if you wish to install Concurrent's recommended set of updates for the CentOS 7.0 distribution. Concurrent recommends that you simply press ENTER to install the full set of Concurrent-recommended package updates.

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

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

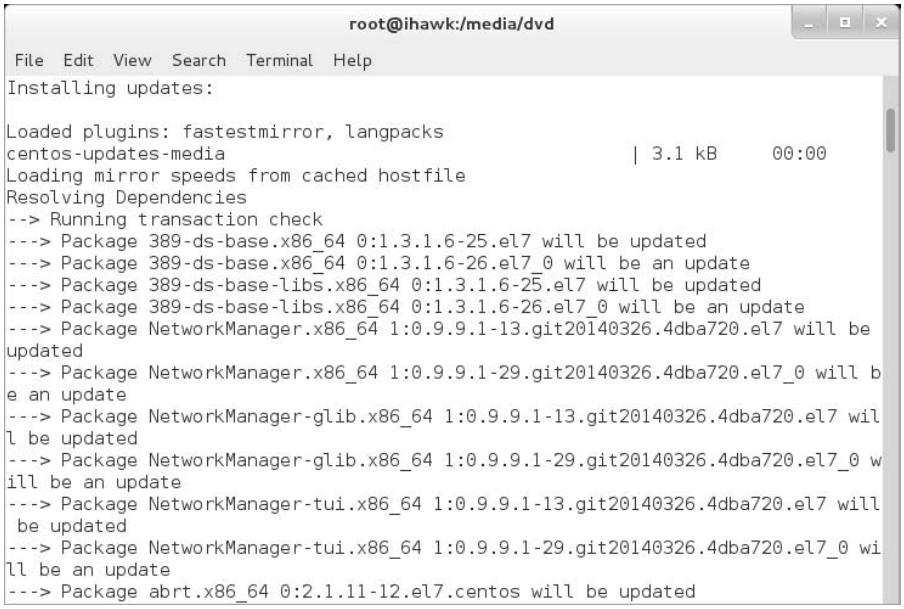

Once the package dependency transaction check has completed, updating of packages is executed. No action is required of you during this step.

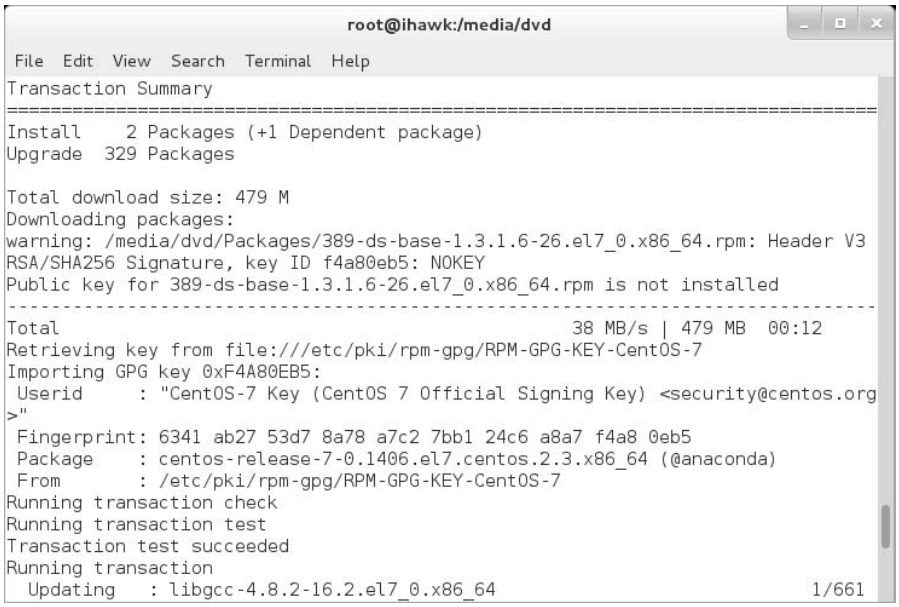

NOTE

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

6. Once all package updating has completed, many new packages are then installed.

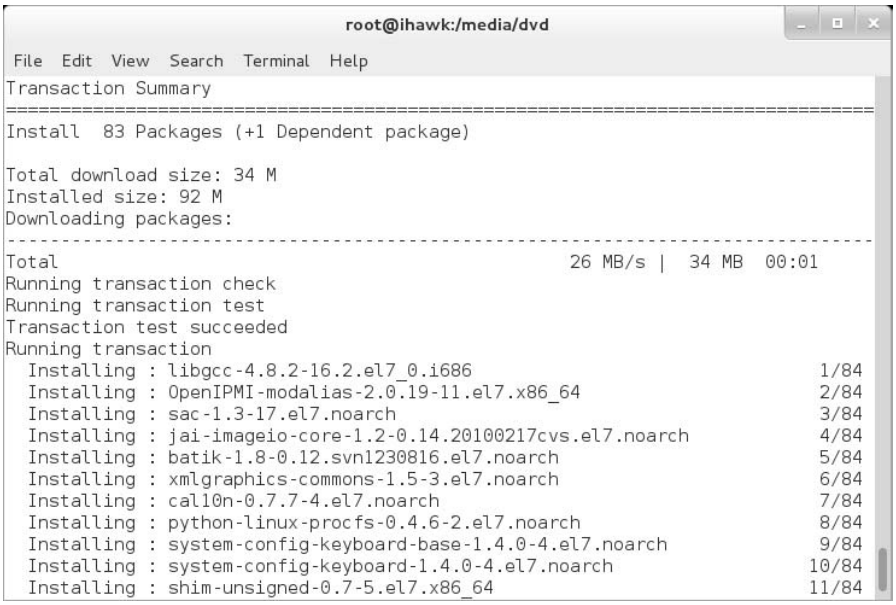

7. Upon completion of installation the entire update has completed.

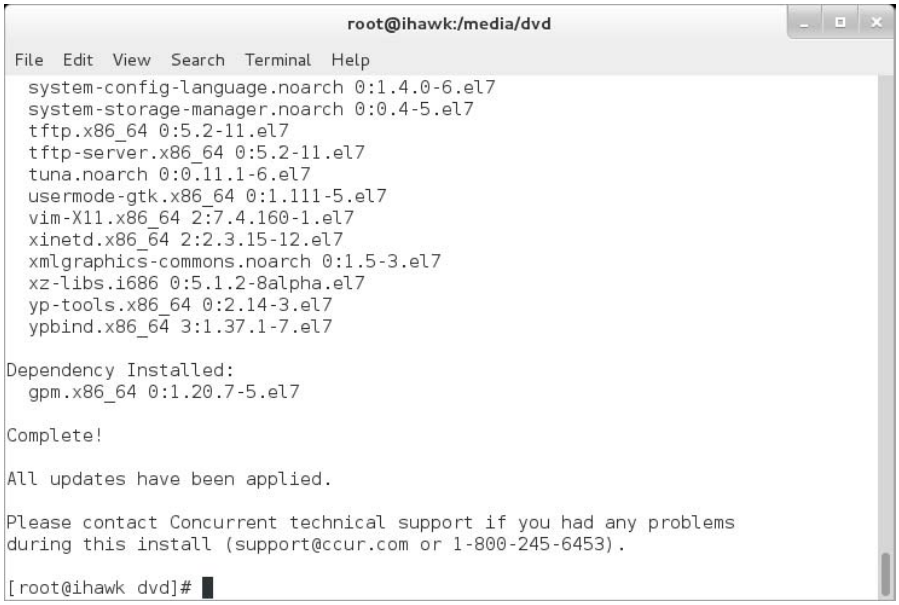

8. When the update completes, execute the following commands:

```
cd ..
eject
```

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

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

10. Type “reboot” at the system prompt and verify that the new CentOS kernel boots.

NOTE

During the first graphical boot after CentOS installation you will be presented with screens to accept the CentOS End-User License Agreement.

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

5.2 Installing RedHawk Linux

After completing the installation of CentOS 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:

NOTE

During each user's first graphical log in, the user will be presented with various screens to configure localization, create users, connect to external resources, etc. before they will be able to use the system.

   a. If in GUI mode, right click on the desktop and select Open Terminal.
   b. At the system prompt, type **`init 1`**.
   c. Enter the root password at the rescue prompt to access the shell.

2. Locate the disc labeled "RedHawk Linux OS" appropriate to your system's architecture and insert it into the optical media drive.

3. If the device does not mount automatically, execute the **`mount`** command appropriate for your optical media device.

4. Execute the following commands to install RedHawk Linux, substituting the actual mount point used for your optical media device for **/media/dvd** if necessary:

   cd /media/dvd
   ./install-redhawk

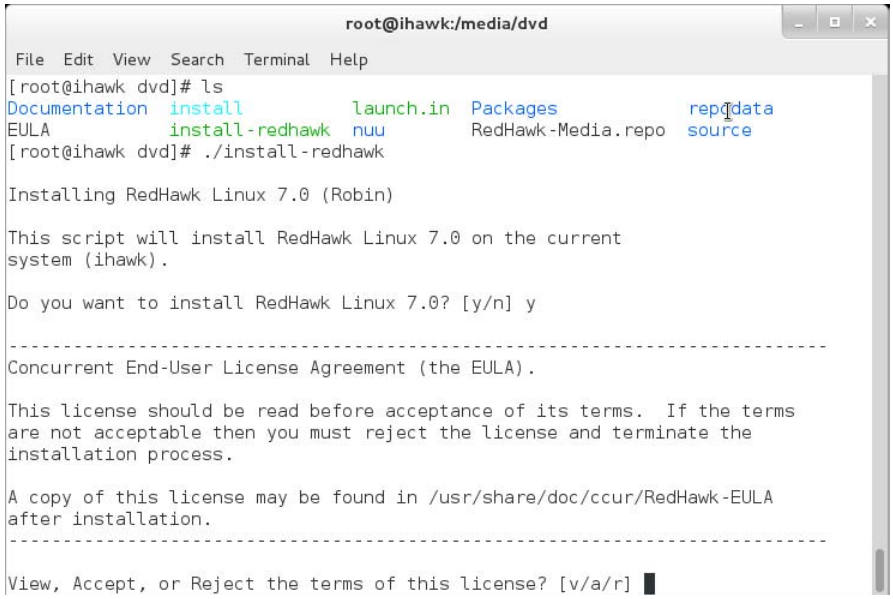

NOTES

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

NOTE 4: While the messages displayed during RedHawk installation appear to indicate that package *downloading* is occurring, no actual network access is performed when using the RedHawk DVD.

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

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

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

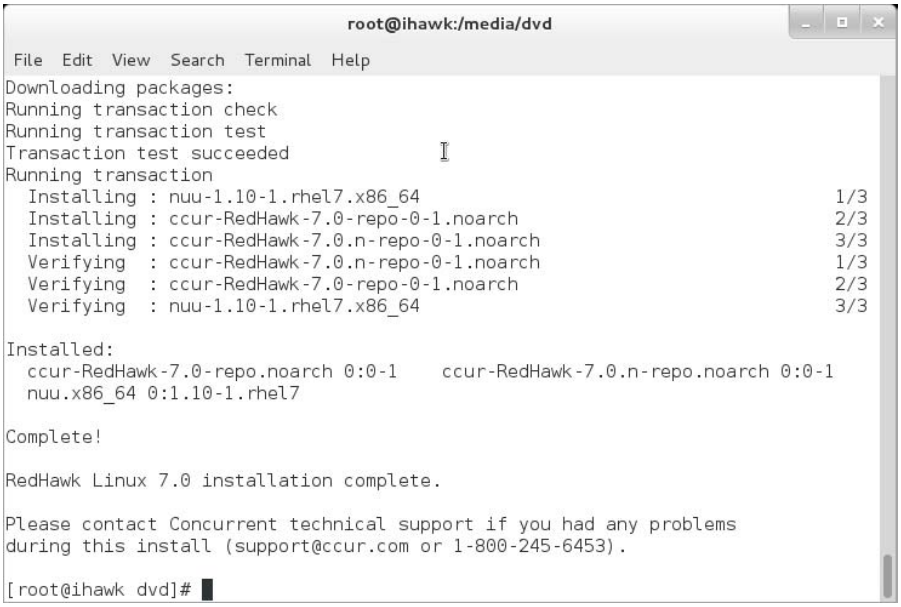

6. When the installation completes, execute the following commands:

```
cd ..
eject
```

7. Remove the disc from the drive and store.

8. Reboot the system choosing a new RedHawk kernel and verify that the system boots properly.

NOTE

During the first graphical boot after RedHawk installation you will be presented with screens to accept the Concurrent End User Agreement.

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

5.3 Installing the RCIM

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

5.3.1 Quick Hardware Installation Checklist

1. Before installing the RCIM, determine if you will be using the RCIM to accept or deliver external interrupts and the mode in which the RCIM will run (see "RCIM Connection Modes" on page 59).
2. Verify that the **`ccur-rcim`** RPM has been installed by executing the command **`rpm -q ccur-rcim`**. The system will inform you if it is not installed. It is a standard package in the RedHawk installation.
3. Power down the system and remove all power cords.

NOTE

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

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

5.3.2 Verifying RCIM Operation

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

   `cat /proc/driver/rcim/status`

   You should see an output similar to that shown below:

   ```
   RCIM-III board 0 is at revision 1 eeprom 2.0.
   RCIM firmware version 10
   This is a standalone (isolated) rcim.
   Has IRQ 108 and major number 253
   GPS present, RCIM firmware version 10
   ```

NOTE

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

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 and it is also included in the *extras* directory on the RedHawk product media. Execute the following steps to install this RPM:

   d. Mount the RedHawk product media and `cd` to the *extras* directory or download the **ccur-ntp** RPM from Concurrent's web site as described in section **9.2.2** on page 69.

   e. Install the RPM via the following command:

```
# rpm -Uvh 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 7.0 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 7.0" appropriate to your system's architecture and insert it into the optical media drive.
3. If the optical media disc device does not mount automatically, execute the **mount** command; for example:

```
mount /media/dvd
```

4. To install, execute the following commands:

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

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

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

```
cd ..
eject
```

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

5.5 Installing the PCI-to-VME Bridge

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

5.5.1 Installing the Hardware

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

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

1. To configure the jumpers on the VME adapter card correctly, refer to Chapter 10 of the *SBS Technologies Model 618-3 Hardware Manual* included with the RedHawk Linux documentation. The PDF file for this manual is available for viewing by clicking on the “Documents” icon on the desktop.
2. To install the PCI adapter in your iHawk system:
   a. Power down the system.
   b. Locate a vacant PCI card slot in the chassis that supports a bus master.
   c. Remove the metal plate that covers the cable exit at the rear of the chassis.
   d. Insert the PCI adapter card into the connector.
   e. Fasten the adapter card in place with the mounting screw.
   f. Replace the cover.
3. To install the VME adapter card:
   a. Ensure that the VME chassis is powered down.
   b. Decide whether the VME adapter card is the system controller. If so, it must be installed in slot 1. If not, locate an unoccupied 6U slot in the card cage for the adapter.
   c. Insert the card into the connector of the selected slot.

4. To install the fiber optic cable:
    a. Ensure that both systems are powered down.
    b. Remove the rubber boots on the fiber-optic transceivers as well as the ones on the fiber-optic cables. Be sure to replace these boots when cables are not in use.
    c. Plug one end of the fiber-optic cable into the PCI adapter card's transceiver.
    d. Plug the other end of the cable into the VME adapter card's transceiver.
    e. Power on both systems.
    f. Ensure that the READY LEDs on both adapter cards are lit.

5.5.2 Installing the Software

1. With RedHawk Linux Version 7.0 running on the system, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk Linux PCI-to-VME Bridge Library Version 7.0" appropriate to your system's architecture and insert it into the optical media drive.
3. If the optical media device does not mount automatically, execute the **mount** command; for example:

    ```
    mount /media/dvd
    ```

4. To install, execute the following commands:

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

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

    ```
    cd ..
    eject
    ```

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

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

5.6 Installing the PCIe-to-VME Bridge

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

5.6.1 Installing the Hardware

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

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

- A PCIe host adapter card that should be installed in the iHawk host. The PCIe host adapter card contains no jumpers or switches that require configuration.
- A VME bridge card that should be installed in the VME chassis. The VME bridge card contains 3.3V/5V key posts for the two on-board PMC sites. These posts are used to select the PCI signaling voltage to be used. Both posts must be set to the same position.
- A PCIe bridge cable that is used to connect these two boards.

1. To install the PCIe host adapter card in your iHawk system:
   a. Power down the system.
   b. Locate a vacant PCIe x4 (or greater) card slot in the chassis that supports a bus master.
   c. Remove the metal plate that covers the card slot at the rear of the chassis.
   d. Insert the PCIe adapter card into the slot.
   e. Fasten the adapter card in place with the mounting screw.
2. To install the VME bridge card:
   a. Ensure that the VME chassis is powered down.
   b. The VME bridge card should be installed in the left-most (system controller) slot in your VME chassis.
3. To install the cable:
   a. Ensure that both the iHawk system and VME chassis are still powered down.
   b. Plug one end of the cable into the PCI host adapter card's transceiver.
   c. Plug the other end of the cable into the VME bridge card's transceiver.
   d. Power on the iHawk system and VME chassis.

5.6.2 Installing the Software

1. With RedHawk Linux Version 7.0 running on the system, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk Linux PCIe-to-VME Bridge Version 7.0" appropriate to your system's architecture and insert it into the optical media driver.
3. If the optical media device does not mount automatically, execute the mount command; for example:

   ```
   mount /media/dvd
   ```

4. To install, execute the following commands:

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

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

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

```
cd ..
eject
```

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

7. Optionally configure the VME windows as described in the next section and then exit single-user mode (Ctrl-D).

5.6.3 Configuring the PCIe-to-VME Windows

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

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

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

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

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

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

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

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

0:BAR0: 0xF8400000 - 0xF8400FFF
0:SLAVES: 8

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

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

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

5.7 Installing Additional RedHawk Products

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

5.8 Installing RedHawk Updates from Disc

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

1. Ensure that the system is in multi-user mode.
2. Locate the disc labeled "Software Updates for RedHawk" appropriate to your system's architecture, insert it into the drive.
3. If the disc device does not mount automatically, execute the **mount** command; for example:

   ```
   mount /media/dvd
   ```

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

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

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

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

   NOTE

   You should disable the external CentOS repos (Base, Updates and Contrib) before continuing. Select Edit Configuration in the Repositories menu to disable these repositories.

   With NUU up-to-date, from the NUU Main window, click on the Select Update of All Out-of-Date button, then click on the Apply Actions button to install all listed product updates.

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

```
cd ..
eject
```

6. Remove the disc from the optical media drive and store.
7. If there are more update discs, repeat these steps to install all updates.
8. Reboot the system choosing the updated RedHawk kernel and verify that the system boots properly.
9. Check for additional RedHawk updates available through the RedHawk Updates web site as outlined in the following section.

5.9 Installing RedHawk Updates from the Updates Web Site

1. Additional updates may have been issued (that are not included on any "Software Updates for RedHawk" disc you may have received). They should be downloaded and installed now. Refer to "Software Updates" on page 68 for details.
2. When complete, or if no updates are available, reboot the system choosing the new RedHawk kernel and verify that the system boots properly.

6.0. NVIDIA Graphics Configuration

The Unified NVIDIA Linux Display Driver (version 346.35) is included in the RedHawk Linux installation and built automatically for all pre-built and custom RedHawk Linux kernels. The RedHawk installer probes for the type and model of video adapters present in the system, and recommends installing the unified **`ccur-nvidia-glx`** rpm when it detects one or more NVIDIA video cards present.

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

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

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

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

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

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

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

6.1 Supported NVIDIA GPUs

If the **`ccur-nvidia-glx`** rpm has been installed on your system, you can view the list of all NVIDIA GPUs supported by the 346.35 driver in “Appendix A” in the following file:

`/usr/share/doc/ccur-nvidia-glx-346-35/README.txt`

Alternatively, you can examine the list of supported GPUs for driver 346.35 at the NVIDIA web site which is located at:

http://www.nvidia.com

Note that if your NVIDIA GPU is not supported by the 346.35 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 346.35 driver
- go to the NVIDIA web site mentioned above and download a legacy NVIDIA driver that supports your NVIDIA card, and follow the installation instructions described on that driver’s web page

NOTE

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

6.2 Multiple NVIDIA Video Card Considerations

Note that you may not use the unified 346.35 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 346.35 driver as your display driver.

6.3 Display Configurations

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

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

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

6.3.1 Single

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

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

6.3.2 Xinerama

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

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

- This mode causes a little extra overhead for X.

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

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

6.3.3 MultiDesktop

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

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

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

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

6.3.4 Twinview

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

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

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

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

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

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

See **/usr/share/doc/ccur-nvidia-glx-346-35/README.txt** for a detailed discussion on Twinview.

6.3.5 Twinview-Xinerama

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

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

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

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

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

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

6.3.6 Twinview-MultiDesktop

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

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

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

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

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

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

6.4 Using nvidia-xconfig

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

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

Follow this procedure exactly to generate the file:

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

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

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

```
nvidia-xconfig
```

5. Reboot to runlevel 5.

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

6.5 Using nvidia-settings

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

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

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

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

```
nvidia-settings
```

The GUI will start:

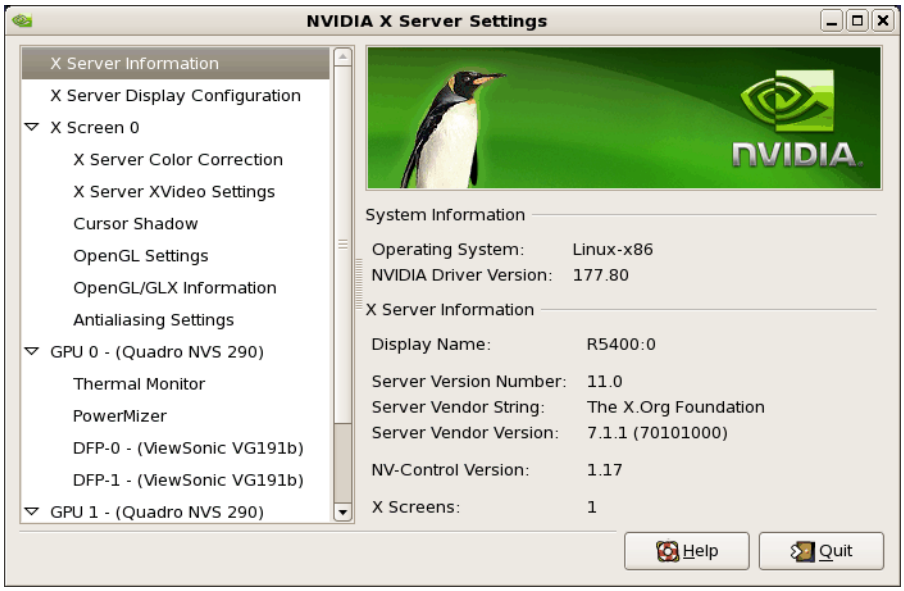

Here we see that we are viewing “X Server Information”, as indicated by the list of options on the left.

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

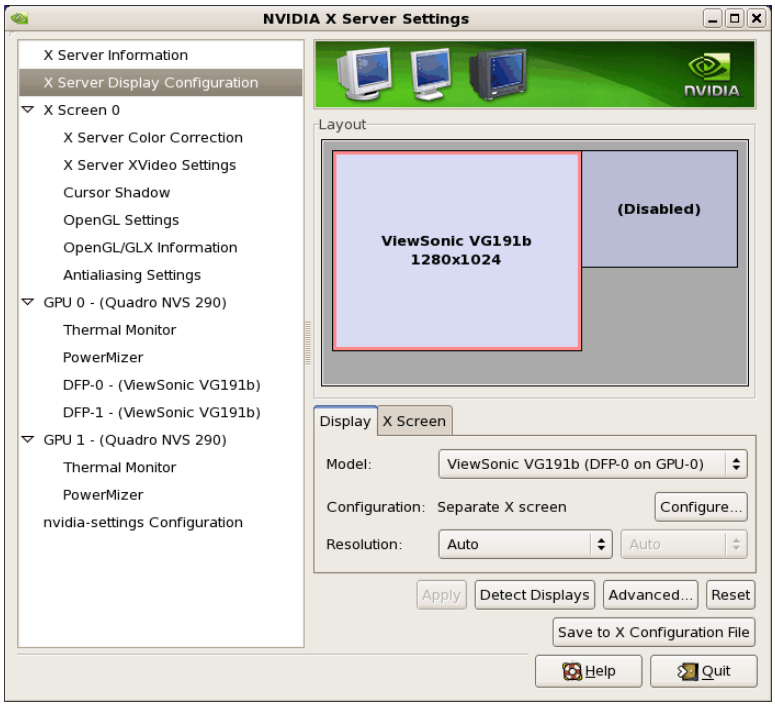

In the Layout window, we see a graphical depiction of the number of displays detected and configured. Notice that the large box is the primary monitor already configured by **`nvidia-xconfig`**. The smaller box to the right (Disabled) is a display that has not yet been configured.

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:

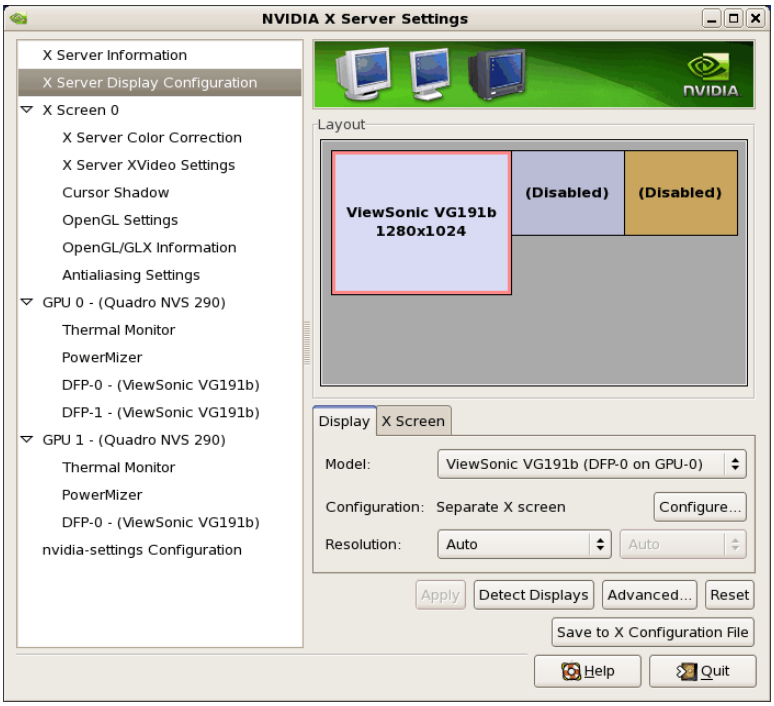

3. To configure two detected displays as a MultiDesktop, we would select the second monitor in the Layout window by clicking on it. We can see the second monitor is selected for configuration by the red outline around the box:

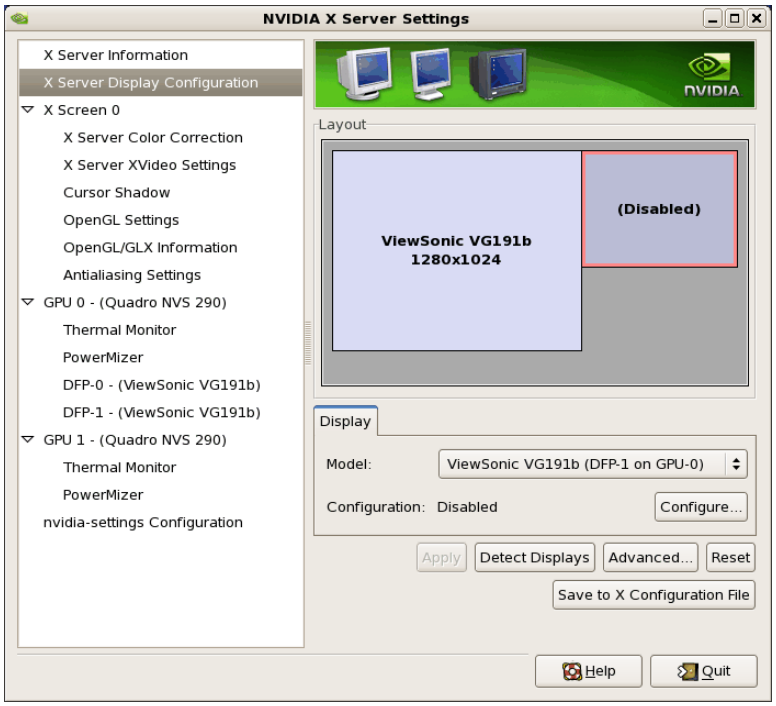

4. Click on the Configure button below the Layout window and select the "Separate X screen" option, then click OK:

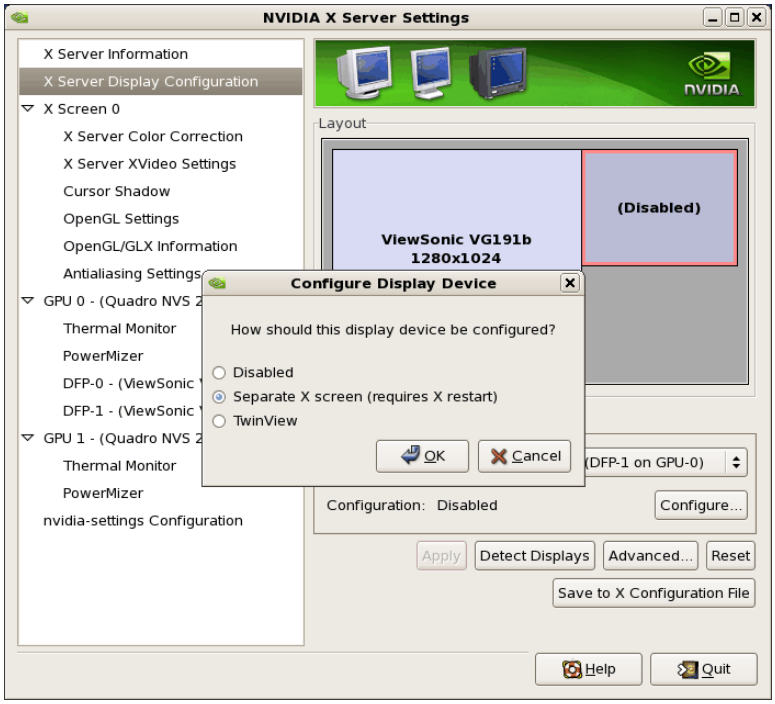

5. If we wanted to enable Xinerama we could do so at this point by selecting the "Enable Xinerama" check box below the Layout window:

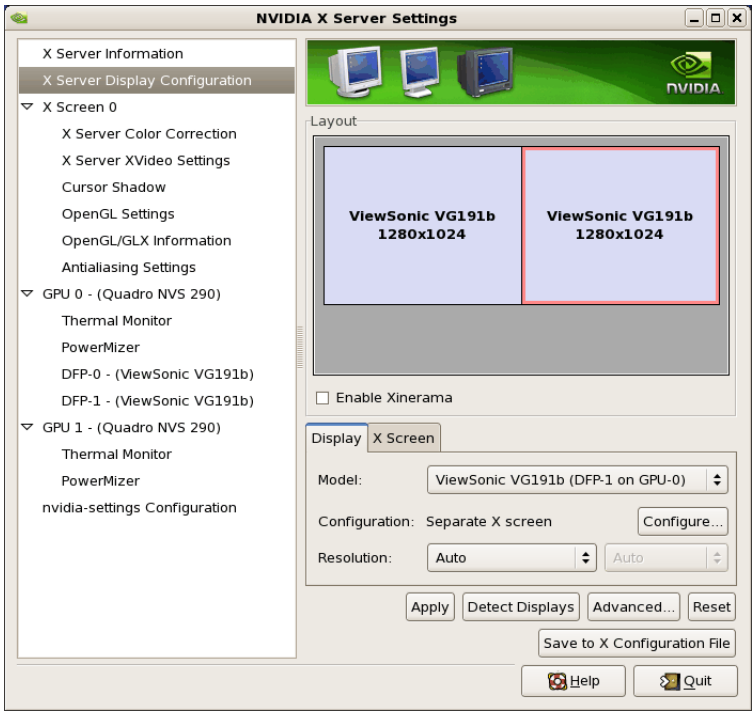

6. To save the configuration to **`xorg.conf`**, click on the Save to X Configuration File button near the bottom of the GUI:

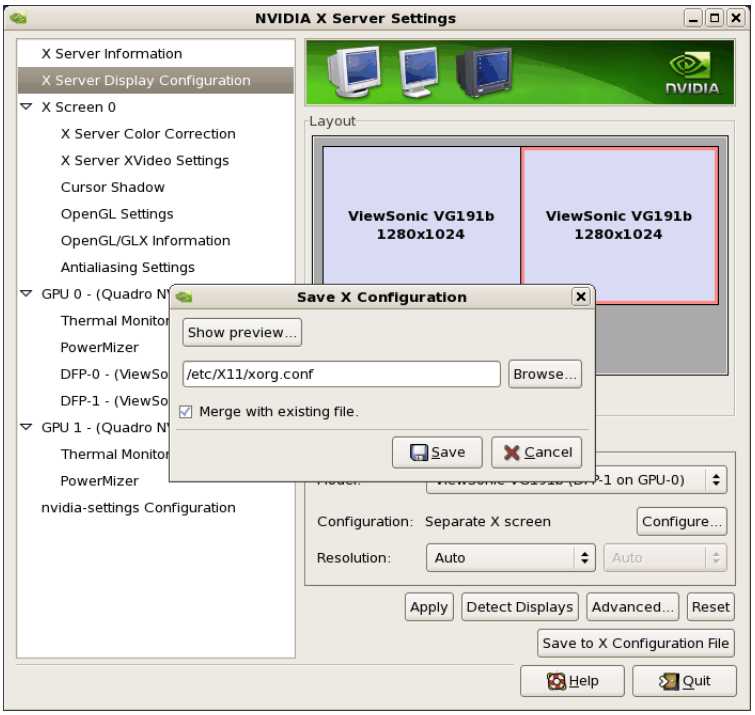

Then click Save on the Save X Configuration dialog box as shown above. (Always leave the Merge with existing file box checked).

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

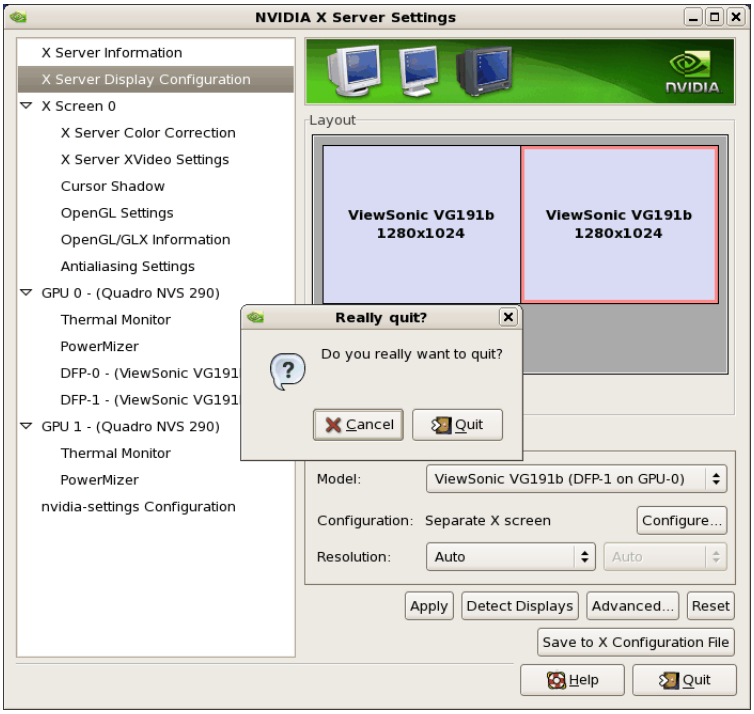

8. Reboot to runlevel 5. You should now be up on two X screens. You can restart **`nvidia-settings`** to further examine or modify your configuration.

Below shows an example configuring the Xinerama option:

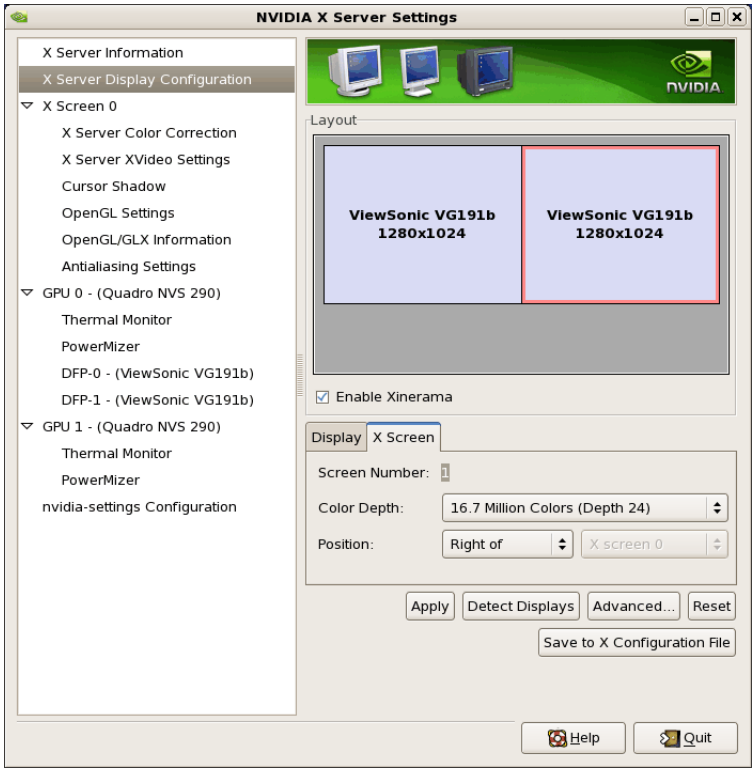

Below shows an example configuring the Twinview option:

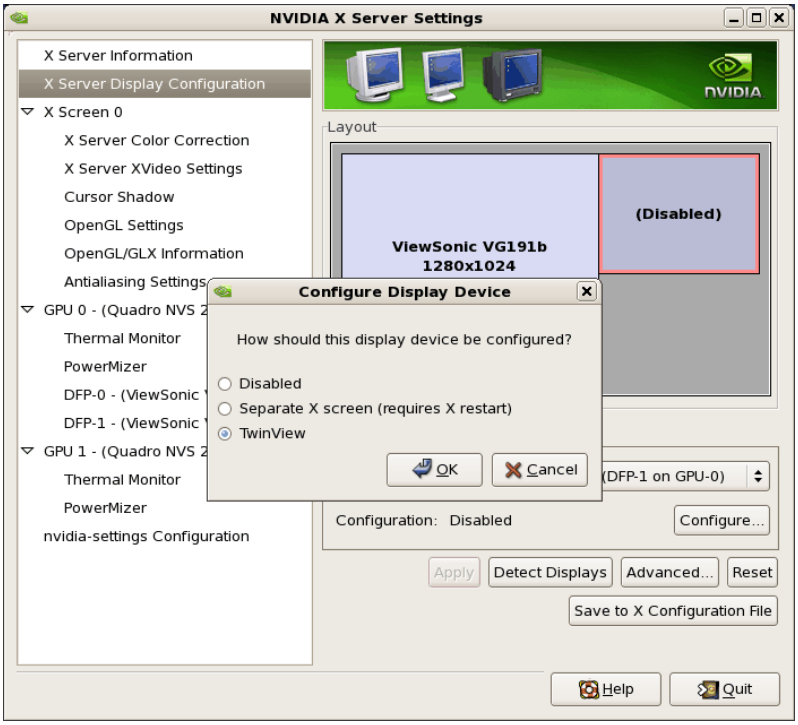

Below shows some stages configuring Twinview-Xinerama or Twinview-MultipleDesktop (the difference would be in selecting or not selecting the Xinerama check box):

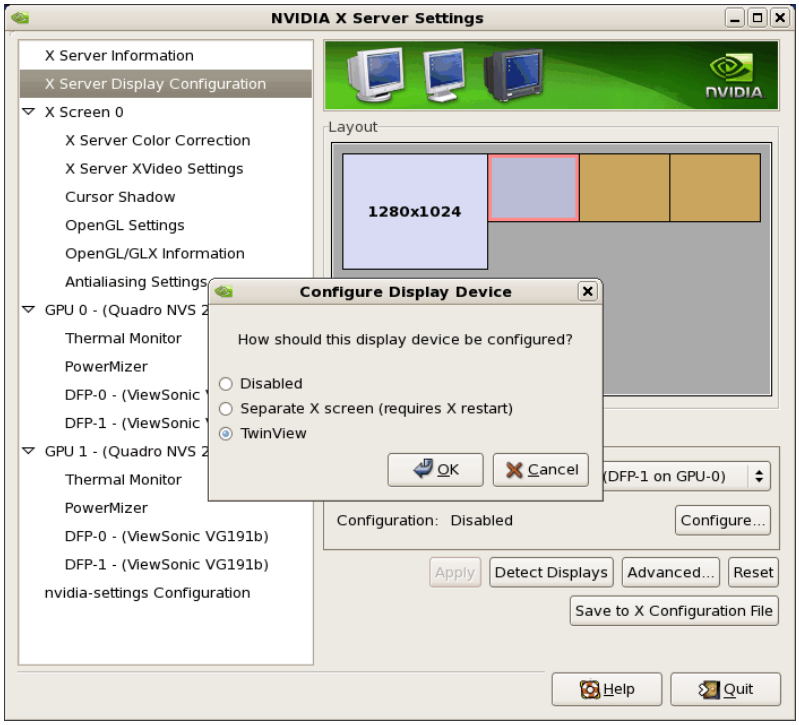

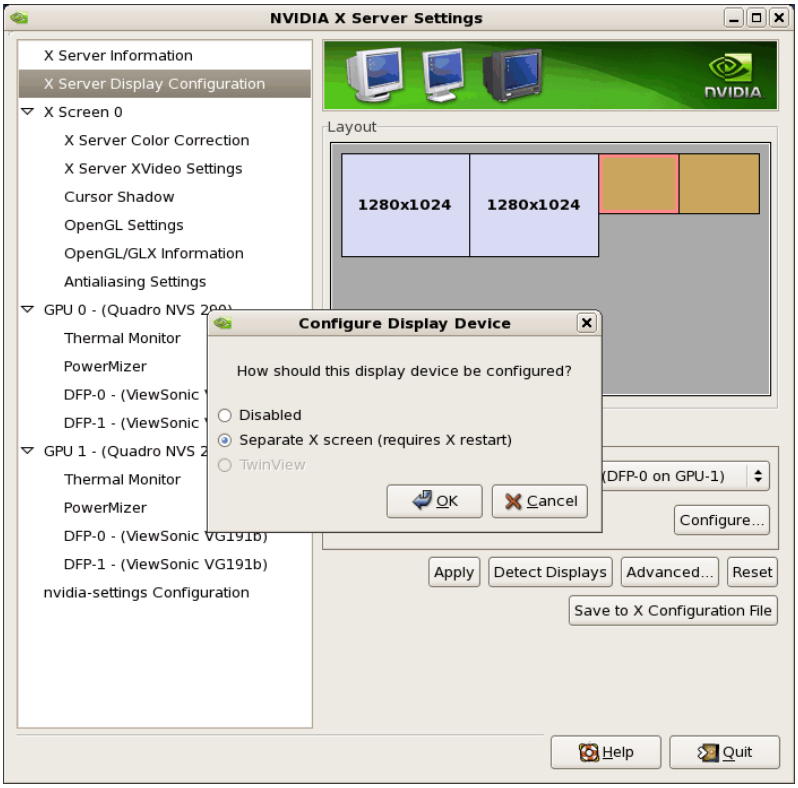

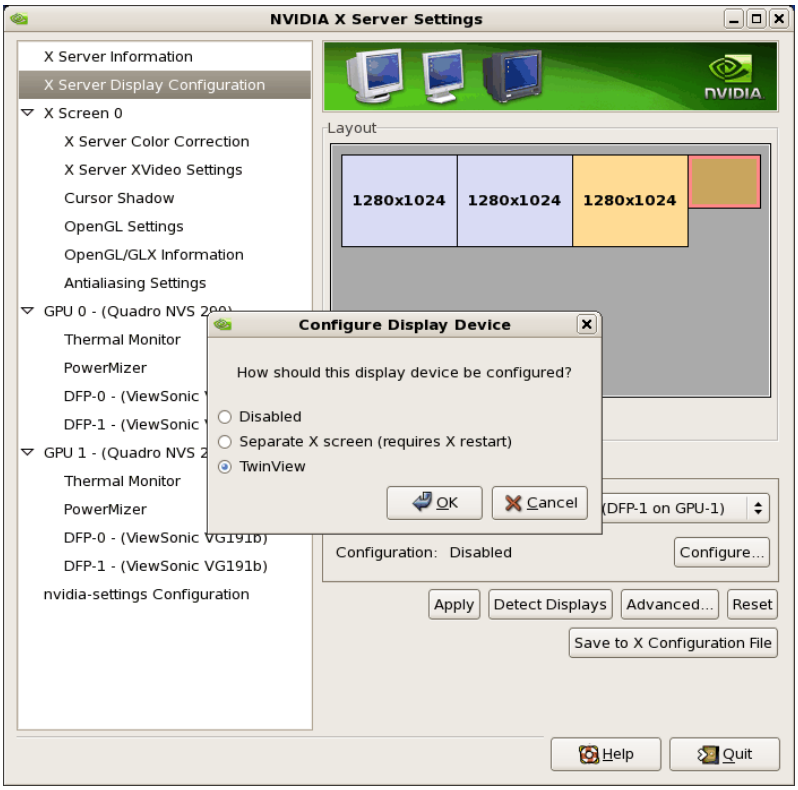

Note that **`nvidia-settings`** will allow you to do much more than covered here.

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

6.6 Modified xorg.conf file generated by nvidia-xconfig

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

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

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

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

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

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

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

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

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

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

Section "Monitor"
     Identifier        "Monitor0"
EndSection

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

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

6.7 Troubleshooting Problems

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

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

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

```
init 1; init 5
```

or

```
reboot
```

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

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

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

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

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

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

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

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

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

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

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

Typical drivers in iHawk systems include:

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

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

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

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

WARNING

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

7.0. Additional Information

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

7.1 Installation/Configuration Issues

7.1.1 File System Backup Recommendations

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

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

Table 7-1 Files Recommended for Backup

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

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

7.1.2 File System Swap Size Guidelines

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

Table 7-2 Swap Partition Size Guidelines

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

7.1.3 RCIM Connection Modes

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

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

7.1.4 How to Use Linux Rescue

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

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

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

It is possible to recover from these situations by using “linux rescue”. “Linux rescue” is a boot option on the CentOS 7.0 installation media.

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

1. Boot a CentOS 6.x 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 CentOS 7.0 installation.

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

Use the following procedures to make corrections:

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

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

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

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

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

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

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

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

  tune2fs /dev/sda2 -L *newlabelname*
- It is possible to configure grub menus to boot correctly off multiple disk(s).

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

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

title RedHawk Linux 7.0-20091211 (Trace=Yes, Debug=No)
        root (hd1,0)
        kernel /vmlinuz-2.6.31.6-RedHawk-5.4-trace ro
root=/dev/sdb2
```

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

7.1.5 How to Install CentOS Remotely Over VNC

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

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

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

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

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

Perform the following steps to install using VNC.

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

1. Follow the procedure outlined in section **5.0. Installation Procedures** on page 13 until you reach **5.1.1** step 4.

   Select “Install CentOS 7” from the grub menu.

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

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

   Hit the Enter key.

NOTE

If the VNC service fails to start, it may be necessary to specify the `ip=`*`xxx.xxx.xxx.xxx`* and `netmask=`*`xxx.xxx.xxx.xxx`* grub options to ensure proper NIC address configuration.

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

```
Starting VNC ...

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

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

Starting graphical installation.
```

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

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

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

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

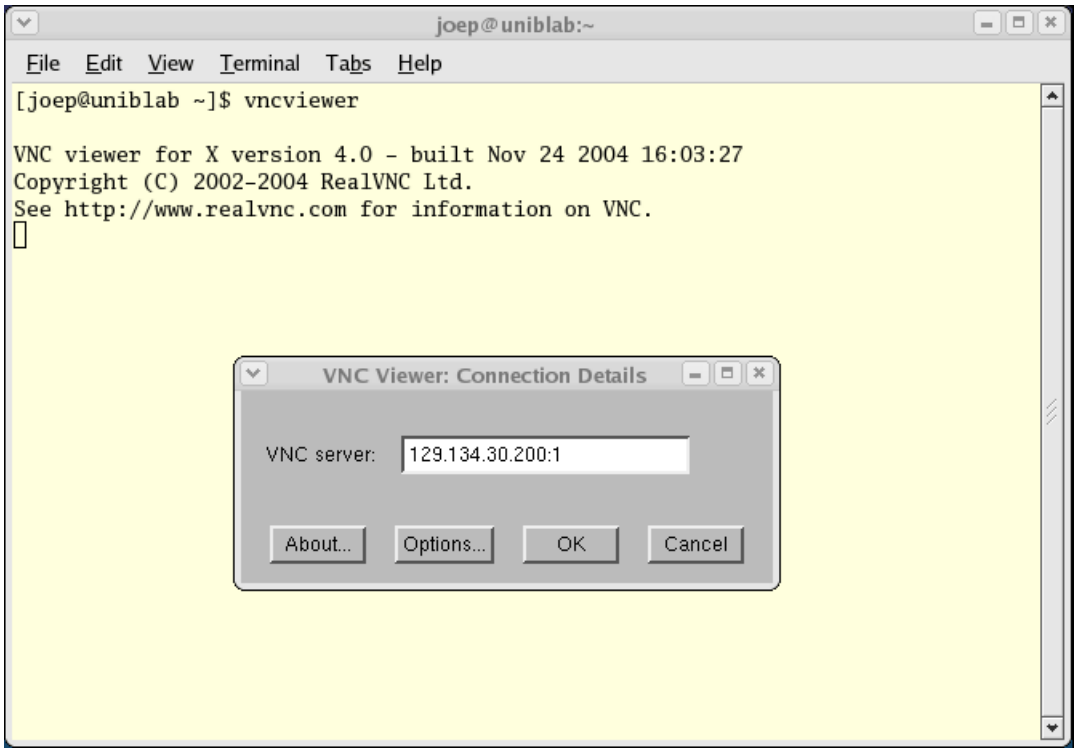

9. Select OK and press Enter to continue. At this point, you are able to perform a graphical installation remotely.

10. Continue with the graphical installation process as described in section **5.1.1** steps 6 - 21, beginning on page 15. After completing step 23 on page page 27, perform the steps below.

11. The reboot puts the system in runlevel 3 and a text-based Initial setup prompt is displayed.

12. Type 1 and press Enter to proceed to the License information prompt.

13. Type 2 and press Enter to select I accept the license agreement.

14. Type c and press Enter to complete license acceptance and return to the Initial setup prompt.

15. Type c and press Enter again to complete the Initial setup; the system will now proceed to the customary login prompt.

16. Log in as the root user and issue the following commands to enable graphical startup and reboot the system:

```
# systemctl set-default graphical.target
# reboot
```

17. After the system reboots, continue with the installation as documented.

7.1.6 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:

  Video problems can often be resolved simply by forcing Linux to use a generic graphics mode with the basic vesa driver. This can be accomplished via the **nomodeset** option. For exampe:

  ```
  > vmlinux ... nomodeset
  ```

  Alternatively, specify a option **driver=***xxx*, where *xxx* is the name of the driver matching the system's video hardware. For example:

  ```
  NVIDIA:    > vmlinux ... driver=nv
  ATI:       > vmlinux ... driver=radeon
  Generic:   > vmlinux ... driver=vesa
  ```

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

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

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

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

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

7.2 Considerations

7.2.1 Compiler Requirements

If you plan to build a kernel from source residing on your iHawk system on a different system, that system must be using **gcc-4.4.4** or later.

7.2.2 Hyper-thread Activation

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

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

8.0. Known Issues

Special consideration should be given to the following areas.

BIOS “Console Redirection”

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

NMI Button

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

PAM Capabilities

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:

```
systemctl disable krb5-telnet
```

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 CentOS 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:

systemctl {**enable**|**disable**} **irqbalance**

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

systemctl {**start**|**stop**} **irqbalance**

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 7.0 is based on CentOS Enterprise Linux 7.0. Binary and source compatibility issues exist for MAXAda programs between CentOS 3.0 and CentOS 7.0.

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

Inactive Virtual Terminal State After Changing Run Levels

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

Excessive ksoftirqd Runs Affecting Determinism

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

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

Pings to Multicast Addresses Disabled by Default

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

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

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

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

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

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

     net.ipv4.icmp_echo_ignore_broadcasts = 0
  ```

Grub Option Holdovers from Previous Kernel Installations

Be aware that when kernel packages are installed, they add a grub entry with options associated with that kernel (e.g., trace and debug kernels add “crashkernel=128M@64M,” other kernels add “quiet,” etc.). In addition, all kernel packages also copy the default grub options into their grub entries. These default grub options are taken from the kernel entry that is currently marked as the default kernel to boot in **`grub.conf`**.

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

USB Ports Non-functional on Supermicro Boards

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

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

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

Configure Quad Port Ethernet

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

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

NFS v4 Intermittent Connectivity

On rare occasions, attempts to use NFS version 4 protocol to connect to an older NFS server may be denied intermittently, with the connection also subsequently failing to negotiate an NFS version 3 protocol connection. If your system experiences this behavior with any NFS servers, addding a `vers=3` mount option to the corresponding entry in `/etc/fstab` will allow the NFS filesystem to be mounted properly.

9.0. Software Updates and Support

9.1 Direct Software Support

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

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

You may also submit a request for assistance at any time by using the Concurrent Computer Corporation web site at **http://real-time.ccur.com/home/customer-support/customer-support-form**.

9.2 Software Updates

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

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

9.2.1 Updating via NUU

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

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

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

NOTE

It is recommended that all CentOS repositories should be disabled when checking for Concurrent updates. In NUU, select the Repositories -> Edit Configuration menu item and ensure that the *base*, *updates* and *extras* repositories are disabled.

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 68.

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.

9.2.2 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:

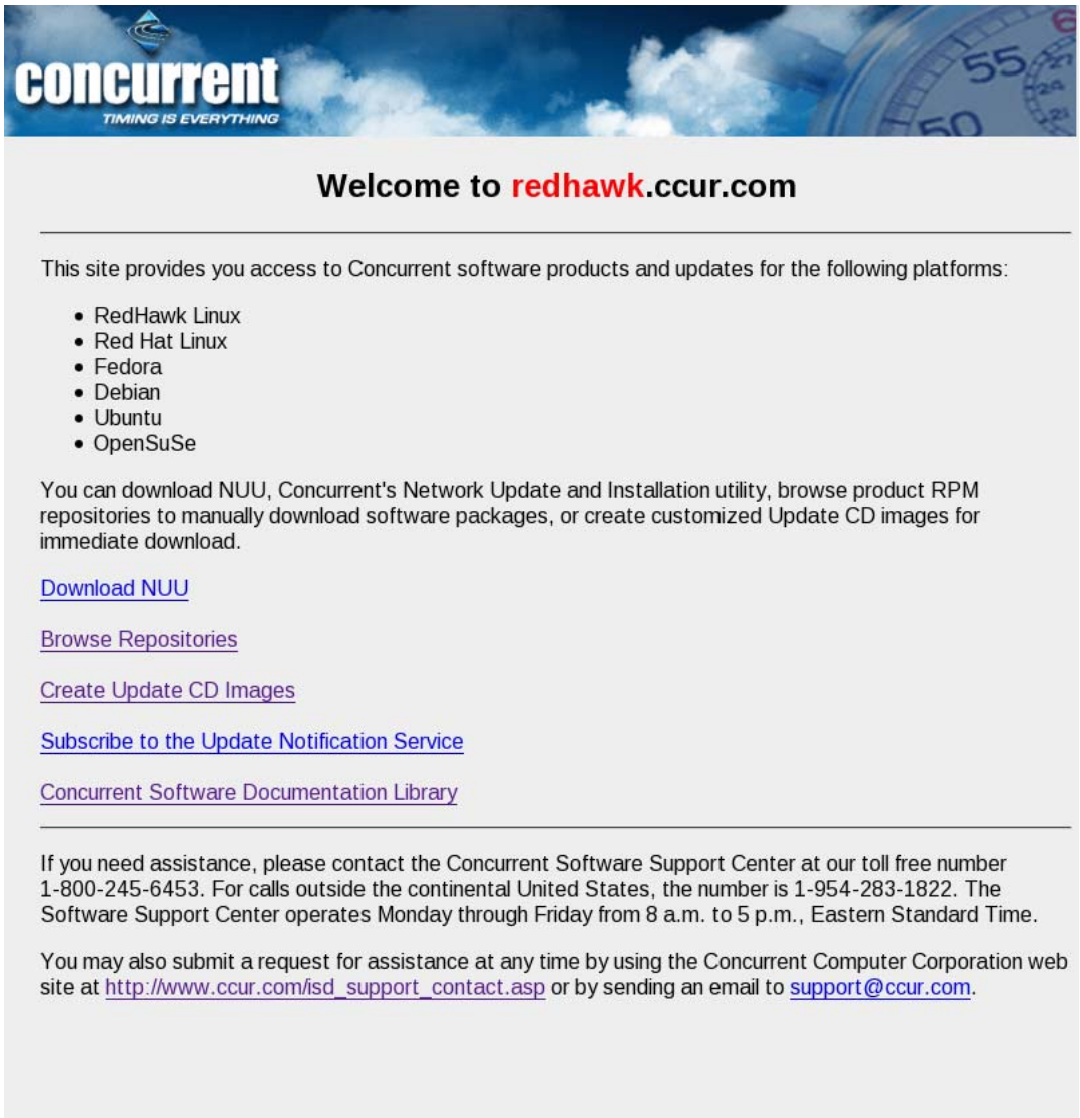

Clicking on the **Browse Repositories** link takes you to an authentication page.

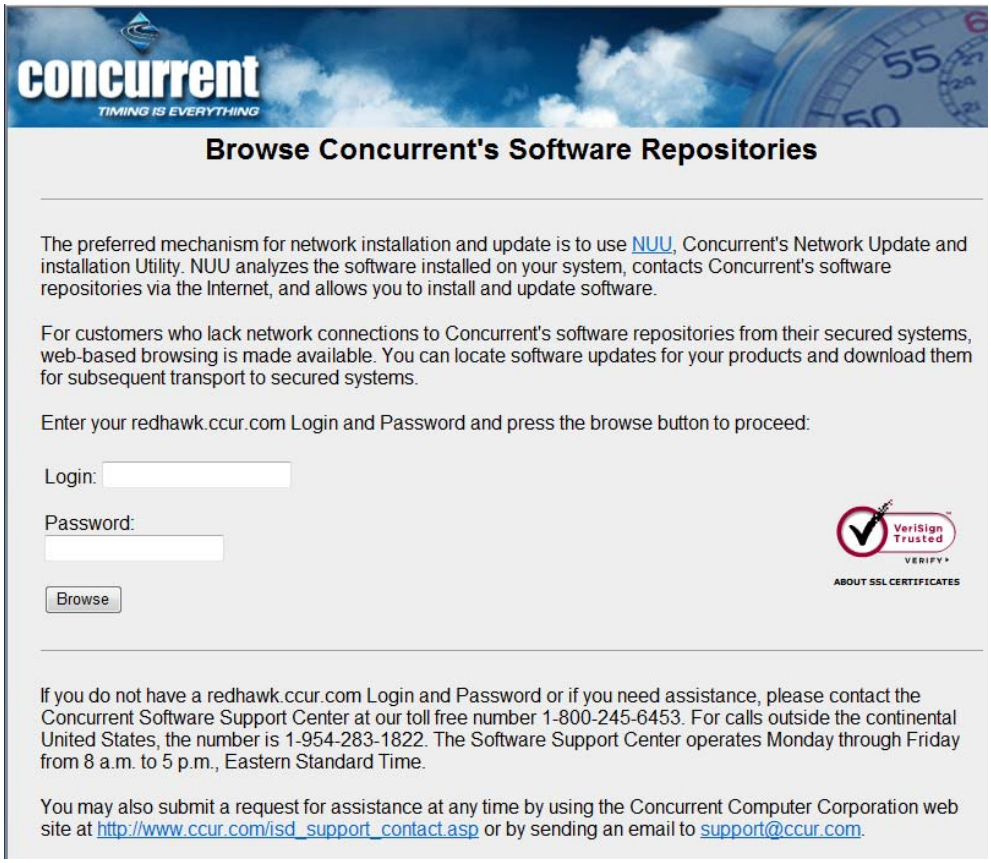

Enter your redhawk.ccur.com Login ID and Password and click the Browse button.

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

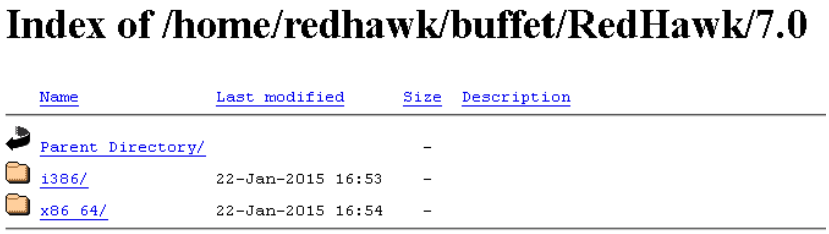

To locate the latest RPMs in the repository, click on the Last modified column header to sort by date. You may need to click twice to set the sort order to newest-to-oldest.

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

To install newly downloaded packages, follow these steps:

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

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

The time it takes to install will vary depending on the number of updates being installed.

3. When complete, exit single-user mode (Ctrl-D).

NOTE

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

9.2.3 Customized Update 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.

9.3 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**.

You can also view the latest RedHawk FAQ at **http://redhawk.ccur.com/docs/root/1redhawk/FAQ/RedHawk-FAQ.pdf**.