

REDHAWK Linux®
Version 8.2
Release Notes

November 2020

0898003-8.2



Disclaimer

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

License

Duplication of this manual without the written consent of Concurrent Real-Time, Inc. is prohibited. Any copy of this manual reproduced with permission must include the Concurrent Real-Time, Inc. copyright notice.

Trademark Acknowledgments

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

Contents

1.0	Introduction	1
1.1	Product Description	1
1.2	Product Media	2
1.3	Version Numbers	3
1.4	System Updates	4
1.5	Syntax Notation	4
2.0	Documentation	5
2.1	RedHawk Linux Documentation	5
2.2	Vendor Documentation	6
3.0	Prerequisites	7
3.1	Software	7
3.2	Hardware	7
4.0	Changes in this Release	8
4.1	Enhancements in RedHawk Version 8.2	8
4.1.1	Based on CentOS Version 8.2 and Ubuntu 20.04	8
4.1.2	Based on kernel.org 5.4.66-rt38	8
4.1.3	Upgrades from Version 8.0	8
4.1.4	New RedHawk System Call IDs Enforced	8
4.1.5	Changes in Managed Interrupts	8
4.1.6	New RCIM Driver	9
4.1.7	Support for encryption of system memory (AMD)	9
4.1.8	New RedHawk Kernel Preservation Service	9
4.1.9	<code>b1scfg</code> grub Utility enhancements	9
4.1.10	Secure Boot support added to custom kernel builds	9
4.1.11	Real-time virtualization enhancements	9
4.1.12	New tracepoints added to xtrace	10
4.1.13	Updated NVIDIA® Driver	10
4.1.14	Updated CUDA SDK 11.02	10
4.1.15	Updated OpenOnload® (renamed Onload)	10
4.2	Enhancements in RedHawk Version 8.0	10
4.2.1	Based on CentOS 8.0	10
4.2.2	Based on kernel.org 4.14	11
4.2.3	UEFI Secure Boot Support	11
4.2.4	Updated RCIM driver	11
4.2.5	Updated RCIM GPS	11
4.2.6	<code>b1scfg</code> Grub Utility	11
4.2.7	Updated NVIDIA® Driver	11
4.2.8	Updated CUDA SDK 10.1	11
4.2.9	Updated Network OpenOnload®	12
4.2.10	Ubuntu Supported	12
4.3	Modifications in Version 8.0	12
4.3.1	Upgrades to version 8.0 not supported	12
4.3.2	Tracing of PERFMON events deprecated	12
4.3.3	RedHawk system calls have new IDs	12
4.4	Changes in CentOS 8.0	13
5.0	Installation Procedures	14
5.1	Installing CentOS Software	15
5.1.1	Installing CentOS	15
5.1.2	Installing CentOS Updates	31
5.2	Installing RedHawk Linux	34

5.3	Installing the RCIM	38
5.3.1	Quick Hardware Installation Checklist	38
5.3.2	Verifying RCIM Operation	38
5.3.3	Installing NTP Updates for GPS Module	39
5.4	Installing the Frequency-Based Scheduler (FBS)	40
5.5	Installing Additional RedHawk Products	40
5.6	Installing RedHawk Updates from Disc	40
5.7	Installing RedHawk Updates from Web Site	41
6.0	Upgrade Procedures	42
6.1	Pre-Upgrade Procedure	42
6.2	RedHawk Pre-Upgrade Steps	43
6.3	Upgrade RedHawk to RedHawk 8.2	44
7.0	NVIDIA Graphics Configuration	46
7.1	Supported NVIDIA GPUs	46
7.2	Multiple NVIDIA Video Card Considerations	46
7.3	Display Configurations	47
7.3.1	Single	47
7.3.2	Xinerama	47
7.3.3	MultiDesktop	47
7.3.4	Twinview	48
7.3.5	Twinview-Xinerama	48
7.3.6	Twinview-MultiDesktop	49
7.4	Using nvidia-xconfig	49
7.5	Using nvidia-settings	49
7.6	Modified xorg.conf file generated by nvidia-xconfig	57
7.7	Troubleshooting Problems	58
8.0	Additional Information	60
8.1	Installation/Configuration Issues	60
8.1.1	File System Backup Recommendations	60
8.1.2	File System Swap Size Guidelines	61
8.1.3	RCIM Connection Modes	61
8.1.4	How to Use Linux Rescue	61
8.1.5	How to Install CentOS Remotely Over VNC	64
8.1.6	If the Installation GUI Fails to Start	66
8.2	Considerations	67
8.2.1	Compiler Requirements	67
8.2.2	Hyper-thread Activation	68
9.0	Known Issues	69
10.0	Software Updates and Support	72
10.1	Direct Software Support	72
10.2	Software Updates	72
10.2.1	Updating via NUU	72
10.2.2	Installing Manually Downloaded RPMs	73
10.2.3	Customized Update Discs	75
10.3	Documentation Updates	75

Appendix A

Ubuntu Support

Overview	A-1
Installing Ubuntu	A-1
Updating Ubuntu Software	A-5
Installing RedHawk Linux	A-7

Updating RedHawk Software	A-9
Ubuntu-Specific Notes	A-10
Building and Installing Custom Kernels	A-11
Installing Ubuntu in Safe graphics mode	A-12

Appendix B

ARM64 Support

Overview	B-1
Installation	B-1

Appendix C

UEFI Secure Boot

Overview	C-1
Configuration	C-1

1.0. Introduction

This document provides release information and installation instructions for Concurrent Real-Time, Inc.'s RedHawk™ Linux® operating system Version 8.2.

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 5.4 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 8.2. However, note that RedHawk can also be installed on *any* CentOS-compatible version 8.2 distribution. In addition, RedHawk 8.2 now supports the popular Ubuntu distribution; see Appendix A, *Ubuntu Support*, for more information.

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 Real-Time developed features derived from the real-time UNIX® implementations that Concurrent Real-Time has supported in its over 45 years of experience developing real-time operating systems.

RedHawk is included with each Concurrent Real-Time 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 8.2 as it is only available in a 64-bit version.

Support for SMP 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 5.4 series of kernels. Linux on the Intel® x86 and ARMv8-A architectures has defined a defacto binary standard which allows shrink-wrapped Linux applications that are designed to run on the Intel x86 and ARMv8-A platforms to run on Concurrent Real-Time's iHawk platforms.

NightStar RT, RedHawk Architect and RedHawk KVM-RT are optional products available to run on RedHawk Linux.

NightStar RT is Concurrent Real-Time'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.

RedHawk Architect is a powerful tool with an easy-to-use GUI that lets a developer choose the Linux and application modules to be included in target images. Target images can be scaled from complete workstations to dedicated servers and even down to small embedded applications. RedHawk Architect provides deployment tools for installing onto hard drives, flash memories, DVD or even USB flash.

RedHawk KVM-RT is a Real-Time Hypervisor solution that utilizes QEMU/KVM and RedHawk real-time features to extend RedHawk's real-time determinism to guest RedHawk virtual machines.

1.2 Product Media

The tables below list the optical media discs that are delivered with RedHawk Linux Version 8.2 for each Concurrent Real-Time 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]
	iHawk ARM-based systems	[HAxxx]
	ImaGen Intel-based systems	[HQ0G1]
	ImaGen AMD-based systems	[HR0G1]
Operating System Discs:		
CentOS 8.2 (64-bit)	(1 DVD)	[WA-CENT, WA-CENT-ARM64 Version 8.2]
CentOS 8.2 Source (64-bit)	(2 DVDs)	[WA-CENT, WA-CENT-ARM64 Version 8.2]
CentOS 8.2 Updates (64-bit)	(1 DVD)	[LXUPDATE Version 8.2]
RedHawk Linux OS (64-bit)	(1 DVD)	[WA-EM64T-AMD64, WA-ARM64 Version 8.2]
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 8.2]
NightStar RT for RedHawk		[WU1220-LDV, WU1220-LDVF Version <i>x.x</i> (<i>x.x</i> =latest version)]
RedHawk Architect (64-bit)		[W-RHLE-E64] Version 8.2
RedHawk KVM-RT (64-bit)		[WU1602-KVM-RT] Version 1.1

1.3 Version Numbers

All RedHawk Linux version numbers take the following form:

$$major.minor[.update]$$

where:

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

For example, 8.0 is the initial release of major version 8, and 8.2 is a subsequent release to version 8.0. Both 8.0 and 8.2 constitute complete product releases and do not require a previous release to be currently installed on the system; both are available on optical media from Software Distribution.

Individual kernel updates available via the RedHawk Updates website (see the "System Updates" section below) are not a complete product and can only be installed as an update to the matching *major.minor* release. For example; a 8.2.1 kernel update can only be installed on a system currently running RedHawk Linux version 8.2.

1.4 System Updates

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

NOTE

Concurrent Real-Time does not recommend arbitrarily 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 Real-Time may destabilize the system and is not recommended. Security updates from outside sources may be installed freely if desired.

Instructions for downloading updates from the website are provided in the section "Software Updates" on page 72.

1.5 Syntax Notation

The following notation is used throughout this document:

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

2.0. Documentation

2.1 RedHawk Linux Documentation

Table 2-1 lists the documentation provided with RedHawk Version 8.2 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 Real-Time’s documentation web site at <http://redhawk.concurrent-rt.com/docs>

Table 2-1 RedHawk Linux Documentation

RedHawk Linux Operating System Documentation	Document Number
<i>RedHawk Linux Version 8.2 Release Notes</i>	0898003-8.2
<i>RedHawk Linux User’s Guide</i>	0898004-780
<i>Real-Time Clock & Interrupt Module (RCIM) User’s Guide</i>	0898007-620
<i>RedHawk Linux FAQ</i>	N/A
Optional RedHawk Product Guides	
<i>RedHawk Linux Frequency-Based Scheduler (FBS) User’s Guide</i>	0898005-350
Optional RedHawk Product Documentation Sets	
<i>RedHawk Architect Release Notes</i>	0898600-8.2
<i>RedHawk Architect User’s Guide</i>	0898601-8.2
<i>KVM-RT Release Notes</i>	0898603-1.1
<i>KVM-RT User’s Guide</i>	0898604-1.1

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

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

Printed copies can be ordered by contacting the Concurrent Real-Time Software Support Center. Refer to “Direct Software Support” on page 72 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 8.2 with updates
- RedHawk Linux operating system (with optional updates)
- NightStar RT for RedHawk tools (optional)
- RedHawk Linux Frequency-Based Scheduler (FBS) (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 8.2, real-time extensions and features are optionally available to 32-bit compatibility-mode applications running under a 64-bit operating system (i.e. x86_64).

3.2 Hardware

- Any Concurrent Real-Time iHawk or ImaGen system
- Real-Time Clock and Interrupt Module (RCIM)
- A dual-layer optical media reader is required to install the system

4.0. Changes in this Release

This section describes enhancements and other changes introduced in RedHawk Version 8.2 (and any subsequent kernel updates if applicable).

4.1 Enhancements in RedHawk Version 8.2

4.1.1 Based on CentOS Version 8.2 and Ubuntu 20.04

RedHawk 8.2 can be used with CentOS Community ENTerprise Operating System 8.2 as well as Ubuntu 20.04 as its base.

4.1.2 Based on kernel.org 5.4.66-rt38

RedHawk 8.2 uses the `kernel.org` Linux 5.4.66-rt38 kernel as the basis for its core kernel. For more information examine the kernel change summaries at the <http://kernelnewbies.org> website.

4.1.3 Upgrades from Version 8.0

Systems installed with RedHawk 8.0, can be upgraded to RedHawk 8.2 without having to do a full reinstall of the system. See "Upgrade Procedures" on page 42 for more information.

4.1.4 New RedHawk System Call IDs Enforced

In this release, the `old_redhawk_syscall` kernel boot option has been removed and the new system call IDs must be used. The new IDs are described in 4.3.3, "RedHawk system calls have new IDs" on page 12.

4.1.5 Changes in Managed Interrupts

- Managed Interrupts, by default, are now migratable. The default behavior of MSI and MSIX interrupts has been changed to allow these interrupts to be moved from their assigned CPU. This allows the user to migrate all interrupts away from a shielded CPU. The old behavior can be retained by passing the kernel boot parameter `enable_managed_irqs`.
- New per-IRQ /proc file for Managed Interrupts. A new file under each IRQ's `/proc/irq/[irq-number]` directory can be used to enable or disable managed interrupts on a per IRQ basis. The file is named `managed`. Writing a 0 to this file will allow that particular IRQ to be migrated while writing a 1 will prevent it from being migrated.
- CPUs can be assigned to handle Managed Interrupts. The `msi_affinity_mask` kernel boot parameter can be set to allocate some CPUs to handle managed interrupts. At boot time, all the MSI and MSIX interrupts will be assigned to the CPUs in the mask, leaving the other CPUs available for shielding. It takes a list of CPUs. For example, `msi_affinity_mask=0-5,8,16`.

4.1.6 New RCIM Driver

New RCIM driver now fully supports the latest hardware revision of RCIM III.

4.1.7 Support for encryption of system memory (AMD)

For the first time RedHawk 8.2 has enabled support for the encryption of system memory. It requires an AMD processor that supports Secure Memory Encryption and the corresponding setting must be enabled in the BIOS. Note, that with this option, you will be unable to perform I/O mappings from user-space. If this change presents problems, this feature can be disabled in the BIOS.

4.1.8 New RedHawk Kernel Preservation Service

New `ccur-default-kernel` service tries to ensure that the kernel chosen as the default is not overwritten by kernel updates. When the system is installed, the RedHawk trace kernel is the default. After that it can be changed via the `-d` option of `blscfg(1)` in CentOS systems and `ccur-grub2(1)` in Ubuntu. The `ccur-default-kernel` service remembers the last default kernel setting and restores it before a reboot. Note that if you want to boot a custom kernel you must set it as the default kernel before rebooting or select it from the kernel menu during the boot.

4.1.9 `blscfg` grub Utility enhancements

- A new option, `-V`, returns 0 if grub is using a BLS configuration, 1 otherwise.
- For kernel options that have a setting, the `--kopt-del=` option no longer requires you to specify the setting, just the option name. For example, the command `blscfg --kopt-del=crashkernel=` will remove the kernel boot option “`crashkernel=256@64M`”.

4.1.10 Secure Boot support added to custom kernel builds

The custom kernel building tool, `ccur-config(8)`, will now preserve any existing kernel module signing key pairs. A new `ccur-config` option, `--nopreserve`, is now available to force key removal.

4.1.11 Real-time virtualization enhancements

The following features were added in this release.

- Two new boot options to allow non-IOMMU capable drivers to work when IOMMU is enabled.
 1. Kernel boot option is `intel_iommu=on,plx_off`.
The `plx_off` setting prevents the Intel-specific IOMMU driver from assigning any device (i.e. RCIM, ADO cards), located behind a PLX bridge chip, from being placed in an IOMMU group and used for passthrough.
 2. Kernel boot option is `intel_iommu.exception_ids`.
This option exempts the listed devices from being placed in an IOMMU group. Requires that the kernel boot parameter `intel_iommu` be set to `on` and that either the kernel boot parameter `iommu` is set to `pt` or that kernel configuration parameter

IOMMU_DEFAULT_PASSTHROUGH is enabled. It takes a comma separated list of devices specified by *vendor_id:vendor_device_id*. The values are specified in hex with no 0x prefix. For example, `intel_iommu.exception_ids=10de:1eb1,10de:10f8`

- A new kernel option `vfio-pci.addr=bus:slot:func[,...]`
This option allows a user to assign one or more devices to the Virtual Function I/O (VFIO) driver. The VFIO driver loads first and claims the specified device(s). This option takes a comma separated list in which devices are specified by bus, slot and function as displayed by `lspci (8)`.
- Support for AMD's Secure Encrypted Virtualization that secures a VM's RAM from the host and other VMs.
- Real-time support for KVM clock time source in VMs running RedHawk.
- The ability to assign specific CPUs to handle managed interrupts (see bullet item in section 4.1.5).

4.1.12 New tracepoints added to xtrace

Two new trace points have been added, `TASK_WORK` and `TASK_WORK_DONE`. These display when a system call exit has been given extra work to do. That work is in the form of a collection of functions to be called. xtrace displays the name of each such function executed and how long each took to execute.

4.1.13 Updated NVIDIA® Driver

RedHawk 8.2 includes the NVIDIA Linux Display Driver version 450.57. Legacy driver support for older NVIDIA GPUs is available for download from the NVIDIA website. Refer to "NVIDIA Graphics Configuration" on page 46 for details.

4.1.14 Updated CUDA SDK 11.02

RedHawk 8.2 includes pre-installed NVIDIA CUDA 11.02 libraries and header files. See the many resources at <http://www.nvidia.com> for more information.

4.1.15 Updated OpenOnload® (renamed Onload)

RedHawk 8.2 comes with Onload version 7.1.0.265. Onload is a high performance network stack from Solarflare.

4.2 Enhancements in RedHawk Version 8.0

4.2.1 Based on CentOS 8.0

RedHawk 8.0 uses CentOS Community ENTERprise Operating System 8.0 as its base. CentOS 8.0 is the world-wide standard enterprise-class Linux distribution that provides a rock-solid foundation. There have many changes in the CentOS 8.0 release which are highlighted in section 4.4 below.

4.2.2 Based on kernel.org 4.14

RedHawk 8.0 uses the `kernel.org` Linux 4.14 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.

4.2.3 UEFI Secure Boot Support

Secure Boot technology ensures that the system firmware checks whether the system boot loader is signed with an authorized cryptographic public key. A chain of trust is then established from the firmware to the signed drivers and kernel modules. See Appendix C for information on configuring UEFI Secure Boot.

4.2.4 Updated RCIM driver

The RCIM's emulator and uart drivers were merged into the RCIM driver freeing up two major numbers. The RCIM device nodes are no longer static but are created and removed when the RCIM driver is loaded and unloaded. Also note that when a virtual machine is configured with an RCIM, the device nodes on the host are removed.

4.2.5 Updated RCIM GPS

Both RCIM II and RCIM III have been modified to support the Trimble SQ GPS receiver. Note that while the `ntpd` service has been removed and replaced by `chronyd` in the CentOS 8.0 release, the RCIM GPS optional product will continue to provide `ccur-ntpd`. You must stop and disable `chronyd` and enable and start `ntpd` using the `systemctl` command in order to use `ccur-ntpd`.

4.2.6 `blscfg` Grub Utility

`blscfg` is a new tool for viewing and managing Boot Loader Specification (BLS) configuration files. It can also be used to set one of the boot entries as the default for Grub to boot. `blscfg` provides similar functionality as the existing `ccur-grub2` command except that `blscfg` displays and manages BLS entries configured in `/boot/loader/entries`, while `ccur-grub2` displays and manages menu entries configured directly within the grub2 configuration file.

`blscfg` should be used when the grub2 configuration uses BLS; otherwise `ccur-grub2` should be used. The CentOS 8 release use BLS by default. Older CentOS installations and Ubuntu installations do not use BLS by default. See `blscfg(1)` for more information.

4.2.7 Updated NVIDIA® Driver

RedHawk 8.0 includes the NVIDIA Linux Display Driver version 430.34. Legacy driver support for older NVIDIA GPUs is available for download from the NVIDIA website. Refer to "NVIDIA Graphics Configuration" on page 46 for details.

4.2.8 Updated CUDA SDK 10.1

RedHawk 8.0 includes pre-installed NVIDIA CUDA 10.1 libraries and header files. New features in CUDA 10.1 include Turing architecture support (including RTCore support), asynchronous task graphs, NSight restructuring and new `libcublasLt`. See the many resources at <http://www.nvidia.com> for more information.

4.2.9 Updated Network OpenOnload®

RedHawk 8.0 comes with OpenOnload version 201811-u1. OpenOnload is a high performance network stack from Solarflare.

4.2.10 Ubuntu Supported

In addition to CentOS and CentOS-compatible distributions, RedHawk 8.0 is available for the Ubuntu 18.04 distribution. See Appendix A, *Ubuntu Support*, for more information.

4.3 Modifications in Version 8.0

Users should be aware of modifications that are being introduced in the RedHawk 8.0 release; these changes are described in the following subsections.

4.3.1 Upgrades to version 8.0 not supported

RedHawk Version 8.0 is a major release. Upgrades from previous releases to version 8.0 is not supported.

4.3.2 Tracing of PERFMON events deprecated

Xtrace PERFMON events will no longer be traced by the RedHawk trace and debug kernels in future releases.

4.3.3 RedHawk system calls have new IDs

Previously, RedHawk system IDs started with ID number 420. Due to conflicts, the IDs have been changed to start at 1900. See new IDs in the table below. Programs using RedHawk system calls will be affected as follows:

- Statically linked binaries must be relinked or switch to dynamic.
- Programs calling syscalls by name (i.e. `syscalls(__NR_postwait)`) must be recompiled.
- Programs calling syscalls directly by number (i.e. `syscalls(422)`) must be updated with new ID and recompiled.

System call	New ID	Old ID
<code>fbsched</code>	1900	420
<code>sched_setscheduler_head</code>	1901	435
<code>resched_cntl</code>	1902	421
<code>postwait</code>	1903	422
<code>cpustat</code>	1904	424
<code>cpucntl</code>	1905	444
<code>procstat</code>	1906	452
<code>local_irq</code>	1907	434
<code>mlockall_pid</code>	1908	437
<code>munlockall_pid</code>	1909	454
<code>watchpoint_ctl</code>	1910	510

4.4 Changes in CentOS 8.0

Many improvements and additions have been implemented in the CentOS 8.0 release and also a lot of functionality has been removed or deprecated. Following are some of the most notable ones. For a more complete description see the CentOS release notes at <https://wiki.centos.org/Manuals/ReleaseNotes/>.

- For security reasons, static libraries are no longer provided.
- Default python version is now 3.6 invoked as **python3**. The `/usr/bin/python` link has been removed. Python 2.7 is still available but not installed.
- **GCC** version jumps from 4.8.5 to 8.2.
- **Qt4** is no longer provided, only **Qt5** is available.
- **Yum** package manager is replaced by **DNF**.
- The Virtual Machine Manager (`virt-manager`) application has been deprecated and the web console **Cockpit** is intended as its replacement.
- Network Time protocol **ntpd** is replaced by **chrony**.
- GNOME has a new default graphical interface (GNOME Shell) and display server (Wayland). Other options can be configured by the user. KDE is no longer provided.
- **ifconfig** (`net-tools`) is deprecated, replaced by **ip** (`iproute`).
- NFS over UDP is no longer supported.
- **nftables** replace **iptables**.

5.0. Installation Procedures

Use the upgrade procedures provided in section 6.0. if you are upgrading from Version 8.0 to Version 8.2.

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

NOTE: Upgrades from RedHawk 7.x to RedHawk 8.x are *not* supported.

NOTE: Resources to assist users supplying their own Red Hat software are available in the Concurrent Real-Time Software Documentation Library <http://redhawk.concurrent-rt.com/docs/>. Look under the corresponding RedHawk Linux Release for the link labeled : Installing on RHEL.

Review the following points before installing the software:

- Concurrent Real-Time 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.

Before you start:

- **Main Memory Size** — You need to know the system’s main memory size in order to properly set up the file system swap size. If you are not sure of the main memory size, access the system’s BIOS to ascertain this information.
- **Network Addresses** — If you do not plan to configure networking via DHCP, you will need to have the following information for your system before beginning the installation:
 - IP address
 - gateway address
 - netmask address
 - primary and secondary DNS addresses
- **Firmware Settings** — Firmware settings should be reviewed before installation and whenever problems involving non-working devices are encountered. Exercise caution before changing default settings. Concurrent Real-Time 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 firmware 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 firmware 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 8.2 version level you can skip directly to "Installing CentOS Updates" on page 31.

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 highlight Install CentOS 8 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 **8.1.6** on page 66 for some boot options that may allow 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.

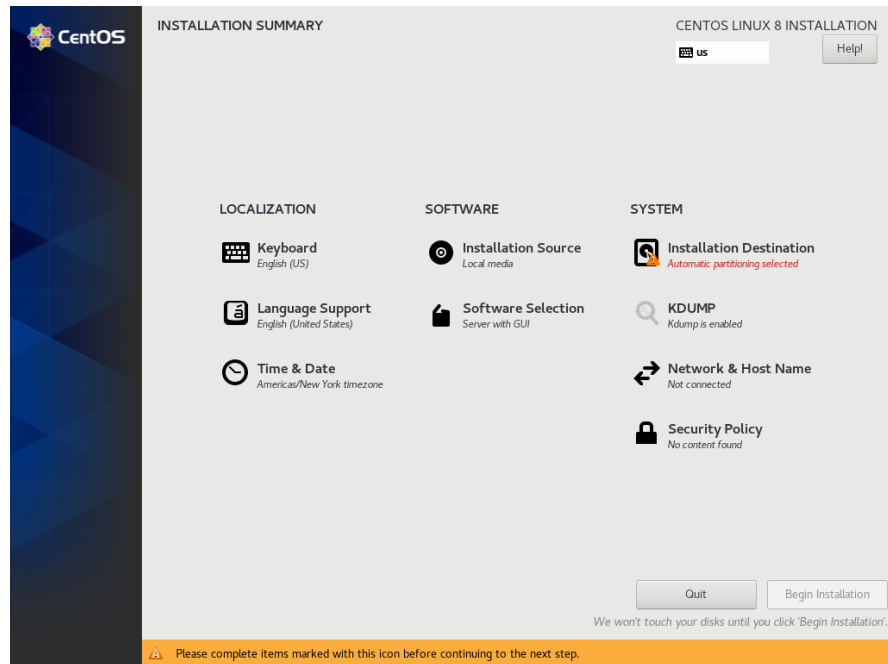
If you wish to use VNC for the installation instead of the GUI or text modes, refer to section **8.1.5** on page 64 for full instructions.

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 Real-Time Support (see "Direct Software Support" on page 72).

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 welcome screen. Select your desired language and press **Continue** to proceed.
6. 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.

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

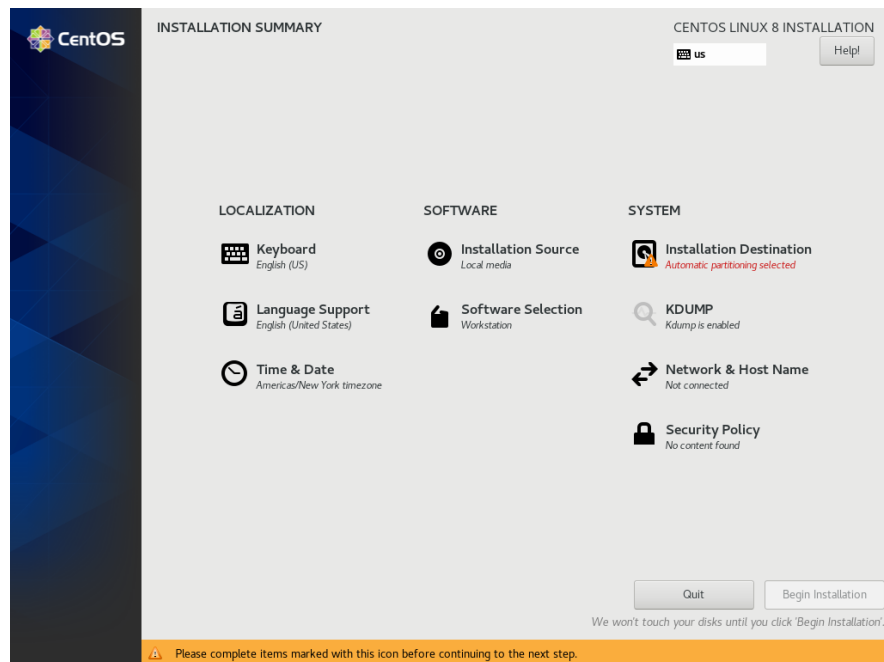


Choose **Workstation** and select all of the **Add-Ons** categories displayed on the right side of the screen.

NOTE

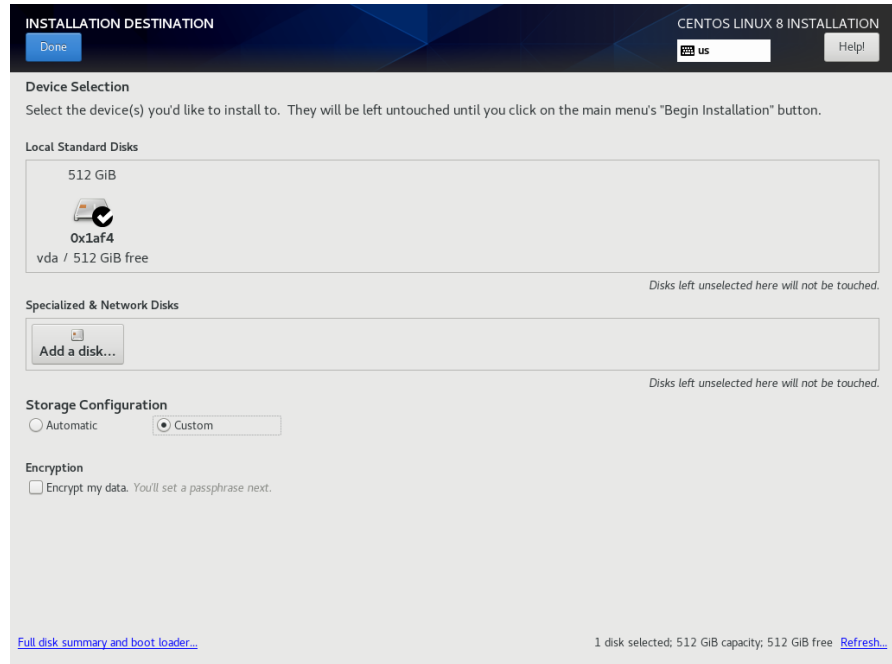
Click on the top add-on (Add-ons for Selected Environment) and then repeat pressing the down arrow on the keyboard, followed by the space bar, until all add-ons are selected. Note that the list will continue to scroll across multiple pages of add-ons and will stop scrolling only after reaching the last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



8. 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.

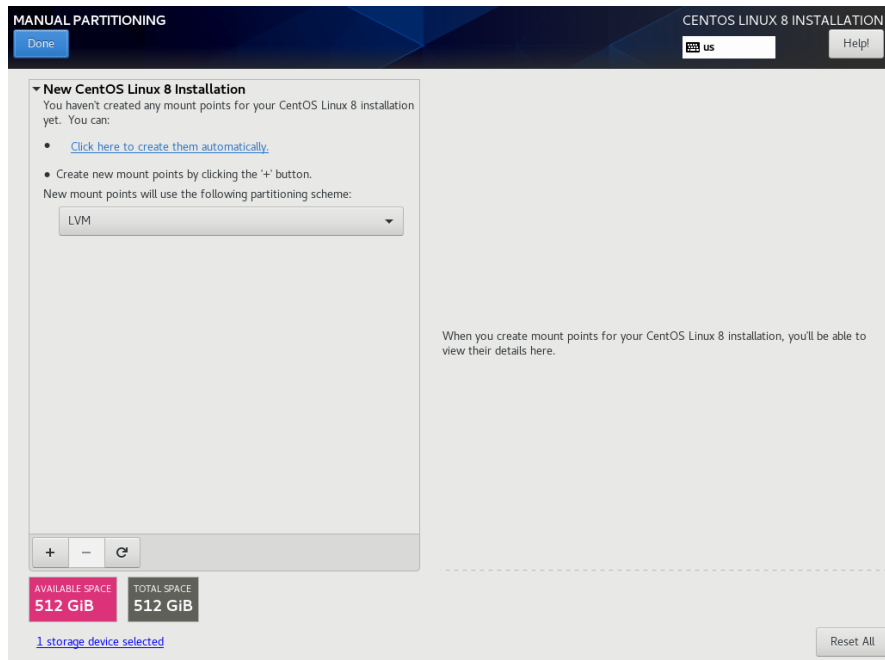


NOTE

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.

9. On this page, select **Custom** under **Storage Configuration** 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 / partition, clicking the “-” (minus) icon, checking the Delete all other filesystems checkbox, and finally pressing the Delete It button to delete an entire existing partitioning scheme, if necessary.



10. 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.

11. On the Manual Partitioning page, select the `/home` partition, change the **Desired Capacity** to 10 GiB, and then press the **Update Settings** button.

NOTE

For smaller hard drives with 128 GB capacity or less, the separate `/home` partition should be eliminated by selecting it and pressing the `-` (minus) key at the bottom of the partition list. Then, skip to the next step and ignore `/home` in all of the examples.

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.

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

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

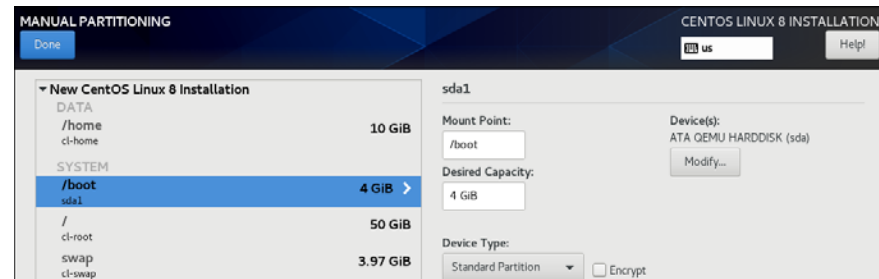


NOTE

For smaller hard drives with 128 GB capacity or less you can leave the default for the / partition and skip to the next step to create the boot partition.

- Next, select the /boot partition, change the Desired Capacity from 1024 MiB to 4096 MiB, and then press the Update Settings button.

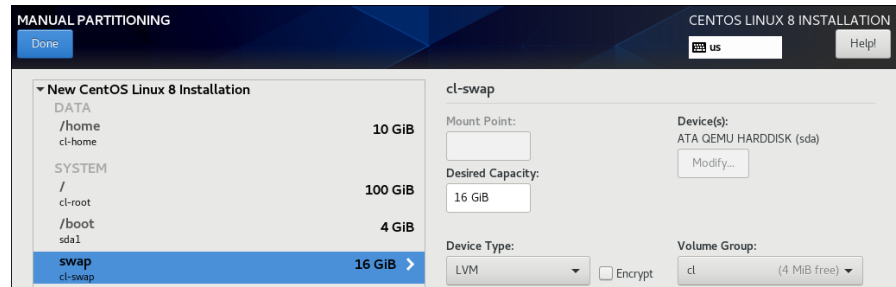
The following snippet shows an example of what should now be displayed:



- 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
0 - 3	2048 MiB
4 - 15	4096 MiB
16 - 63	8192 MiB
64 - 255	16 GiB
256 - 511	32 GiB
512 - 1024	64 GiB

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




15. 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.


NOTE

For smaller hard drives with 128 GB capacity or less, the `/home` partition should have been deleted above, so you can perform this step with the `/` partition instead; thus, the `/` partition size will be set to the maximum size available using all remaining disk space.

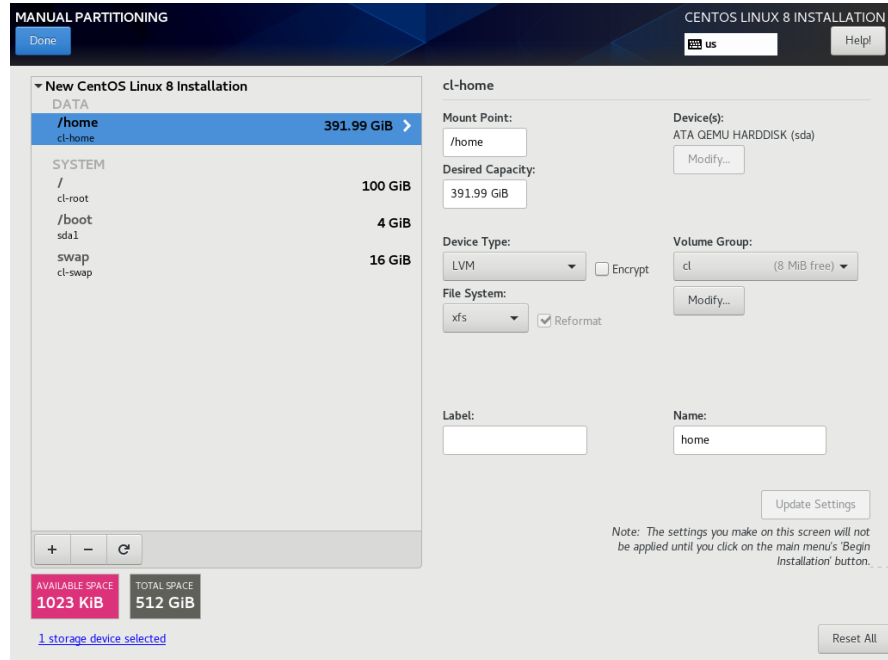
The following snippet shows an example of what should now be displayed:



Alternatively, here is an example of what should now be displayed for systems with smaller hard drives with 128 GB capacity or less that should not have a `/home` partition:




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



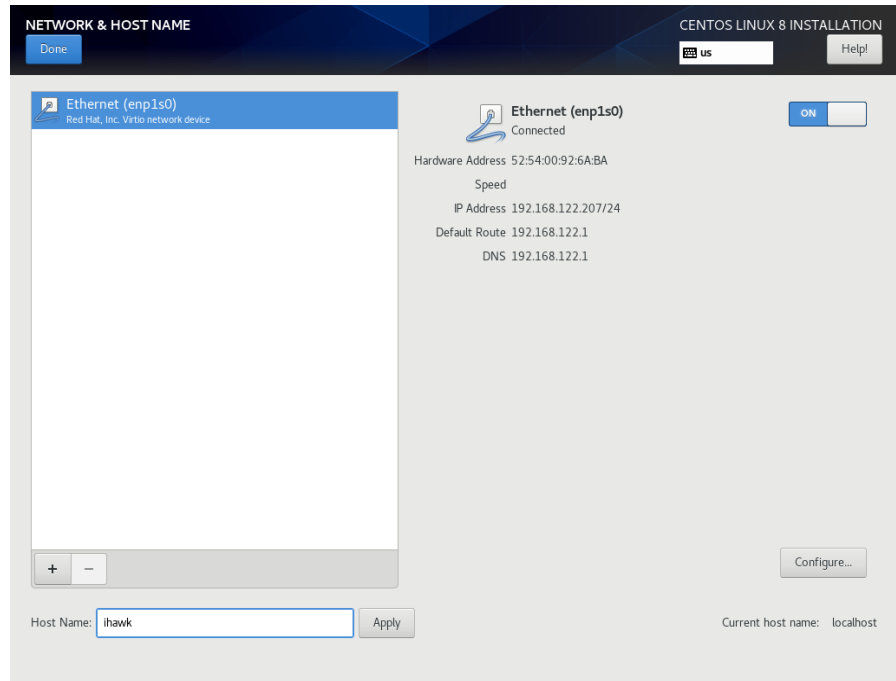
Finally, press the Done button to use this layout for the system's partitioning.

17. 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.


18. Back on the Installation Summary page, click on **Network & Hostname** to set the system's hostname and networking configuration. Type in the hostname on the box at the bottom of the page and click on the **Apply** button next to it. For each network connection device shown on the left, toggle the connection button to the right to the **ON** position to connect to that network automatically. An example page follows:



Ensure the connected device is set to **ON** and press **Done** when you have finished configuring all of the networking details.

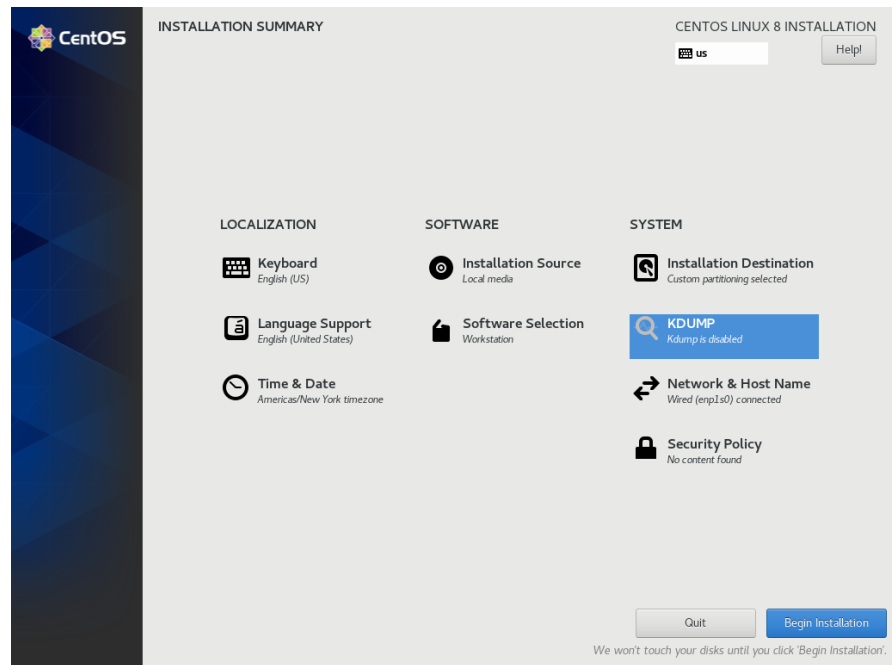
19. Back on the Installation Summary page, click on **Enable kdump** to deconfigure the default kernel crash dump settings.

The following screen shows the default kernel crash disabled:




Uncheck the Enable kdump check box and then press Done.

20. The Installation Summary page should now appear as shown in the following example screen:




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

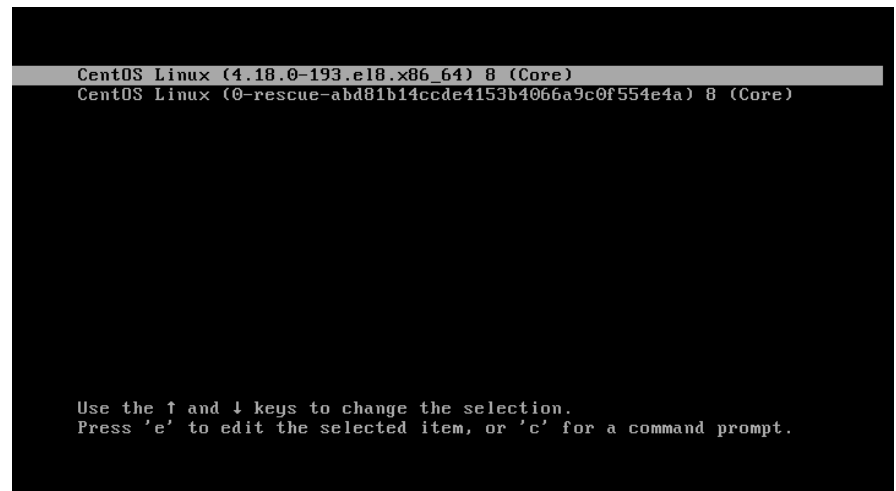
During the system installation you must set the system's root password and also create a non-root user.

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.

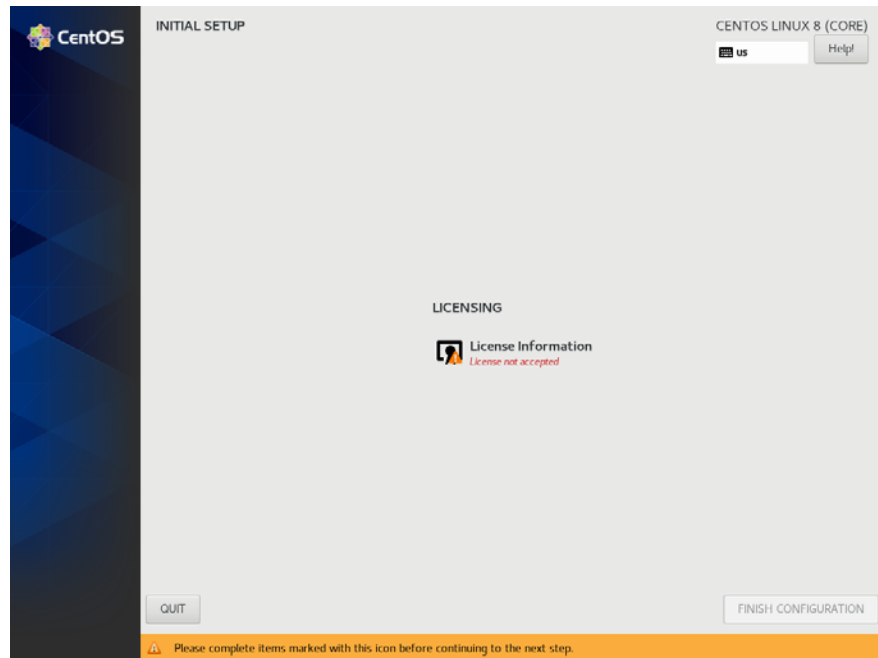
22. You should see the following grub screen displayed after the system restarts.



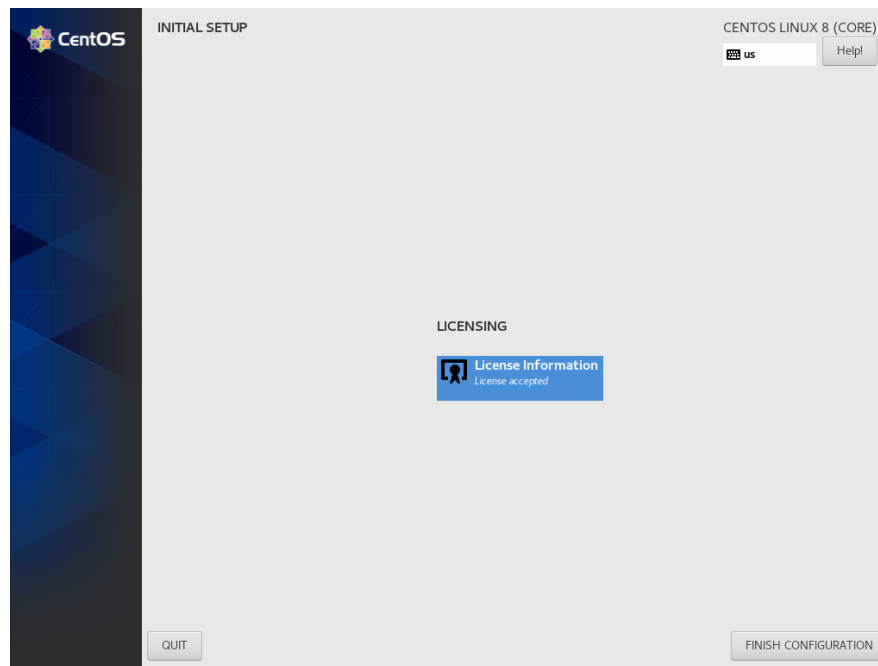
Press the **Enter** key or allow the grub screen to timeout to boot the newly installed system.

23. Once the system has almost fully booted, the boot process will pause and require you to accept the CentOS license agreement, as shown in the

following screen:



Click on the License Information icon. On the resulting screen check the I accept the license agreement checkbox at the bottom of the screen and then press the Done button at the top left. The following screen should now be displayed:



Finally, press the Finish Configuration button and the system will boot up to a graphical login. Login as root.

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.

Next, to bring up a terminal window, click on the **Activities** menu at the top of the screen and then on the terminal icon that appears on the left side of the screen.

The 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 Updates Version 8.2. This contains the latest updates released by CentOS which have been validated by Concurrent Real-Time 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. At the system prompt, type **init 1**.
 - b. 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.

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

```
cd /path-to-actual-mount-point
./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/CentOS8.2-Updates type iso9660 (ro,noexec)

> mount -o remount,exec /media/CentOS8.2-Updates
```

- The installation script will first ask you to confirm that you want to proceed with updating CentOS packages. Type **Y** and press **Enter** to continue.

The installation script will then ask if you wish to install additional CentOS packages required by RedHawk. Concurrent Real-Time recommends that you simply press **Enter** to install the full set of Concurrent Real-Time-recommended package updates.

The CentOS operating system will be updated from packages in the Update DVD. No action is required of you during this step.

```
root@ihawk:/run/media/root/CentOS-8.2-Updates-x86_64
File Edit View Search Terminal Help
[root@ihawk CentOS-8.2-Updates-x86_64]#
[root@ihawk CentOS-8.2-Updates-x86_64]# ./install-updates

This script will apply updates for the CentOS Linux packages currently
installed on this system.

You will also have the option of installing additional CentOS Linux
packages to prepare for installing RedHawk Linux and optional Concurrent
Real-Time software.

Do you want to continue? [y/n] y

In addition to installing updates, would you like to also install additional
packages to prepare for installing RedHawk Linux and optional Concurrent
Real-Time software? [y/n] (yes) y

Installing updates:

Last metadata expiration check: 0:00:03 ago on Wed 16 Sep 2020 10:33:31 AM EDT.
Dependencies resolved.
=====
Package      Arch  Version                Repository              Size
=====
Installing:
kernel       x86_64 4.18.0-193.14.2.el8_2
```

NOTE

During the update, a failure to restore the java.security file can be safely ignored.

NOTE

If the update appears to stall for several minutes shortly after pressing the **Enter** key, it is very likely that the CentOS package update checking service has locked the RPM package database to check for any available CentOS package updates. Simply issue a **killall packagekitd** at a bash shell prompt to safely allow the installation to continue.

Once the update of the CentOS operating system has completed, packages required by RedHawk will now be installed. No action is required of you during this step.

```

root@ihawk:/run/media/root/CentOS-8.2-Updates-x86_64
File Edit View Search Terminal Help
Complete!
Determining additional packages to install.

Installing additional packages:

Last metadata expiration check: 0:04:22 ago on Wed 16 Sep 2020 10:33:31 AM EDT.
Package accountsservice-0.6.50-8.el8.x86_64 is already installed.
Package alsa-lib-1.2.1.2-3.el8.x86_64 is already installed.
Package atkmm-2.24.2-6.el8.x86_64 is already installed.
Package atk-2.28.1-1.el8.x86_64 is already installed.
Package audit-libs-3.0-0.17.20191104git1c2f876.el8.x86_64 is already installed.
Package bash-4.4.19-10.el8.x86_64 is already installed.
Package bc-1.07.1-5.el8.x86_64 is already installed.
Package binutils-2.30-73.el8.x86_64 is already installed.
Package cairo-gobject-1.15.12-3.el8.x86_64 is already installed.
Package cairomm-1.12.0-7.el8.x86_64 is already installed.
Package cairo-1.15.12-3.el8.x86_64 is already installed.
Package centos-logos-80.5-2.el8.x86_64 is already installed.
Package chkconfig-1.11-1.el8.x86_64 is already installed.
Package coreutils-8.30-7.el8.2.1.x86_64 is already installed.
Package cpio-2.12-8.el8.x86_64 is already installed.
Package cpp-8.3.1-5.el8.0.2.x86_64 is already installed.

```

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.

- Upon completion of the system update as shown in the image below, the system will be ready for the RedHawk Linux installation.

```

root@ihawk:/run/media/root/CentOS-8.2-Updates-x86_64
File Edit View Search Terminal Help
qt5-qtX11extras-5.12.5-1.el8.x86_64
recode-3.6-47.el8.x86_64
systemtap-sdt-devel-4.2-6.el8.x86_64
texlive-lib-7:20180414-14.el8.x86_64
tftp-server-5.2-24.el8.x86_64
ttmkfsdir-3.0.9-54.el8.x86_64
xinetd-2:2.3.15-24.el8.x86_64
xorg-x11-fonts-Type1-7.5-19.el8.noarch
xorg-x11-proto-devel-2018.4-1.el8.noarch
xz-devel-5.2.4-3.el8.x86_64
zlib-devel-1.2.11-13.el8.x86_64

Complete!

All updates have been applied.

Please contact Concurrent Real-Time technical support if you had any
problems during this installation (http://www.concurrent-rt.com/support
or 1-800-245-6453).

Note that you may rerun this script at any time. You can also mount
this DVD and use dnf(8) to access its contents. See the RedHawk Linux
Release Notes for more information.

[root@ihawk CentOS-8.2-Updates-x86_64]#

```

- 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/sr1**.

- Remove the disc from the optical media drive and store.
- Type “reboot” at the system prompt and verify that the new CentOS kernel boots.
- 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:

- Log in as root and take the system down to single-user mode:
 - At the system prompt, type **init 1**.
 - Enter the root password at the rescue prompt to access the shell.
- Locate the disc labeled “RedHawk Linux Version 8.2” appropriate to your system’s architecture and insert it into the optical media drive.
- 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 if necessary:

```
cd /path-to-actual-mount-point
./install-redhawk
```

```
root@ihawk:/run/media/root/RedHawk-8.2-x86_64
File Edit View Search Terminal Help
[root@ihawk RedHawk-8.2-x86_64]# ./install-redhawk
This process will install the RedHawk Linux core files
for RedHawk Linux 8.2 (Taal) on the current system (ihawk).
Do you want to install RedHawk Linux? [y/n] y
-----
Concurrent Real-Time End-User License Agreement (the EULA).
This license should be read before acceptance of its terms. If the terms
are not acceptable then you must reject the license and terminate the
installation process.
A copy of this license may be found in /usr/share/doc/ccur/RedHawk-EULA
after installation.
-----
View, Accept, or Reject the terms of this license? [v/a/r] a
Would you like to install 32-bit compatibility packages? [y/n] (y)
This system does not appear to have an NVIDIA graphics card.
We do not recommend installing RedHawk NVIDIA graphics packages.
Install RedHawk NVIDIA graphics packages anyway? [y/n] (n)
```

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. Two installation options are then presented. The first option is to install 32-bit compatibility packages. The default action is to install.

The second installation option is to install the RedHawk NVIDIA graphics packages. The recommended action is dependent on whether an NVIDIA graphics card is detected on the system. If detected, the default action is yes to install; otherwise the default is no, do not install. When opting to install the RedHawk NVIDIA packages, you are also then prompted with an option to install NVIDIA CUDA packages.

NOTE

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

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.

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

```

root@ihawk:/run/media/root/RedHawk-8.0-x86_64
File Edit View Search Terminal Help
Installing      : nuu-2.0.1.beta-1.rhel8.x86_64          1/3
Running scriptlet: ccur-RedHawk-8.0-repo-1-1.noarch     2/3
Installing      : ccur-RedHawk-8.0-repo-1-1.noarch     2/3
Running scriptlet: ccur-RedHawk-8.0-repo-1-1.noarch     2/3
Running scriptlet: ccur-RedHawk-8.0.n-repo-1-1.noarch  3/3
Installing      : ccur-RedHawk-8.0.n-repo-1-1.noarch  3/3
Running scriptlet: ccur-RedHawk-8.0.n-repo-1-1.noarch  3/3
Verifying       : ccur-RedHawk-8.0.n-repo-1-1.noarch  1/3
Verifying       : ccur-RedHawk-8.0-repo-1-1.noarch     2/3
Verifying       : nuu-2.0.1.beta-1.rhel8.x86_64        3/3

Installed:
  ccur-RedHawk-8.0.n-repo-1-1.noarch      ccur-RedHawk-8.0-repo-1-1.noarch
  nuu-2.0.1.beta-1.rhel8.x86_64

Complete!

The RedHawk Linux installation is complete.

Please contact Concurrent Real-Time technical support if you had any
problems during this installation (http://www.concurrent-rt.com/support
or 1-800-245-6453).

[root@ihawk RedHawk-8.0-x86_64]#

```

- When the installation completes, execute the following commands:

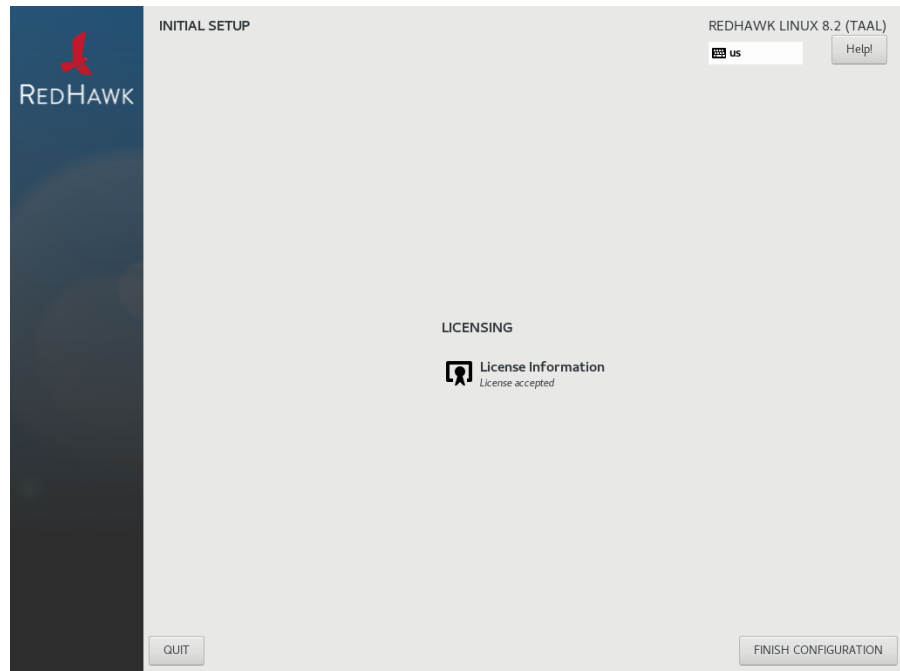
```

cd ..
eject

```

- Remove the disc from the drive and store.
- Reboot the system choosing a new RedHawk kernel and verify that the system starts booting properly.
- Once the system has almost fully booted, the boot process will again pause and allow you to accept or reject the RedHawk license agreement, as

shown in the following screen:



If you wish to accept the RedHawk license agreement simply press the **Finish Configuration** button and allow the system to continue to boot up to the graphical login screen.

NOTE

Desktops icons are not supported on the default Standard (Wayland display server) but are supported in some of the other options, such as the Classic (Wayland display server).

To switch to another display server, simply log off and, from the login screen, click on your username. At this point, click on the gear to the left of the Sign In button, and select a display server.

If you are running on a display server that supports desktop icons (see note above), you must right-click and select “Allow Launching” in order to launch the RedHawk icons.

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

5.3 Installing the RCIM

The following serves as an installation checklist for installing 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 61).
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 Real-Time, Inc. strongly recommends that you use an anti static 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.
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 11
This is a standalone (isolated) rcim.
```

Has IRQ 59 and major number 243
Board options: none

NOTE

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

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

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

You should see output similar to the following:

```
h/Not_Configured, sync/ptr, clock
pig0|out0, pig1|out1, pig2|out2, pig3|out3, pig4|out4, pig5|out5,
pig6|out6, pig7|out7, pig8|out8, pig9|out9, pig10|out10, pig11|out11,
none|di0/f, none|di1/f, none|di2/f, none|di3/f, none|di4/f, none|di5/f,
none|di6/f, none|di7/f, none|di8/f, none|di9/f, none|di10/f, none|di11/f,
eti0/f, eti1/f, eti2/f, eti3/f, eti4/f, eti5/f,
eti6/f, eti7/f, eti8/f, eti9/f, eti10/f, eti11/f,
```

5.3.3 Installing NTP Updates for GPS Module

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

- The **ccur-ntp** RPM is available via Concurrent Real-Time's Software Repository system and it is also included on the "RedHawk NTP Version 8.2" product media that was included with your system. Execute the following steps to install this RPM:
 - Disable the **chronyd** (8) service:


```
# systemctl stop chronyd
# systemctl disable chronyd
```
 - Mount the "RedHawk NTP Version 8.2" product media and **cd** to the mounted directory or download the **ccur-ntp** RPM from Concurrent Real-Time's web site as described in section 10.2.2 on page 73.
 - Install the RPM via the following command:


```
# rpm -Uvh ccur-ntp*.rpm
```
 - Enable the ntpd daemon


```
# systemctl enable ntpd
```
 - Reboot the system for changes to take effect.


```
# reboot
```

NOTE

If you see the error that follows, verify the **chronyd** service is not running and the **ntpd** service is running and that the corresponding **/dev/rcim_uart** device file is owned by **ntp** and is readable by all and writable by **ntp** (644 access permissions).
ntpq: read: Connection refused

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 the Frequency-Based Scheduler (FBS)

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 8.2 running, log in as root and take the system down to single-user mode.
2. Locate the disc labeled "RedHawk FBS Version 8.2" 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 /path-to-actual-mount-point  
./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 Additional RedHawk Products

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

5.6 Installing RedHawk Updates from Disc

Discs labeled "RedHawk Update Disc" 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 "RedHawk Update Disc" 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
```

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

```
cd /path-to-actual-mount-point  
./install-updates
```

When the installation is complete, the Main window for Concurrent Real-Time's Network Update Utility (NUU) displays, listing product updates that complete your installation of RedHawk. Depending on the installation state of your Concurrent Real-Time 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 repositories (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.

- When the installation completes, execute the following commands:

```
cd ..  
eject
```

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

5.7 Installing RedHawk Updates from Web Site

- Additional updates may have been issued (that are not included on any "RedHawk Update Disc" you may have received). They should be downloaded and installed now. Refer to "Software Updates" on page 72 for details.
- When complete, or if no updates are available, reboot the system choosing the new RedHawk kernel and verify that the system boots properly.

6.0. Upgrade Procedures

Use the upgrade procedures provided here only if you are upgrading from Version 8.0 to Version 8.2.

If generating or regenerating a RedHawk 8.2 system, migrating a system from an early Version 7.x or earlier, or converting a system from 32-bit to 64-bit, follow the steps in “Installation Procedures” on page 14.

NOTE: Resources to assist users supplying their own Red Hat software are available in the Concurrent Real-Time Software Documentation Library <http://redhawk.concurrent-rt.com/docs/>. Look under the corresponding RedHawk Linux Release for the link labeled : Installing on RHEL.

Review the following points before upgrading the software:

- It is recommended that the installer first read through the entire procedure once before actually performing the upgrade.
- It is **strongly recommended** that installation and updates of software be performed in single-user mode to guard against the possibility of corrupting a running application.
- Install all available upgrades. No quality claims are made by Concurrent Real-Time unless all appropriate version levels of all RedHawk Linux software are installed.
- The media in use is accessed at various times and must be kept in the drive during the installation until you are instructed to install different media.
- Prior to upgrade, remove any floppy media that may have been left in the machine.
- The CentOS kernel should be running when upgrading to RedHawk Linux Version 8.2.

6.1 Pre-Upgrade Procedure

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

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

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

In this example, the root partition on which this CentOS was installed is **/dev/sda2**.

Resolve any discrepancies before continuing with the upgrade.

- Avoid adding or removing drives to a system before or after an install or upgrade to prevent having difficulties with grub, the BIOS, and the operating system correctly identifying various drives and partitions.
- If you cannot reboot your system after the upgrade, refer to the section “How to Use Linux Rescue” on page 61.

6.2 RedHawk Pre-Upgrade Steps

Upgrading a system from RedHawk 8.0 to RedHawk 8.2 is a fairly straightforward process, however some manual steps need to be performed before you can begin the upgrade.

NOTE

RedHawk 8.2, unlike some previous versions of RedHawk, does *not* require the use of the original CentOS Enterprise Linux 8.2 installation disc to perform an upgrade.

To upgrade the system follow these steps, performing all commands as the *root* user.

1. Relocate several CentOS and RedHawk DNF repository files that will no longer be needed with RedHawk 8.2 as these may confuse NUU:

```
# mkdir /etc/yum.repos.old
# cd /etc/yum.repos.d
# mv *redhawk*8.* /etc/yum.repos.old
```

Note that after the upgrade has completed successfully you can safely remove the `/etc/yum.repos.old` directory as these repository files can no longer be used with RedHawk 8.2.

2. Update the `systemd` (1) service manager. Note that the update instructions differ depending on whether your system has access to the internet.

If the system to be upgraded has access to the internet, issue the following command and skip to next numbered instruction.

```
# dnf update systemd
```

If the system does not have access to the internet, follow these steps:

- a. Install the Centos 8.2 update media and verify it automounted under the path `/run/media/root/CentOS-8.2-Updates-x86_64`. Use the `mount` (8) command to verify where it is mounted.

If it was mounted under a different path (happens when the user does not login to the desktop as root), then follow these steps to mount it under the repo supported path:

```
# umount /dev/cdrom
# mkdir -p /run/media/root/CentOS-8.2-Updates-x86_64
# mount /dev/cdrom \
/run/media/root/CentOS-8.2-Updates-x86_64
```

- b. Once the media is mounted, save the system's repo file and copy the repo file in the CD to the `/etc/yum.repos.d` directory:

```
# cd /etc/yum.repos.d
# mv ./CentOS-Updates-Media.repo \
/etc/yum.repos.old
# cd /run/media/root/CentOS-8.2-Updates-x86_64
# cp ./CentOS-Updates-Media.repo \
/etc/yum.repos.d
```

- c. Update the systemd service manager from the Centos 8.2 update media:

```
# dnf --disablerepo='*' \
--enablerepo=centos-8.2-updates-media-BaseOS \
update systemd
```

3. If the system console is using an NVIDIA graphics card, you must rename the X Server's `xorg.conf` file so that it will not conflict with the *nouveau* driver in the CentOS kernel:

```
# mv -f /etc/X11/xorg.conf /etc/X11/xorg.conf.old
```

Once the system has been upgraded to RedHawk 8.2 a new `xorg.conf` file will be generated. If you previously made manual changes to the `xorg.conf` file you may wish to restore it to the version saved above after the upgrade is complete.

4. Remove the nuu package from the system:

```
# rpm -e nuu
```

5. Configure the system to boot a CentOS kernel by default using the `blscfg -d` command and reboot the system. Invoke the `blscfg` command first with no options to see the id number of the CentOS kernel. See the `blscfg(1)` man page for more information. |

```
# blscfg
# blscfg -d <CentOs-id-#>
# reboot
```

6.3 Upgrade RedHawk to RedHawk 8.2

Once the pre-upgrade steps have been completed in the previous section, the rest of the upgrade steps are identical to that of a full installation proceeding *after* the base CentOS has been installed.

Continue with the full installation instructions starting with section **Installing CentOS Updates** on page 31 and continuing through section **Installing RedHawk Updates from Web Site** on page 41. Once you complete those instructions the upgrade is complete.

NOTE

Installing the CentOS 8.2 Updates on an existing 8.0 system will produce warnings throughout the process; these warnings can be safely ignored.

7.0. NVIDIA Graphics Configuration

The Unified NVIDIA Linux Display Driver (version 450.57) is included in the RedHawk Linux installation and built automatically for all pre-built and custom RedHawk Linux kernels. The RedHawk installer probes for the type and model of video adapters present in the system, and recommends installing the unified `ccur-nvidia-glx` rpm when it detects one or more NVIDIA video cards present.

See sections **7.1 Supported NVIDIA GPUs** and **7.2 Multiple NVIDIA Video Card Considerations** for more information on NVIDIA driver compatibility.

After you have completed the basic installation for your display configuration, you can edit the resulting `/etc/x11/xorg.conf` for final adjustments.

7.1 Supported NVIDIA GPUs

If the `ccur-nvidia-glx` rpm has been installed on your system, you can view the list of all NVIDIA GPUs supported by the 450.57 driver in “Appendix A” in the following file:

```
/usr/share/doc/ccur-nvidia-glx-450-57/README.txt
```

Alternatively, you can examine the list of supported GPUs for driver 450.57 at the NVIDIA web site which is located at:

<http://www.nvidia.com>

Note that if your NVIDIA GPU is not supported by the 450.57 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 450.57 driver
- go to the NVIDIA web site mentioned above and download a legacy NVIDIA driver that supports your NVIDIA card, and follow the installation instructions described on that driver’s web page

NOTE

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

7.2 Multiple NVIDIA Video Card Considerations

Note that you may not use the unified 450.57 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 450.57 driver as your display driver.

7.3 Display Configurations

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

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

Use the **nvidia-settings** utility to create all multiple display configurations. See section **7.5 Using nvidia-settings**.

7.3.1 Single

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

Use the **nvidia-xconfig** utility to create a single NVIDIA display. See section **7.4 Using nvidia-xconfig**.

7.3.2 Xinerama

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

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

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

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

7.3.3 MultiDesktop

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

- A separate X Screen for each display is used.
- The mouse moves freely between displays.
- If you maximize a window, it fills a single display.

- You can *not* drag objects from one display to another.
- This causes a little extra overhead for X.

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

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

7.3.4 Twinview

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

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

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

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

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

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

See `/usr/share/doc/ccur-nvidia-glx-450-57/README.txt` for a detailed discussion on Twinview.

7.3.5 Twinview-Xinerama

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

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

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

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

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

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

7.3.6 Twinview-MultiDesktop

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

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

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

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

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

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

7.4 Using nvidia-xconfig

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

The **nvidia-xconfig** utility will generate some superfluous entries which you may choose to remove. See the example **xorg.conf** file in section 7.6 for reference.

Follow this procedure exactly to generate the file:

1. Read sections 7.0. - 7.3 and complete any actions required.
2. Reboot to runlevel 3.
3. Move **/etc/X11/xorg.conf** to **/etc/X11/xorg.conf.orig**:


```
mv -i /etc/X11/xorg.conf /etc/X11/xorg.conf.orig
```
4. Run the **nvidia-xconfig** utility (**/usr/bin/nvidia-xconfig**):


```
nvidia-xconfig
```
5. Reboot to runlevel 5.

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

7.5 Using nvidia-settings

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

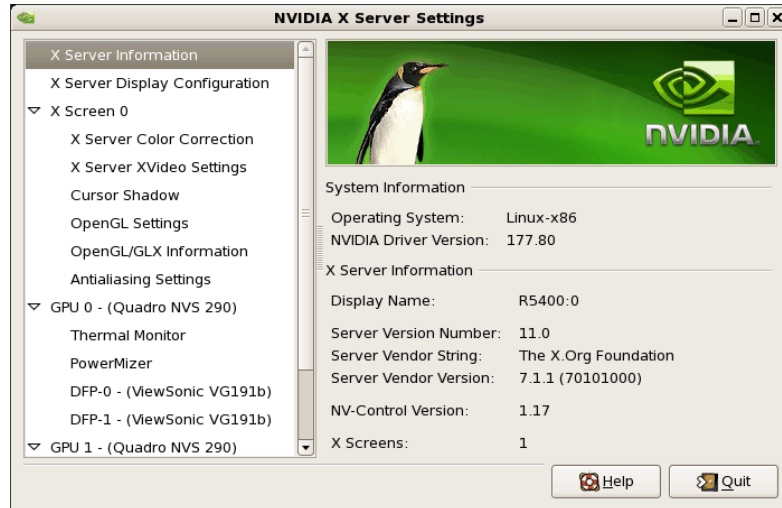
Before using **nvidia-settings**, complete the **nvidia-xconfig** procedure outlined in section 7.4. Failure to do so may result in poor or unstable performance.

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

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

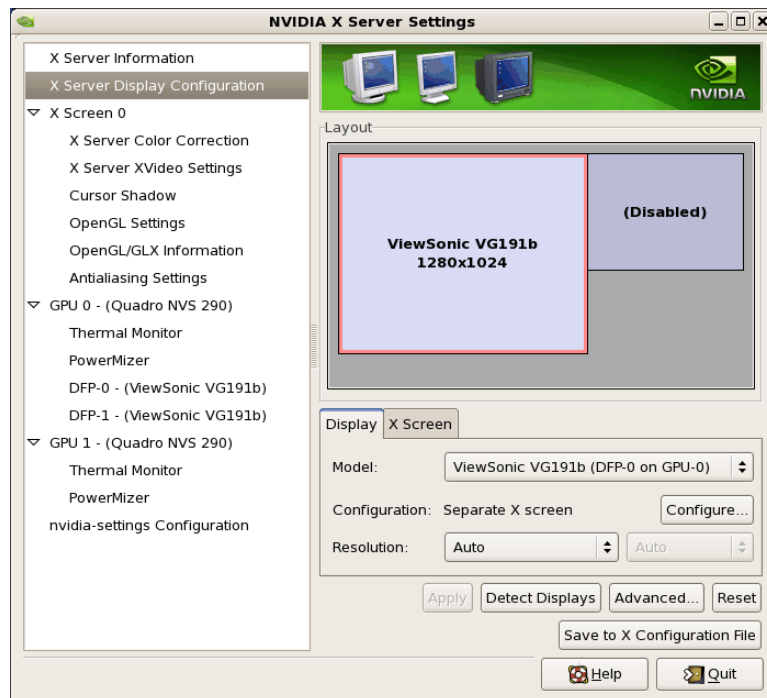
nvidia-settings

The GUI will start:



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

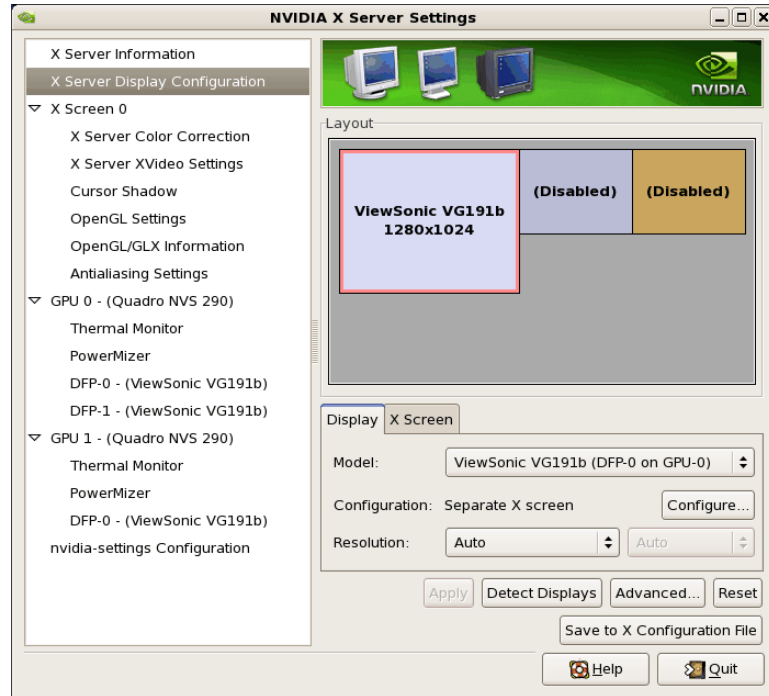
2. To create a new “X Server Display Configuration” select that option from the list on the left as shown below:



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:

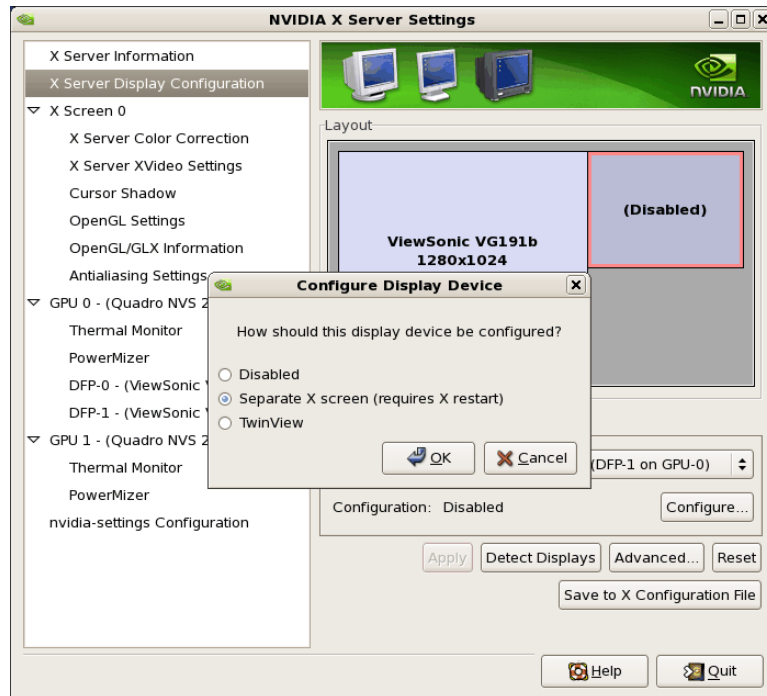


3. To configure two detected displays as a MultiDesktop, we would select the second monitor in the Layout window by clicking on it. We can see the


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**:



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

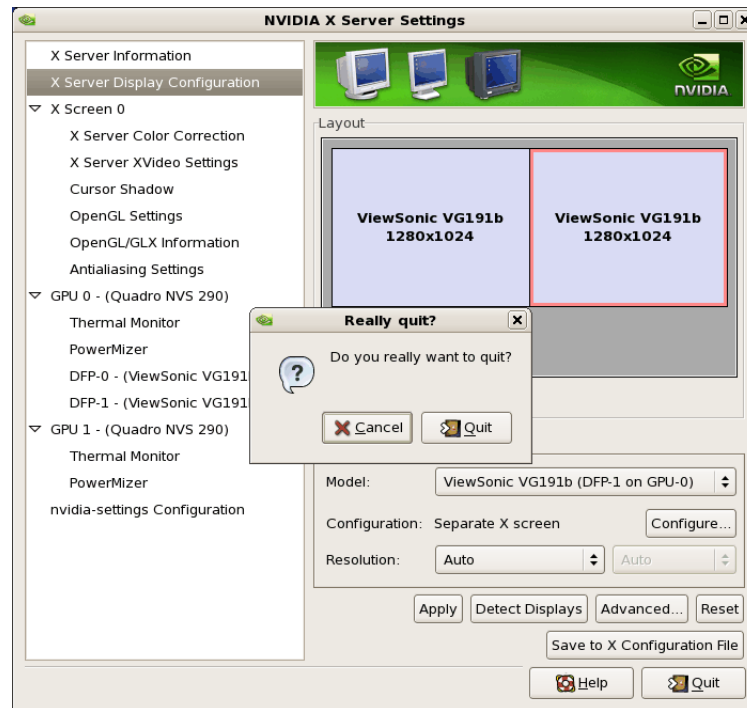


- To save the configuration to `xorg.conf`, click on the Save to X Configuration File button near the bottom of the GUI:




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

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

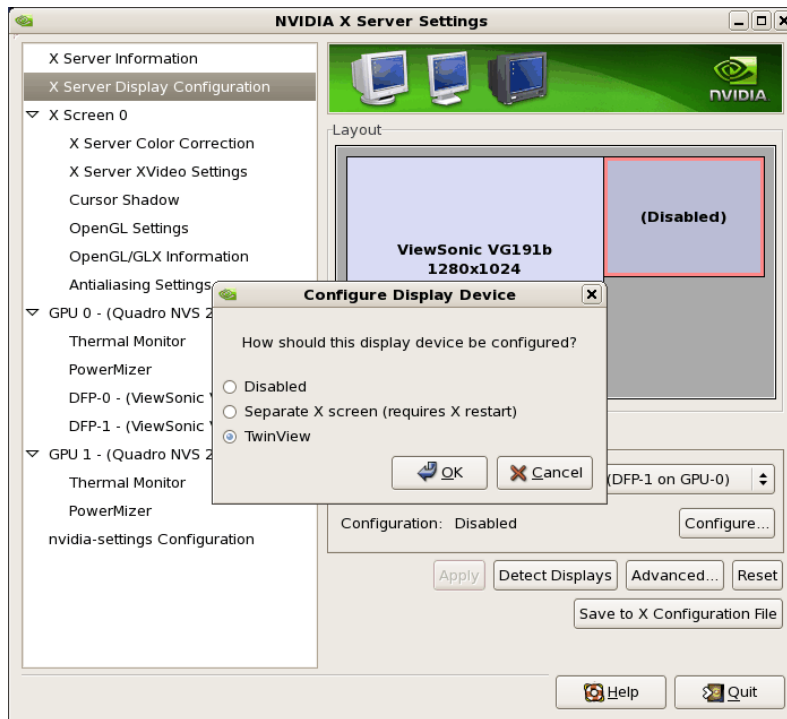


- Reboot to runlevel 5. You should now be up on two X screens. You can restart `nvidia-settings` to further examine or modify your configuration.

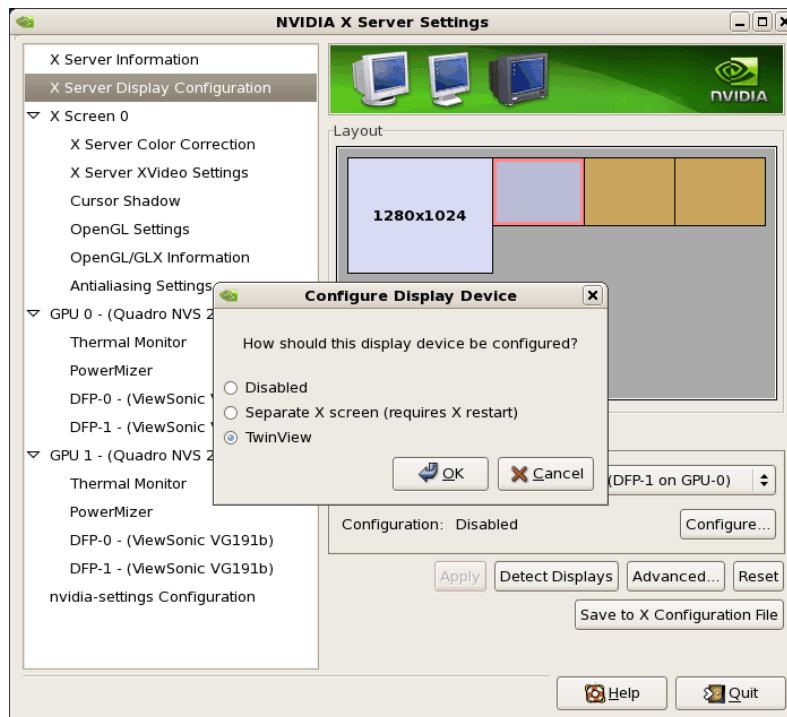
Below shows an example configuring the Xinerama option:

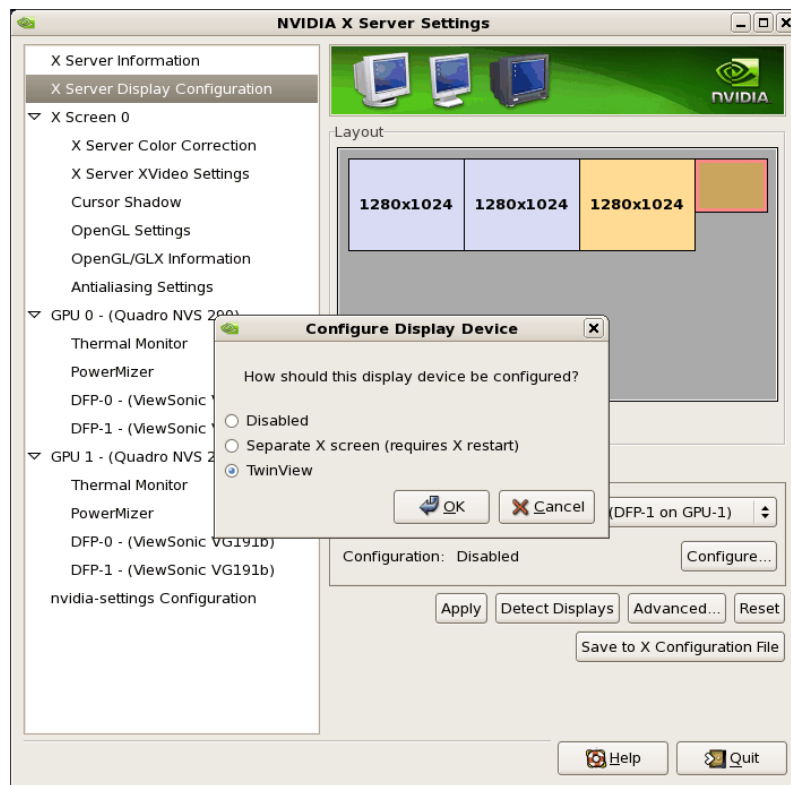
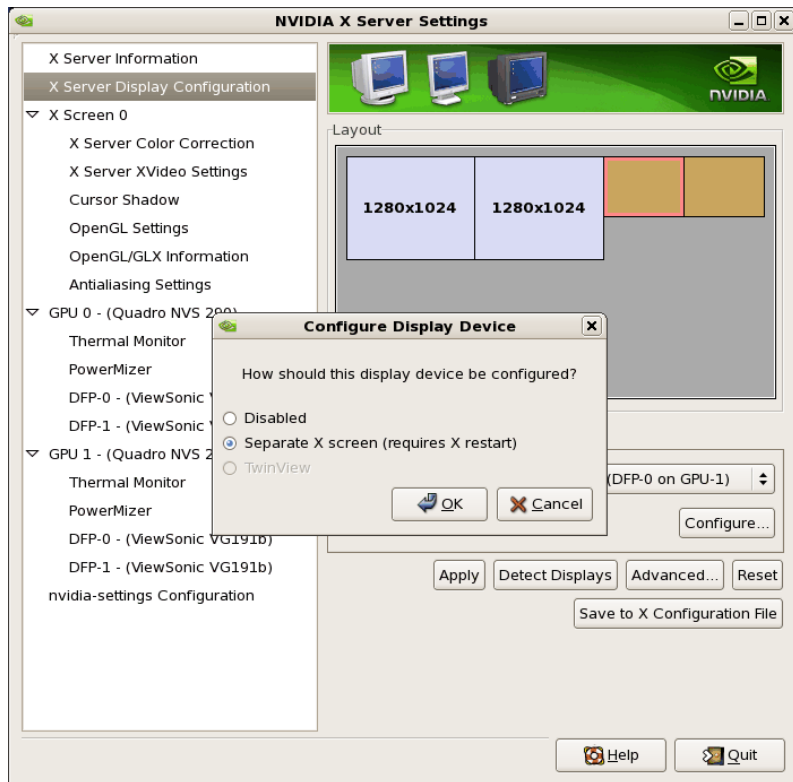


Below shows an example configuring the Twinview option:



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





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

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

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

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

7.6 Modified xorg.conf file generated by nvidia-xconfig

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

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

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

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

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

```

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

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

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

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

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

Section "Monitor"
    Identifier     "Monitor0"
EndSection

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

```

```

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

```

7.7 Troubleshooting Problems

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

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

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

```
init 1; init 5
```

or

```
reboot
```

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

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

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

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

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

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

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

```

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

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

```



```

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

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

```

Typical drivers in iHawk systems include:

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

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

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

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

WARNING

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

8.0. Additional Information

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

8.1 Installation/Configuration Issues

8.1.1 File System Backup Recommendations

All file systems should be archived or preserved prior to commencing a re-installation. This can be accomplished using normal backup methods. Concurrent Real-Time 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	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 > <i>save_file</i>	For reference to configure new system.
/etc/*.conf	Customized configuration files.
/etc/sysconfig/*	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 <code>/bin/rpm -qal --last more</code> to view.

8.1.2 File System Swap Size Guidelines

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

Table 7-2 Swap Partition Size Guidelines

Main Memory Size (GB)	Recommended Swap Partition Size
0 - 3	2048 MiB
4 - 15	4096 MiB
16 - 63	8192 MiB
64 - 255	16 GiB
256 - 511	32 GiB
512 - 1024	64 GiB

8.1.3 RCIM Connection Modes

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

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

8.1.4 How to Use Linux Rescue

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

- the disk LABEL scheme is not correct
- there are mismatching LABEL or partition entries in `/etc/fstab` and/or `/etc/grub2.cfg` -> `/boot/grub2/grub.cfg`
`/etc/grub2-efi.cfg` -> `/boot/efi/EFI/centos/grub.cfg`

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


It is possible to recover from these situations by booting into a Linux rescue image via the *Rescue a CentOS Linux system* boot option on the CentOS 8.2 installation media.

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

1. Boot a CentOS 8.2 installation disc.
2. Once the grub menu is displayed use the arrow keys to select Troubleshooting as shown in the following screen.



3. Press **Enter** and then select Rescue a CentOS Linux system as shown in the following screen.



- Press **Enter** and the system will boot into rescue mode from the live DVD. Once fully booted the following screen will be displayed.

```
Starting installer, one moment...
anaconda 29.19.2.17-1.el6 for CentOS Linux 8 started.
* installation log files are stored in /tmp during the installation
* shell is available on TTYZ
* when reporting a bug add logs from /tmp as separate text/plain attachments
=====
Rescue

The rescue environment will now attempt to find your Linux installation and
mount it under the directory : /mnt/sysimage. You can then make any changes
required to your system. Choose '1' to proceed with this step.
You can choose to mount your file systems read-only instead of read-write by
choosing '2'.
If for some reason this process does not work choose '3' to skip directly to a
shell.

1) Continue
2) Read-only mount
3) Skip to shell
4) Quit (Reboot)

Please make a selection from the above: 1

anaconda11:main* 2:shell 3:log 4:storage-log 5:program-log          Switch Tab: R11+Tab-1 Help: F1
```

- Press **1** to continue and the following text should be displayed.

```
=====
Rescue Shell
=====

Your system has been mounted under /mnt/sysimage.

If you would like to make the root of your system the root of the active system,
run the command:

    chroot /mnt/sysimage

When finished, please exit from the shell and your system will reboot.
Please press ENTER to get a shell:
```

- Press **Enter** to get to the root shell.
- You can now examine the files under `/mnt/sysimage` to correct issues identified in error messages before the system failed to boot properly.
- If you wish to transfer files or packages to the system over the network you can start networking via DHCP with the following command. Invoke the `ifconfig` command to find out the network device name if it is not `eth0` as used in this example.

```
ifup eth0
```

- After networking is enabled, you can also issue the following command to better simulate running commands on the actual system:

```
chroot /mnt/sysimage
```

At this point the system files on the hard drive appear in the usual locations in the file system, and you can run additional commands like **blscfg** in CentOS-based operating systems and **ccur-grub2** in Ubuntu, to change the default kernel to boot or modify kernel boot options. See the **blscfg(1)** and **ccur-grub2(1)** man page for more information.

8.1.5 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 -> Internet -> Remote Viewer**. Alternatively, you can install the *tigervnc* package and then run the **vncviewer** application.

Perform the following steps to install using VNC.

On the system you wish to install:

1. Follow the procedure outlined in section **5.0. Installation Procedures** on page 14 until you reach **5.1.1** step 4.

Select “Install CentOS 8” from the grub menu.

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

```
> vmlinuz 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 VNC Viewer.

On the gnome menu: Applications -> Internet -> Remote Viewer

Or from the command line using **vncviewer** if you previously installed the *tigervnc* package.

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 16. After completing step 23 on page 29, perform the steps below back on the actual system's graphical console.
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.

8.1.6 If the Installation GUI Fails to Start

If during installation the GUI fails to start, it may be caused by graphics driver compatibility problems or with certain graphics cards and/or monitors not being probed properly. Adding the **nomodeset** boot parameter to the kernel may solve this problem.

To modify boot parameters first stop the grub automatic boot countdown by pressing the arrow keys on the keyboard. You should now see a screen similar to the following.



```
RedHawk Linux (5.4.61-rt37-RedHawk-8.2-prt-trace) 8.2 (Taal)  
RedHawk Linux (5.4.61-rt37-RedHawk-8.2-prt-debug) 8.2 (Taal)  
RedHawk Linux (5.4.61-rt37-RedHawk-8.2-trace) 8.2 (Taal)  
RedHawk Linux (5.4.61-rt37-RedHawk-8.2-prt) 8.2 (Taal)  
RedHawk Linux (5.4.61-rt37-RedHawk-8.2-kdump) 8.2 (Taal)  
RedHawk Linux (5.4.61-rt37-RedHawk-8.2-debug) 8.2 (Taal)  
RedHawk Linux (5.4.61-rt37-RedHawk-8.2) 8.2 (Taal)  
CentOS Linux (4.18.0-193.14.2.el8_2.x86_64) 8 (Core)  
CentOS Linux (4.18.0-193.el8.x86_64) 8 (Core)  
CentOS Linux (0-rescue-abd81b14ccde4153b4066a9c0f554e4a) 8 (Core)
```

Use the ↑ and ↓ keys to change the selection.
Press 'e' to edit the selected item, or 'c' for a command prompt.
The selected entry will be started automatically in 4s.

Next press the **e** key to edit the grub boot entry for the highlighted kernel (this example is using the trace kernel) and a screen similar to the following will appear.

```
load_video
set gfx_payload=keep
insmod gzio
linux ($root)/vmlinuz-5.4.61-rt37-RedHawk-8.2-trace root=/dev/mapper/cl_ihawk-\
root ro resume=/dev/mapper/cl_ihawk-swap rd.lvm.lv=cl_ihawk/root rd.lvm.lv=cl_\
ihawk/swap rhgb quiet
initrd ($root)/initramfs-5.4.61-rt37-RedHawk-8.2-trace.img $tuned_initrd

Press Ctrl-x to start, Ctrl-c for a command prompt or Escape to
discard edits and return to the menu. Pressing Tab lists
possible completions.
```

Press the down arrow to scroll down until the line beginning with `linux ($root)` appears. Use the arrow keys to append the `nomodeset` boot option to the end of this line, as shown in the following screen.

```
load_video
set gfx_payload=keep
insmod gzio
linux ($root)/vmlinuz-5.4.61-rt37-RedHawk-8.2-trace root=/dev/mapper/cl_ihawk-\
root ro resume=/dev/mapper/cl_ihawk-swap rd.lvm.lv=cl_ihawk/root rd.lvm.lv=cl_\
ihawk/swap rhgb quiet nomodeset
initrd ($root)/initramfs-5.4.61-rt37-RedHawk-8.2-trace.img $tuned_initrd

Press Ctrl-x to start, Ctrl-c for a command prompt or Escape to
discard edits and return to the menu. Pressing Tab lists
possible completions.
```

Once the option has been added press **Ctrl-x** to boot the kernel and hopefully the new option will improve the graphics startup situation.

8.2 Considerations

8.2.1 Compiler Requirements

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

8.2.2 Hyper-thread Activation

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

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

9.0. Known Issues

Special consideration should be given to the following areas.

Problems executing programs that use RedHawk system calls

The RedHawk system call IDs have been changed to start at 1900. Verify that the changes outlined in section 4.3.3, RedHawk system calls have new IDs, have been made.

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 8.2 is based on CentOS Enterprise Linux 8.2. Binary and source compatibility issues exist for MAXAda programs between CentOS 3.0 and CentOS 8.2.

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-F1** to switch to virtual terminal 1.

Excessive ksoftirqd Runs Affecting Determinism

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

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

Pings to Multicast Addresses Disabled by Default

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

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

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

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

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

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

Grub Option Holdovers from Previous Kernel Installations

Be aware that when kernel packages are installed, they add a grub entry with options associated with that kernel (e.g., trace and debug kernels add “`crashkernel=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, adding a `vers=3` mount option to the corresponding entry in `/etc/fstab` will allow the NFS filesystem to be mounted properly.

Unable to Login After Enabling SELinux

If you experience problems logging into RedHawk systems after enabling SELinux (e.g. after entering login and password the system immediately returns to the login screen) then you will need to relabel the filesystem to be compatible with SELinux. Reboot the system using the CentOS kernel and issue the following command as the root user:

```
# touch /.autorelabel
```

Then reboot the system to the RedHawk kernel and you should be able to login successfully.

10.0. Software Updates and Support

10.1 Direct Software Support

Software support is available from a central source. If you need assistance or information about your system, please contact the Concurrent Real-Time 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 Real-Time, Inc. web site at <http://concurrent-rt.com/support>

10.2 Software Updates

Updates to Concurrent Real-Time RedHawk software can be obtained via Concurrent Real-Time'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 Real-Time's software repositories
- Building a customized Update disc using Concurrent Real-Time's web site for immediate download

10.2.1 Updating via NUU

NUU supports installation and updating of software products from Concurrent Real-Time software repositories over a network. NUU utilizes DNF 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 Real-Time software products installed on your system.

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

NOTE

It is recommended that all CentOS repositories should be disabled when checking for Concurrent Real-Time 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 when you click on the NUU link on the redhawk.concurrent-rt.com website or directly via this link <http://redhawk.concurrent-rt.com/network/QuickStart.pdf>


The first time you invoke NUU you will need to specify your redhawk.concurrent-rt.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 72.


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


10.2.2 Installing Manually Downloaded RPMs

You can browse Concurrent Real-Time’s software repositories to locate updated RPMs and download them for manual installation.

Access the RedHawk Updates web site (<http://redhawk.concurrent-rt.com/>) by clicking on the “Concurrent Real-Time Software Portal” icon on the desktop. Below is the top half of the screen displayed when you access this web site.




concurrent
REAL-TIME

 **redhawk.concurrent-rt.com**

This site provides you access to Concurrent Real-Time software products and updates to those products that execute on the following platforms:


- RedHawk Linux
- CentOS
- Red Hat Enterprise Linux
- Ubuntu

You can download [NUU](#), Concurrent Real-Time's Network Update and Installation utility, browse product repositories to manually download software packages, or create customized Update CD images for immediate download.


Services:

- [Document Library](#)
- [Network Repositories](#)
- [Browse Repositories](#)
- [Create Update Images](#)
- [Manage NightStar Licenses](#)
- [Customer Survey](#)

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



concurrent
REAL-TIME

 **redhawk.concurrent-rt.com**

Browse Software Repositories

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

For Ubuntu-based systems, there is a similar mechanism using standard Ubuntu utilities. See [APT Repositories](#) for more information.

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

Enter your redhawk.concurrent-rt.com Login and Password and press the browse button to proceed:





Login:

Password:

Enter your redhawk.concurrent-rt.com Login ID and Password and click the **Browse** button.

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

Index of /eportal/repos/RedHawk/8.2

Name	Last modified	Size	Description
 Parent Directory/		-	
 i386/	23-Dec-2020 12:14	-	
 isos/	23-Dec-2020 12:47	-	
 x86_64/	23-Dec-2020 12:58	-	

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

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

To install newly downloaded packages, follow these steps:

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

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


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

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

NOTE

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

10.2.3 Customized Update Discs

You can use Concurrent Real-Time'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.concurrent-rt.com>) by clicking on the "Concurrent Real-Time Software Portal" icon on the desktop, then click on **Create Update Images**.

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

10.3 Documentation Updates

For the latest documents, including updated Release Notes and User Guides, go to Concurrent Real-Time's documentation web site at <http://redhawk.concurrent-rt.com/docs>.

You can also view the latest RedHawk FAQ at <http://redhawk.concurrent-rt.com/docs/root/1Linux/1RedHawk/RedHawk-FAQ.pdf>

A

Ubuntu Support


This appendix contains additional information on using RedHawk 8.2 with the open source Ubuntu Linux distribution.

Overview

RedHawk 8.2 is compatible with the Ubuntu 20.04 LTS (Long Term Support) distribution. This appendix provides instructions for installing RedHawk 8.2 onto a brand new Ubuntu 20.04 installation. Installing RedHawk 8.2 onto an existing Ubuntu 20.04 system should also work correctly, however that is beyond the scope of this document.

Installing Ubuntu

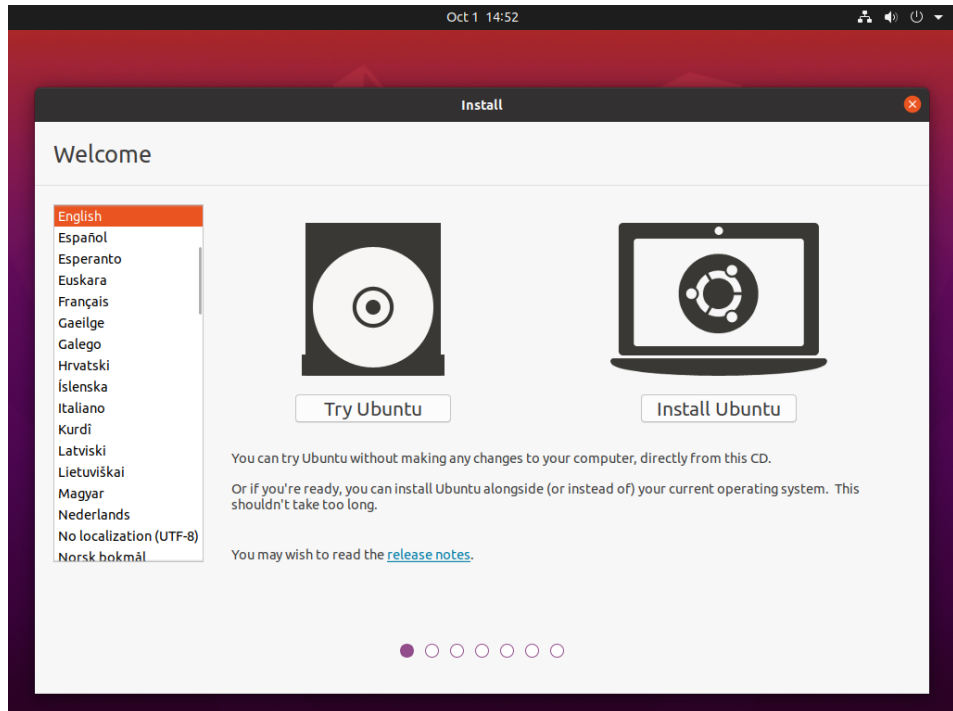
First, download the latest Ubuntu 20.04 LTS Desktop ISO from the [Ubuntu.com](https://ubuntu.com) website and burn it onto a DVD. Boot the DVD in your target system. The following page will appear.



NOTE

If your system is using UEFI firmware, you will come up on a page with a boot menu. Double click on **Ubuntu** at the top and on the next page click on the **Install Ubuntu** icon. Starting with the **Keyboard Layout** page, the pages should be the same as shown below.

Shortly after, the **Welcome** page will be displayed as below:




Press the **Install Ubuntu** button to begin the installation.


NOTE

If instead of the **Welcome** page you see a persisting dark screen, it may be that the installer system is incompatible with the graphics card on the target system. In that case, you may still be able to install in a safe graphics mode. See “Installing Ubuntu in Safe graphics mode” on page A-12.

Next, select the desired layout in the Keyboard layout page shown below.

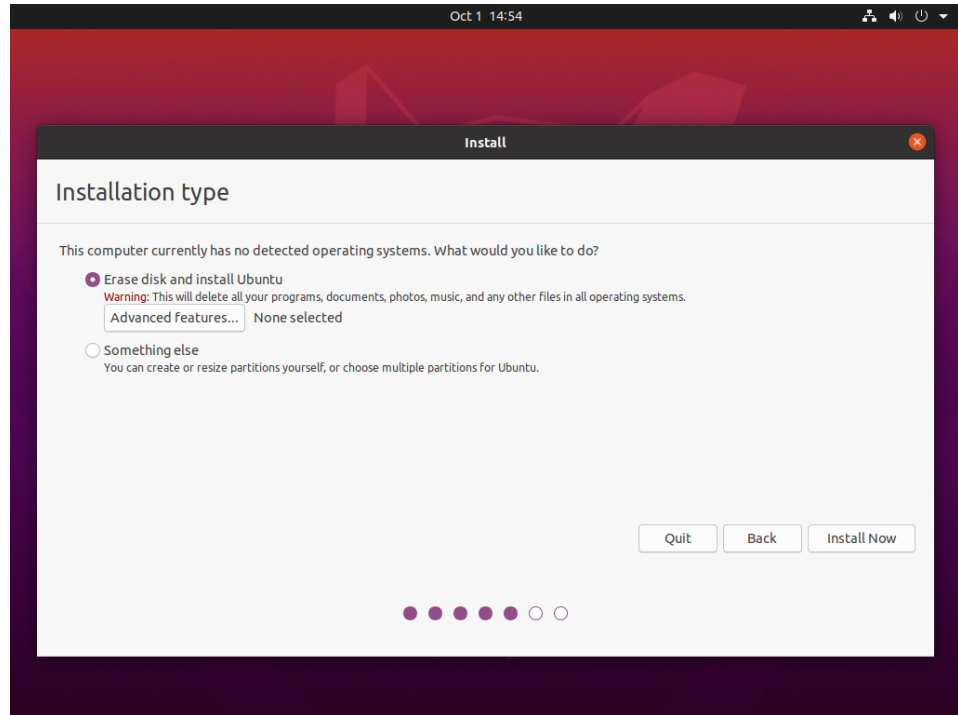


Press Continue and the Updates and other software page should be displayed as follows:



Check the boxes for **Normal installation** and **Download updates while installing Ubuntu**. The download option will download, but not install, updates.


Press **Continue** and the **Installation type** screen should be displayed as follows:



Choose **Erase disk and install Ubuntu** and then press **Install Now** to continue. You will then be prompted to confirm the action. Press **Continue** to confirm and continue.

In the next few pages you will be prompted to configure various settings including **Timezone**, **Username**, **System Name** and **Password**. After configuring these settings, press **Continue** to continue with the installation.

Depending on system and network speed, the installation can take up to 60 minutes to complete; once the installation has finished, the following message will be displayed.



Press **Restart Now** to reboot the newly installed system.

Updating Ubuntu Software

While updates were downloaded on the system during the install, the updates were not applied. Follow the instructions below to update the system now.

Once the system has booted to a graphical login screen, log into the system and bring up a terminal window.

Execute the following command to update the cached list of available packages and their versions.

```
# sudo apt update
```

The screen image below shows an example output from the command.

```

me@ihawk: ~/Desktop
To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

me@ihawk:~/Desktop$ ls
me@ihawk:~/Desktop$ sudo apt update
[sudo] password for me:
Hit:1 http://us.archive.ubuntu.com/ubuntu focal InRelease
Get:2 http://us.archive.ubuntu.com/ubuntu focal-updates InRelease [111 kB]
Get:3 http://security.ubuntu.com/ubuntu focal-security InRelease [107 kB]
Get:4 http://us.archive.ubuntu.com/ubuntu focal-backports InRelease [98.3 kB]
Fetched 317 kB in 1s (399 kB/s)
Reading package lists... Done
Building dependency tree
Reading state information... Done
176 packages can be upgraded. Run 'apt list --upgradable' to see them.
me@ihawk:~/Desktop$

```

To install new packages and newer versions of packages, issue the following command:

```
# sudo apt upgrade
```


You will be presented with information about the pending upgrade and then be prompted to continue

```

me@ihawk: ~/Desktop
libproxy1v5 libpulse-mainloop-glib0 libpulse0 libpulsedsp
libreoffice-ogltrans libreoffice-pdfimport libreoffice-style-breeze
libreoffice-style-colibre libreoffice-style-elementary
libreoffice-style-tango libridl-java libsane libsane-common libsmartcols1
libsmbclient libsnmp-base libsnmp35 libssh-4 libuno-cppu3
libuno-cppuhelpergcc3-3 libuno-purpenvhelpergcc3-3 libuno-sal3
libuno-salhelpergcc3-3 libunoloader-java libuuid1 libuv1 libwbclient0
libwebkit2gtk-4.0-37 libwhoopsie0 libwinpr2-2 libx11-6 libx11-data
libx11-xcb1 linux-firmware memtest86+ mount mutter mutter-common packagekit
packagekit-tools ppp pulseaudio pulseaudio-module-bluetooth pulseaudio-utils
python3-apport python3-aptdaemon python3-aptdaemon.gtk3widgets
python3-commandnotfound python3-distupgrade python3-problem-report
python3-software-properties rfkill rhythmbox-plugin-alternative-toolbar
rsyslog samba-libs sane-utils secureboot-db simple-scan snapd
software-properties-common software-properties-gtk sudo ubuntu-desktop
ubuntu-desktop-minimal ubuntu-minimal ubuntu-release-upgrader-core
ubuntu-release-upgrader-gtk ubuntu-standard unattended-upgrades
uno-libs-private ure util-linux uuid-runtime whoopsie xserver-common
xserver-xephyr xserver-xorg-core xserver-xorg-legacy xwayland xz-utils
yaru-theme-gnome-shell yaru-theme-gtk yaru-theme-icon yaru-theme-sound
176 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
Need to get 0 B/284 MB of archives.
After this operation, 7,905 kB of additional disk space will be used.
Do you want to continue? [Y/n]

```


During the upgrade, dialogs also may be displayed asking you to confirm other actions. The example dialog below asks permission to update the grub configuration file. Use the Up Arrow key to select to install the package maintainer's version, use the Tab key to select Ok, then press Enter.



When the upgrade completes, issue the reboot command to reboot the system:

```
# sudo reboot
```

Installing RedHawk Linux

Once the system has booted to a graphical login screen, log into the system. Insert the RedHawk OS Installation DVD into the system's optical media drive and close any file-browser window that opens when the DVD is inserted. Then open a terminal window, and cd to the mounted directory of the DVD.


NOTE

If an issue is encountered when installing from DVD, execute the following command after inserting the DVD.

```
sudo setfacl -m g:::5,o:::5 /media/*
```

Execute the following commands to install RedHawk Linux, substituting the actual mount point used for your optical media device.

```
cd /path-to-actual-mount-point
sudo ./install-redhawk
```



```
me@ihawk: /media/me/RedHawk-8.2-x86_64
me@ihawk: /media/me/RedHawk-8.2-x86_64$ sudo ./install-redhawk
[sudo] password for me:

This process will install the RedHawk Linux core files
for RedHawk Linux 8.2 (Taal) on the current system (ihawk).

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

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

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

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

-----

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

This system does not appear to have an NVIDIA graphics card.
We do not recommend installing RedHawk NVIDIA graphics packages.

Install RedHawk NVIDIA graphics packages anyway? [y/n] (n)
```

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.


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.

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

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.

Messages similar to the ones below will be displayed once all RedHawk package installation has completed successfully.



```
me@ihawk: /media/me/RedHawk-8.2-x86_64
Created symlink /etc/systemd/system/multi-user.target.wants/ccur-default-kernel.service → /lib/systemd/system/ccur-default-kernel.service.
Setting up ccur-xtrace (8.2-20200930) ...
Setting up ccur-rcim (8.2-20200930) ...
Created symlink /etc/systemd/system/multi-user.target.wants/rcim.service → /lib/systemd/system/rcim.service.
Setting up zlib1g-dev:amd64 (1:1.2.11.dfsg-2ubuntu1) ...
Setting up libncurses5-dev:amd64 (6.2-0ubuntu2) ...
Setting up ccur-gather (8.2-20200930) ...
Setting up libelf-dev:amd64 (0.176-1.1build1) ...
Setting up ccur-kernel-source (8.2-20200930) ...
Processing triggers for mime-support (3.64ubuntu1) ...
Processing triggers for gnome-menus (3.36.0-1ubuntu1) ...
Processing triggers for libc-bin (2.31-0ubuntu9.1) ...
Processing triggers for systemd (245.4-4ubuntu3.2) ...
Processing triggers for man-db (2.9.1-1) ...
Processing triggers for desktop-file-utils (0.24-1ubuntu3) ...

The RedHawk Linux installation is complete.

Please contact Concurrent Real-Time technical support if you had any
problems during this installation (http://www.concurrent-rt.com/support
or 1-800-245-6453).

me@ihawk: /media/me/RedHawk-8.2-x86_64$
```

Eject and remove the disc from the drive and reboot the system choosing a newly installed RedHawk kernel and verify that the system boots up to the graphical login.

At this point you have completed the basic RedHawk Linux installation. If you wish to install additional RedHawk software, refer to the following sections in the main chapter of the Release Notes:

“Installing the Frequency-Based Scheduler (FBS)” on page 40

“Installing Additional RedHawk Products” on page 40

Updating RedHawk Software

NOTE

The latest and most up-to-date instructions for installing software via Concurrent’s network package repositories can always be found on the <http://redhawk.concurrent-rt.com/network> web page.

Perform the following steps to allow the Ubuntu package management software to access Concurrent Real-Time's Debian repositories.

1. Obtain Concurrent Real-Time's public key by issuing the following command:

```
sudo wget http://redhawk.concurrent-rt.com/network/  
ccur-public-key
```

2. Add this key to the package management database by issuing the following command:

```
sudo apt-key add ccur-public-key
```

3. Create a new file named `redhawk.list` in the `/etc/apt/sources.list.d` directory with the line below, replacing *Login* and *Pass* with the appropriate values.

```
deb http://redhawk.concurrent-rt.com/ubuntu/Login/  
Pass/redhawk/generic 8.2.n core
```

Note that the Login ID (also referred to as your Site ID) and password are assigned to you from Concurrent Real-Time and are included in a cover letter that is delivered with your system or with your software.

Also note that in addition to the core channel shown in the example above you can add `fbs` and/or `reimgps` to the end of the line if you have purchased those optional products.

4. Issue the following command to update the currently cached package repository information:

```
sudo apt-get update
```

This command will *not* change any software on your system; it merely updates all of the package repository definitions from the `/etc/apt/sources.list` and the files under `/etc/apt/sources.list.d`.

5. Use the following command to bring your system up-to-date with the latest software packages from Concurrent Real-Time:

```
sudo apt-get install 'ccur-*
```

It is important that you use single quotes around the `ccur-*` text so that the command will correctly match all available packages.

Alternatively, you can install or update individual specific Concurrent Real-Time packages using a command similar to the following example:

```
sudo apt-get install ccur-kernel-trace ccur-vmlinux-trace
```

Ubuntu-Specific Notes

This section captures various Ubuntu-specific procedures that differ from CentOS procedures that existing RedHawk customers may be familiar with.

Building and Installing Custom Kernels

To install a newly built RedHawk customer kernel you should follow the usual steps that are described in the `ccur-config(1)` man page. However, you will also need to issue some additional commands to complete the kernel installation. See example below.

Note

The following additional packages are not in the default installation and must be installed on the system in order to build a RedHawk custom kernel: `make`, `gcc`, `flex`, `bison`, `qt5-default`, `g++`.

The following example shows the commands to configure, build and fully install a custom kernel based on the RedHawk trace configuration. Note that warnings can be safely ignored if the subsequent commands successfully complete.


```
# cd /usr/src/  
# cd linux-5.4.66RedHawk8.2  
# sudo ./ccur-config trace  
# sudo make -j 10 bzImage modules  
# sudo make -j 10 modules_install  
# sudo make -j 10 install  
# sudo update-initramfs -c -k 5.4.66-rt38-RedHawk-8.2-custom  
# sudo update-grub
```

To boot from the custom kernel run `ccur-grub2(1)` without arguments to obtain the grub index number of the custom kernel. Then execute the command a second time to make the custom kernel the default kernel as follows:

```
# sudo ccur-grub2 -d <index number>
```

Installing Ubuntu in Safe graphics mode


Boot from DVD and press the Down Arrow key when you see the following page. Pay close attention as the page is up only for a short time.



Select the appropriate language, then press Enter.



Use the Down Arrow key to select Install Ubuntu (safe graphics).



Once selected press **Enter** to continue. Follow the installation instructions at the beginning of this chapter to complete the installation.

B

ARM64 Support

This appendix contains additional information on using RedHawk 8.2 with hardware that utilizes the 64-bit ARMv8-A processor architecture (a.k.a. *ARM64*, a.k.a. *Aarch64*).

Overview

RedHawk 8.2 has been ported to the ARM64 architecture and has currently been qualified on the following development systems:

- Applied Micro X-C1 Server Development Platform (a.k.a. Mustang)
- ARM Juno R1 64-bit Development Platform

Support for additional ARM64 systems is planned. Customers interested in other systems should contact their regional Concurrent Real-Time sales representative for more information.

Installation

Installation of RedHawk 8.2 for ARM64 is supported utilizing both the CentOS 8.2 and Ubuntu 18.04 LTS base distributions. Unfortunately, the installation of these base distributions on all currently available ARM64 hardware varies considerably and is beyond the scope of this document.

Concurrent Real-Time recommends that you purchase pre-installed ARM64 systems from Concurrent Real-Time instead of trying to install them yourself. However, for those up for the challenge, the following guidelines may help guide you during the initial base distribution installation:

- For RedHawk to work the board must be configured with UEFI firmware. Boards that come pre-installed with other firmware (e.g. the *Das U-Boot* boot loader) cannot be used with RedHawk until they are converted to UEFI. See the many resources on the Internet for conversion instructions.
- Few ARM64 systems feature a graphical console, however both CentOS and Ubuntu support installation using a simple text interface over the serial console. In addition, CentOS can still be installed graphically by utilizing a remote VNC connection. Refer to Internet resources for more information.
- Often the easiest method for performing a CentOS or Ubuntu installation is to connect a USB DVD reader to the ARM64 system with the distribution's optical media disc inserted. While it is unlikely that the system will be able

to automatically boot the disc's installer, once the system is booted to a UEFI shell manual commands can be issued to boot the disc's installer.

- Once the base distribution installation has started successfully, refer to “Installing CentOS” on page 15 or “Installing Ubuntu” on page A-1 to complete the base distribution installation.

After performing any necessary initial low-level configuration of the ARM64 system and successfully installing the base distribution, continue with “Installing RedHawk Linux” on page 34 for CentOS systems or “Installing RedHawk Linux” on page A-7 for Ubuntu systems.

C

UEFI Secure Boot

This appendix contains instructions to configure a RedHawk Linux system to boot in UEFI Secure Boot mode.

NOTE

UEFI Secure Boot is supported in RedHawk Linux version 8.0.1 and later releases; it is not supported in 8.0.

Overview

UEFI Secure boot is a verification mechanism ensuring that code launched by firmware is trusted. UEFI Secure Boot requires that each binary loaded at boot is validated against known keys, located in firmware. A chain of trust is then established from the firmware to the signed drivers and kernel modules.


Configuration

Booting a RedHawk Linux system in UEFI Secure Boot mode requires that the BIOS is modified to enable Secure Boot and that the RedHawk machine owner key (MOK) is added to the system's MOK list.

The following instructions will configure a RedHawk Linux system to boot in UEFI Secure Boot mode.

1. If you have not yet installed the RedHawk Linux system, install it now. Note that the system must be installed UEFI-enabled and running a RedHawk Linux version 8.0.1 or later. See “Installation Procedures” on page 14 of this document for help.
2. Reboot the system and enter the BIOS Setup. Verify that the system is in UEFI Boot mode and then enable Secure Boot. Save the BIOS change and reset the system. When the kernel boot menu appears, use the up and down arrow keys to choose a base (CentOS or Red Hat) kernel from the kernel boot menu.

The following screen shows the enabling of Secure Boot in the BIOS.



- Once booted, open a terminal window and execute the following **mokutil (1)** command to verify that the system is booted with Secure Boot enabled.

```
# mokutil --sb-state
SecureBoot enable
```

- Next, issue the following **mokutil (1)** command to request for the enrollment of the RedHawk machine owner key (MOK). The command will prompt to create and confirm a password for this request.

```
# cd /boot/efi
# mokutil --import ./redhawk-ca.cer
input password:
input password again:
```

- Now configure the RedHawk trace kernel as the default boot kernel. Use **blscfg** to find the index of the RedHawk trace kernel. If it is not set as the default boot kernel, then use the **-d** option to set it as the default.
- Reboot the system. Just before loading GRUB, on the system console, the MOK Management menus will prompt you to complete the enrollment request. The screens that follow show an example MokManager key enrollment.

NOTE

The MokManager window displayed on the console must be accessed within 10 seconds or it will time out and resume booting without enrolling the MOK.


NOTE

The enrollment request is remembered only on the boot following when the request is issued. If the MokManager window times out or there is an issue during the enrollment, you must reboot a base system kernel (CentOS or Red Hat) and reissue the request, starting again with step 4 above.


Press any key, within 10 seconds, to start MOK management.




Select 'Enroll Key from disk' and press ENTER.




Press ENTER on the selected key to continue:.




Select the key file, `redhawk-ca.cer`, and press ENTER:




Next, a screen (not shown here) will give the option to verify the key or/and continue. Select view key to view the key and then press ENTER to go back. Select continue and press ENTER to continue. Below is an example of the key verification screen.



Select **Yes** to enroll the key and then press ENTER.



Select **Reboot** and press ENTER.



Note that on some systems, after successfully completing the enrollment, the reboot will take you back into the MokManager menu. If this happens, allow it to time out and continue booting.

7. Once booted, verify that you are booted on the RedHawk Linux trace kernel (`uname (1)`) and that Secure Boot is enabled (`mokutil --sb-state`).