



Linux® OS Version 2.0 Release Notes

December 17, 2003

0898003-2.0

**READ ME BEFORE
INSTALLING
THIS PRODUCT**



Disclaimer

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

License

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

Trademark Acknowledgments

Linux is a registered trademark of Linus Torvalds.

Red Hat is a registered trademark of Red Hat, Inc.

RedHawk, iHawk, NightStar, NightView, NightSim, and NightProbe are trademarks of Concurrent Computer Corporation.

AMD Opteron and AMD64 are trademarks of Advanced Micro Devices, Inc.

POSIX is a registered trademark of the Institute of Electronics and Electrical Engineers.

The X Window System is a trademark of The Open Group.

PowerEdge is a trademark of Dell Computer Corporation.

All other trademarks are the property of their respective owners.

Contents

1.0. Introduction	1
1.1. Product Media	1
1.2. System Updates	2
2.0. Documentation	3
2.1. RedHawk Linux Documentation	3
2.2. Vendor Documentation	3
3.0. Prerequisites	4
3.1. Software	4
3.2. Hardware	4
4.0. Version 2.0 Installation Procedures	5
4.1. Installing Red Hat Software	5
4.2. Installing Red Hat Software Updates	8
4.3. Installing RedHawk Linux	9
4.4. Installing RedHawk Updates	10
4.5. Installing the RCIM	10
4.5.1. Quick Hardware Installation Checklist	10
4.5.2. Check RCIM Operation	11
4.6. Installing Frequency-Based Scheduler (FBS)	
Software	11
4.7. Installing FBS Updates	12
4.8. Installing NightStar Tools	12
5.0. NVIDIA Graphics Display	13
5.1. Configuring Multiple Display Capability	14
5.2. Connecting the Monitors and Booting the System	16
6.0. Additional Considerations	18
6.1. Compiler Requirements	18
6.2. NightView Accessibility	18
7.0. Information on Version 2.0	19
7.1. General	19
7.2. RedHawk Linux Kernels	19
7.3. Real-Time Features in RedHawk Linux	20
7.4. NightStar Development Tools	24
8.0. Changes in this Release	25
8.1. Enhancements	25
8.1.1. Support for 64-bit Processing	25
8.1.2. Dynamic Creation of Custom Kernel Trace Events	25
8.2. Modifications	25
8.2.1. High Resolution Process Accounting	25
8.2.2. Frequency-Based Scheduler and Performance Monitor	26
8.2.3. Real-Time Clock and Interrupt Module (RCIM)	26
8.2.4. Rescheduling Variables	26
8.2.5. POSIX Message Queue Facility	27
9.0. Known Issues	28

- 10.0. Using the Linux Kernel Crash Dump (LKCD) 30
 - 10.1. Documentation 30
 - 10.2. Installation/Configuration Details 30
 - 10.3. Forcing a Crash Dump on a Hung System 31
- 11.0. Useful Information 32
 - 11.1. Man Pages 32
 - 11.2. RCIM Connection Modes 32
 - 11.3. File System Swap Size Guidelines 32
- 12.0. Software Updates and Support 33
 - 12.1. General 33
 - 12.2. How to Download and Install Updates 33
- 13.0. Direct Software Support 34

1.0. Introduction

This document provides release information and installation instructions for Concurrent Computer Corporation's RedHawk™ Linux® operating system Version 2.0. This release contains the initial 64-bit processing capabilities for the AMD Opteron™ processor while maintaining compatibility with the Intel® Pentium® Xeon™ processors used in the iHawk™ Series 860 systems.

RedHawk Linux is a real-time version of the industry-standard open source Linux operating system. The RedHawk Linux kernel is based on the kernel maintained by Linus Torvalds and utilizes the 2.6.0 version. For a complete description of the operating system and its features, see "Information on Version 2.0" on page 19.

The AMD Opteron version of RedHawk Linux is a full 64-bit operating system that executes both 32-bit and 64-bit applications in native mode on the processor.

The RedHawk Linux operating system is included with each Concurrent iHawk system. The iHawk Series 860 features from one to eight Intel Pentium Xeon processors in a single rackmount or tower enclosure. The iHawk 860G is a one or two processor, high-performance AGP/PCI-based platform for real-time imaging applications. The iHawk systems offer leading-edge integrated circuit and packaging technology. iHawks are true symmetric multiprocessors (SMPs) that run a single copy of the RedHawk Linux real-time operating system. All CPUs in a system are linked by a high-speed front-side processor bus and have direct, cache-coherent access to all of main memory.

RedHawk Linux supports Concurrent's powerful set of NightStar™ development tools on iHawk 860 systems. Users can schedule, monitor, debug and analyze their real-time applications on iHawk multiprocessor systems or remotely from a desktop PC or laptop. Each tool runs on the iHawk target system non-intrusively, thus preserving the deterministic characteristics of the real-time application.

1.1 Product Media

The following CDs are delivered with RedHawk Linux Version 2.0:

Label on Disc		Qty
Red Hat Enterprise Linux Version 3		9
Red Hat Update CD for RedHawk Linux	LXUPDATE Version 2.0 WA9086-J	1
RedHawk Linux OS	WA9086-J Version 2.0	1
RedHawk Linux Documentation CD	LXDOCKIT Version 2.0 WA9086-J	1

The following optional CDs may also be provided:

Label on Disc		Qty
RedHawk Linux NightStar Tools	WU1220-LDV NightStar Tools Version <i>x.x</i> (where <i>x.x</i> = latest version)	1
RedHawk Linux Frequency-Based Scheduler	WU1210-JL Version 2.0	1

Warning

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

1.2 System Updates

As RedHawk Linux updates are issued, they are made available for downloading from Concurrent's RedHawk Linux update website. Instructions for downloading updates from the website are included with the installation and update instructions in this Release Notes and are explained in the section "Software Updates and Support" on page 34.

NOTE

It is not recommended that the user download Red Hat updates.

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

2.0. Documentation

2.1 RedHawk Linux Documentation

Table 2-1 lists the documentation provided by Concurrent as PDF files on the RedHawk Linux documentation CD.

NightStar Tools documentation is included as PDF files on the RedHawk Linux NightStar Tools CD.

Updates to Concurrent documentation as PDF files can be downloaded from the RedHawk Updates web site. Refer to “Software Updates and Support” on page 34 for more information. Documentation is also available from Concurrent’s web site at <http://www.ccur.com>.

Printed copies of Concurrent documentation can be ordered by contacting the Concurrent Software Support Center. The toll-free number for calls within the continental United States is 1-800-245-6453. For calls outside the continental United States, the number is 1-954-283-1822.

Table 2-1 RedHawk Linux Documentation

Manual Name	Document Number
<i>RedHawk Linux Online Documentation Roadmap</i>	0898002-400
<i>RedHawk Linux OS Version 2.0 Release Notes</i>	0898003-2.0
<i>RedHawk Linux User’s Guide</i>	0898004-400
<i>Frequency-Based Scheduler (FBS) User’s Guide</i>	0898005-200
<i>Real-Time Clock & Interrupt Module (RCIM) PCI Form Factor User’s Guide</i>	0898007-210

2.2 Vendor Documentation

Documentation provided by Red Hat is included on the Red Hat CDs delivered with RedHawk Linux.

Commercial off-the-shelf (COTS) documentation applicable to your iHawk Series 860 is included with the system. Contact your local Dell sales office to purchase additional copies.

3.0. Prerequisites

3.1 Software

- Red Hat Enterprise Linux Version 3
- RedHawk Linux operating system Version 2.0

3.2 Hardware

- any iHawk Series 860 system or AMD Opteron processor
- Real-Time Clock and Interrupt Module (RCIM)

4.0. Version 2.0 Installation Procedures

The steps provided in this section apply when installing RedHawk Linux under the following conditions:

- installing RedHawk Linux Version 2.0 on an AMD Opteron system
- initially installing or regenerating an iHawk system (iHawk systems are normally shipped fully installed)

Review the following points before installing the software:

- It is recommended that the installer first read through the entire installation procedure once before actually performing the installation.
- It is **strongly recommended** that installation and updates of software be performed in single-user mode to guard against the possibility of corrupting a running application.
- The installation media in use is accessed at various times and must be kept in the drive during the installation until you are instructed to install different media.
- Prior to installation, remove any floppy media that may have been left in the machine.

Before you start:

- If you plan to configure networking using fixed numbers (recommended), you will need to have the following information for your system before beginning the installation:
 - IP address
 - netmask address
 - gateway address
 - primary and secondary DNS addresses
- 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.

4.1 Installing Red Hat Software

The Red Hat Enterprise Linux software must be installed on the system before RedHawk Linux. Complete the following steps to install and configure Red Hat. Refer to Online Help if needed.

1. Power on the system, if needed, to open the CD-ROM drawer.
2. Insert Red Hat Enterprise Linux Installation Disc 1 of 4 in the CD-ROM drive and power down the system.
3. Power on the system with the installation disc installed in order to boot from the disc. If you wish to check the Red Hat media, type `linux mediacheck` at the `boot :` prompt. To start the installation, press **ENTER**.

At the **Welcome** screen, click **Next**.

- At the **Language Selection** screen:
Choose English language. Click **Next**.
- At the **Keyboard** screen:
Accept keyboard defaults. Click **Next**.

- At the Mouse Configuration screen:
Select the appropriate mouse model. **Note:** Click the “Emulate 3 buttons” box at bottom of page if using a 2 button mouse. Click Next.
- At the Disk Partitioning Setup screen:
Choose “Manually Partition with Disk Druid” and click Next to partition the disk manually and create the partition sizes as shown in Table 4-1 below. Note that this table shows recommended partition sizes for a 36 GB disk or larger.

Note

All partitions on SCSI drives must be placed on **/dev/sda**; all partitions on IDE drives must be placed on **/dev/hda**. If the system has multiple drives, you must un-select all other drives as shown in the “Allowable Drives:” section of the Add Partition setup menu.

- At the Partitioning screen, first delete all existing partitions by highlighting each line and clicking DELETE. This must be done one partition at a time. To add partitions, highlight the disk drive and click NEW. Select the Force to be primary partition box at the bottom of each partition setup page. Click OK to continue to set up the next partition.

Table 4-1 Recommended Disk Partition Sizes

Mount Point	Type	Filesystem Type	Partition Size
/boot	primary	ext3	100 MB
/	primary	ext3	16384 MB (16 GB)
swap (See Note below)	primary	swap	Size of system’s main memory plus 5 percent additional. (See Table 11-1 on page 33 for additional information)
/home	primary	ext3	Fill to maximum allowable size
Note: Select File System Type on the Add Partition screen to select the swap file system. Enter the recommended size.			

- Click Next when finished with partitioning.
- At the Boot Loader Configuration screen:
Accept GRUB boot loader defaults. Enter a GRUB password if desired (not recommended). Click Next.

- At the **Network Configuration** screen:
Configure networking as appropriate (fixed network numbers are recommended - you will need address information if using fixed numbers). Select the network device and click **Edit** to configure networking. Refer to Online Help for additional instructions. Click **Next** after completion.
 - At the **Firewall** screen:
Choose “No firewall”, then click **Next**.
 - At the **Additional Language Support** screen:
Choose any additional languages as necessary. Click **Next** at completion.
 - At the **Time Zone Selection** screen:
Select the appropriate time zone, then click **Next**.
 - At the **Set Root Password** screen:
Enter root password and confirmation, then click **Next**.
 - At the **Package Defaults** screen:
Select “Customize the set of packages to be installed”, then click **Next**.
 - At the **Package Group Selections** screen:
Scroll down to the bottom of the list and select “Everything.” Click **Next**.
 - At the **About to install** page, click **Next**.

At this point in the installation, the disk will format and packages on Red Hat Installation disc 1 will be installed. When disc 1 is completed, you will be prompted to insert disc 2. Remove disc 1 and insert disc 2, then click **OK**. Take the same action when prompted for disks 3 and 4 and finally disc 1 again.
 - At the **Graphical Interface (X) Configuration** screen:
Accept defaults, click **Next**.
 - At the **Monitor Configuration** screen:
Accept monitor defaults, click **Next**.
 - At the **Customize Graphical Configuration** screen:
Provide the Color Depth and Screen Resolution values for your monitor, then click **Next**.
4. Remove any media from respective drives to prevent booting the installation disc again and store appropriately.
 5. At the next window prompt, press **Exit** to reboot the system. Verify that Red Hat boots up properly.
 6. Once the system boots up successfully, proceed as follows:
 - At the **Welcome** screen:
Click **Next** to continue.

- At the **License Agreement** screen:
Click “Yes, I agree to the License Agreement”, then click **Next**.
 - At the **Date and Time** screen:
Set the appropriate date and time. Enable NTP if desired. Click **Next** to continue.
 - At the **User Account** screen:
Enter a username and password for each account you wish to set up at this time (this can be done after the full installation). Click **Next** to continue.
 - If your system has a sound card installed, the **Sound Card** screen will appear. At the **Sound Card** screen:
Click **Play test sound** if desired. Click **Next**.
 - At the **Red Hat Network** screen:
Select “No, I do not want to register my system.” Click **Next** to continue.
 - At the **Install Additional CDs** screen:
Three choices appear for installing additional Red Hat CDs. If you wish to install any of them now, click on the **Install** button for the desired CD and follow the instructions that display. Click **Next** to continue.
 - At the **Finished Setup!** screen:
Click **Next** to continue.
7. Initial Red Hat installation is now complete. The login window should appear at this point. Remove any media from respective drives. Continue with the steps in section 4.2 to install Red Hat software updates.

4.2 Installing Red Hat Software Updates

The disc labeled “Red Hat Enterprise Linux Update CD for RedHawk Linux Version 2.0” contains the latest updates released by Red Hat and validated against RedHawk Linux. These updates are important to the proper operation of RedHawk Linux.

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

1. Log in as root and take the system down to single-user mode:
 - a. Right click on the desktop and select **New Terminal**.
 - b. At the system prompt, type **init 1**.
2. Locate the disc labeled “Red Hat Enterprise Linux Update CD for RedHawk Linux Version 2.0” and insert it into the CD-ROM drive.
3. If the cdrom device does not mount automatically, execute the following command:

```
mount /mnt/cdrom
```

- Execute the following commands to install the update software:

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

Follow the on-screen instructions until the installation script completes. This will take some time to complete.

- Upon completion of software installation, execute the following commands:

```
cd /
umount /mnt/cdrom
eject
```

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

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

Execute the following:

```
rpm -e kernel-smp
```

- Exit single-user mode (^D), reboot and verify that the new kernel boots. Continue with the steps in the next section to install RedHawk Linux.

4.3 Installing RedHawk Linux

After installing Red Hat and its updates, perform the following steps to install RedHawk Linux:

- Log in as root and take the system down to single-user mode:
 - Right click on the desktop and select **New Terminal**.
 - At the system prompt, type **init 1**.
- Locate the disc labeled “RedHawk Linux OS Version 2.0” and insert it into the CD-ROM drive.

- If the cdrom device does not mount automatically, execute the following command:

```
mount /mnt/cdrom
```

- Execute the following commands to install RedHawk Linux:

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

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

- Upon completion of software installation, execute the following commands:

```
cd /
umount /mnt/cdrom
eject
```

- Remove the disc from the CD-ROM drive and store. Exit single-user mode (^D). Check for RedHawk updates as outlined in the following section.

4.4 Installing RedHawk Updates

If there are any updates to this release available for downloading from Concurrent's web site, they should be installed now. Follow these steps:

1. In multi-user mode, access the redhawk.ccur.com web site.
2. Log onto the site using the login and password provided by Concurrent.
3. A list of packages that you have not yet downloaded will display. Click on the packages to download them and save them in a suitable location (e.g. `/root/updates`).
4. Return to single-user mode as root.
5. To install the newly downloaded packages, change directory to the location of the updates and issue the following command:

```
rpm -Fvh *.rpm
```

The command should complete successfully, though it may take several minutes depending on the number of updates being installed.

6. Exit single-user mode (^D). Reboot the system* choosing the new RedHawk kernel and verify that the system boots properly. Continue with section 4.5 if an RCIM is to be installed or section 4.6 to install FBS.

* **Note:** If this system is a model 868 8-way iHawk, disable hyper-threading through the BIOS, or on the GRUB boot command line by specifying "noht".

4.5 Installing the RCIM

The following serves as an installation checklist for installing a PCI RCIM in an iHawk system. If an RCIM is not installed in your iHawk system, it should be installed now. See the *Real-Time Clock and Interrupt Module (RCIM) PCI Form Factor User's Guide* for complete details.

4.5.1 Quick Hardware Installation Checklist

1. Before installing the RCIM, decide what mode the RCIM will run in (see "RCIM Connection Modes" on page 32) and if you will be using the RCIM to accept external interrupts.
2. Verify that all pre-requisite software has been installed.
3. 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.
4. Power down the system and remove all power cords.
5. Open the case and identify a PCI slot for installation. (Refer to the chart on the case for help in identifying card slot types.)

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

6. Insert the RCIM into the selected PCI slot and lock it down.

7. Attach synchronization input cable if required (see the RCIM manual for complete details).
8. Close the case and reconnect all power cords.
9. Power up the system and verify the system boots up properly. Refer to the *Real-Time Clock and Interrupt Module (RCIM) PCI Form Factor User's Guide* for configuration options.

4.5.2 Check 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 board revision 1.
This is a standalone (isolated) rcim.
```

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

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

You should see an output similar to the following:

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

4.6 Installing Frequency-Based Scheduler (FBS) Software

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

1. With RedHawk Linux Version 2.0 running, log in as root and take the system down to single-user mode:
 - a. Right click on the desktop and select **New Terminal**.
 - b. At the system prompt, type **init 1**.
2. Locate the disc labeled “RedHawk Linux Frequency-Based Scheduler Version 2.0” and insert it into the CD-ROM drive.
3. To mount cdrom device, execute the following command:

```
mount /mnt/cdrom
```

4. To install, execute the following commands:

```
cd /mnt/cdrom
./install-fbs
```

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

5. Upon completion of software installation, execute the following commands:

```
cd /  
umount /mnt/cdrom  
eject
```

6. Remove the disc from the CD-ROM drive and store. Exit single-user mode (^D). Check for FBS updates as outlined in the following section.

4.7 Installing FBS Updates

If there are any updates to this FBS release available for downloading from Concurrent's web site, they should be installed now. Follow these steps:

1. In multi-user mode, access the redhawk.ccur.com web site.
2. Log onto the site using the login and password provided by Concurrent.
3. A list of packages that you have not yet downloaded will display. Click on the packages to download them and save them in a suitable location (e.g. `/root/updates`).
4. Return to single-user mode as root.
5. To install the newly downloaded packages, change directory to the location of the updates and issue the following command:

```
rpm -Fvh *.rpm
```

The command should complete successfully, though it may take several minutes depending on the number of updates being installed.

6. Exit single-user mode (^D), reboot the system choosing the new RedHawk kernel and verify that the system boots properly.

4.8 Installing NightStar Tools

NightStar Tools are supported on the iHawk 860 systems. If you will be installing the NightStar toolset on an iHawk system, do so now. Refer to the *RedHawk NightStar Tools Release Notes* for installation instructions.

5.0. NVIDIA Graphics Display

This section addresses graphics display configurations on iHawk 860 graphics systems only.

The NVIDIA Linux Display Driver (version 1.0-4496) is included in the RedHawk Linux installation and built automatically for all pre-built and custom RedHawk Linux kernels.

Note that if you install an NVIDIA graphics card into your system after RedHawk Linux has been installed, you may need to install the user module, `ccur-nvidia-glx`, at that time. The RPM for this module can be found on the RedHawk Linux installation CD in the `nvidia` directory.

There are several options for configuring two or more displays. The upper limit on the number of displays that can be driven effectively is application dependent. Theoretically, you can configure as many displays as you have video adapters.

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

The following multiple display modes are available:

- One desktop spanned across two monitors.
NVIDIA calls this TwinView. In this mode, the workspace size is twice as large as with a single monitor display. Windows can be dragged across displays, which can be positioned either vertically or horizontally. Switching between workspaces affects both displays simultaneously so it's easy to switch between data views that are linked together, for example, during demos and presentations. Multiple workspaces can be configured.
- Separate desktops.
This mode provides two or more separate desktops that act independently, for example a NightView debugger on one monitor and an application running on another. Because the monitors work independently, displays are not synchronized as in TwinView mode and you cannot drag windows across monitors; however, it allows you more combinations for viewing data.

The mode in use depends upon the settings in the X server configuration file when the X server is booted. You can easily switch between modes by maintaining more than one configuration file and copying the one you want to use to `/etc/X11/XF86Config` before starting the X server.

Along with the original `XF86Config` configuration file located in `/etc/X11`, there are additional sample configuration files. These samples illustrate configurations for TwinView, dual desktop and quad desktop. Use these as a reference when creating configuration files for use at your site.

To configure and boot your system for multiple displays, follow the steps in the next section.

5.1 Configuring Multiple Display Capability

If you think you will utilize more than one display mode, create a configuration file for each mode you intend to use. Configuration settings are read from `/etc/X11/XF86Config` when the X server is booted. Have the desired configuration file saved to that path before booting the X server.

Note

Parameters required for the multiple display modes are discussed below, but other fields in the configuration file may need to be modified depending upon other factors at your site. For complete information about the `XF86Config` file, see the `XF86Config(5x)` man page. Additionally, you can download the README file from the www.nvidia.com web site.

You will need to know the following information for each monitor (consult the documentation provided with the monitors):

- the range of possible refresh rates
- supported resolutions
- supported color depths

Configuration for One Desktop Across Two Monitors

To create a virtual desktop that displays across two monitors (for example, NVIDIA's TwinView) you need to provide appropriate settings in the "Screen" section of the configuration file. The display mode and monitor characteristics are defined as "Options." Refresh rates and resolutions for each monitor and the location (orientation) of the second monitor used as the extended desktop are supplied. "TwinViewOrientation" can be "LeftOf", "RightOf", "Above" or "Below" to represent where the second monitor will be positioned relative to the first.

Following is an example of the "Screen" section of the configuration file with possible settings for TwinView mode.

```
Section "Screen"
    Identifier "Screen0"
    Device "Videocard0"
    Monitor "Monitor0"
    DefaultDepth 24
    Option "TwinView"
    Option "SecondMonitorHorizSync" "30-97"
    Option "SecondMonitorVertRefresh" "50-180"
    Option "MetaModes" "1280x1024,1280x1024"
    Option "TwinViewOrientation" "LeftOf"
    SubSection "Display"
        Depth 16
        Modes "1280x1024" "1024x768" "800x600" "640x480"
    EndSubSection
    SubSection "Display"
        Depth 24
        Modes "1280x1024" "1024x768" "800x600" "640x480"
    EndSubSection
EndSection
```

Configuration for Separate Desktops

For this mode, the Device, Screen, Monitor and ServerLayout sections in the configuration file are utilized. The number of displays available is limited only by hardware constraints. The examples below are for a dual display. For additional displays, add additional sections with appropriate designations (see **XF86Config.quaddesktop** for examples). Follow the steps below to modify the configuration file.

1. First, update the ServerLayout section to reflect the multiple display configuration. Each screen to be used will be listed here and defined in Screen sections later in the file.

```
Section "ServerLayout"
    Identifier "nvidia dualdesktop"
    Screen 0 "Screen0" 0 0
    Screen 1 "Screen1" LeftOf "Screen0"
    InputDevice "Mouse0" "CorePointer"
    InputDevice "Mouse1" "SendCoreEvents"
    InputDevice "Keyboard0" "CoreKeyboard"
EndSection
```

2. Then, create separate Device sections. In each section, provide the BusID (*bus:device:function*) of the graphics card associated with each display. This information displays when you `cat /proc/pci.` In the case of a dual head adapter, specify separate sections (one for each head) using the same BusID. **Caution:** BusIDs can change when PCI cards are added or removed.

The driver in each section in the example below is identified as "nvidia." A separate screen is assigned in each section.

```
Section "Device"
    Identifier "nvidia0"
    Driver "nvidia"
    BusID "PCI:1:0:0"
    Screen 0
EndSection
```

```
Section "Device"
    Identifier "nvidia1"
    Driver "nvidia"
    BusId "PCI:1:0:0"
    Screen 1
EndSection
```

- Next, create separate Screen sections, each using one of the devices identified in the Device sections; for example:

```
Section "Screen"
    Identifier "Screen0"
    Device      "nvidia0"
    Monitor     "Monitor0"
    DefaultDepth 24
    Subsection "Display"
        Depth 24
        Modes "1600x1200" "1024x768" "800x600" "640x480"
    EndSubsection
EndSection
```

```
Section "Screen"
    Identifier "Screen1"
    Device      "nvidia1"
    Monitor     "Monitor1"
    DefaultDepth 24
    Subsection "Display"
        Depth 24
        Modes "1600x1200" "1024x768" "800x600" "640x480"
    EndSubsection
EndSection
```

- Finally, create separate Monitor sections, each describing the characteristics of one of the monitors to be used; for example:

```
Section "Monitor"
    Identifier "Monitor0"
    VendorName "Viewsonic"
    ModelName  "G90f"
    DisplaySize 370 270
    HorizSync 30.0 - 97.0
    VertRefresh 50.0 - 180.0
    Option      "dpms"
EndSection
```

5.2 Connecting the Monitors and Booting the System

The graphics adapter card identifies multiple outputs as head 1, head 2, etc. With multiple monitors, the console output goes to head 1 during boot or when X is not running. The monitors on subsequent heads only display when X is running.

- Connect monitor 1 to head 1 of the adapter, monitor 2 to head 2 of the adapter, etc.
- To allow a recovery if there is a problem during the installation, set the system's default run level to enable the system to boot to console instead of starting up X.

- Edit **/etc/inittab**:

```
vi /etc/inittab
```

- Locate the following line:

```
id:5:initdefault
```

- c. Change the 5 to 3, save and exit.
 - d. Reboot the system with `init 3`.
3. If you will be using separate desktops and wish to use **gnome**, edit `/etc/X11/gdm/gdm.conf` and uncomment the following line:

```
1=Standard
```

Note that **kde** works without modifications.

4. Make a backup copy of the original configuration file, then copy the file you wish to use to `/etc/X11/XF86Config`; for example:

```
cp XF86Config XF86Config.orig
cp XF86Config.mydual XF86Config
```

5. Startup X with `init 5`.
6. To test the success of the installation, run the following command:

```
startx
```

If successful, you will see the NVIDIA logo and then the desktop appears. Reset the system's default run level back to run level 5 (X11 startup mode).

- a. Using **vi** to edit `/etc/inittab`, locate the line reading:

```
id:3:initdefault
```

Change the 3 to 5, save and exit.

- b. Execute the command `init 6`.

If unsuccessful, the problem may be due to a misconfigured `XF86Config` file. Review the `/var/log/XFree86*.log` files along with the NVIDIA README document to debug your `/etc/X11/XF86Config` file.

7. If you want to test the 3D performance of your card, perform the following:

```
glxgears
```

This should bring up three gears rotating at high speed with frames per second displayed in the associated window.

If the setup is incorrect, you may get the following on-screen messages:

```
Xlib: extension "GLX" missing on display ":0.0".
Error: couldn't get an RGB, Double-buffered visual
```

Refer to the NVIDIA README document for troubleshooting tips.

6.0. Additional Considerations

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

6.1 Compiler Requirements

If you plan to build a kernel from source residing on your iHawk system on a different system, that system must be using **gcc-2.96-98** or later.

6.2 NightView Accessibility

A possible result of upgrading Red Hat software during a RedHawk Linux upgrade is a break in the symbolic link for the NightView debugger (**nview**) making it unavailable. The RedHawk Linux install script will re-establish this link and no further action should be necessary. If you encounter a problem accessing **nview**, ensure that a symbolic link is established linking **/usr/bin/nview** to the current NightView installation; e.g, **/usr/opt/ccur-NightView-5.6-913/usr/bin/nview**.

7.0. Information on Version 2.0

7.1 General

RedHawk Linux is a real-time version of the industry-standard open source Linux operating system. The RedHawk Linux Version 2.0 kernel is based on the kernel maintained by Linus Torvalds and utilizes the 2.6.0 version. Concurrent used Red Hat Enterprise Linux and supplied an alternate multithreaded, fully-preemptible kernel with low latency enhancements for real-time processing. Additional features include:

- features in the real-time operating systems that Concurrent has been deploying for many years
- performance improvements developed by Concurrent to improve worst-case process dispatch latency or to improve determinism in process execution

These kernel enhancements provide the support needed for developing complex real-time applications. Additional user-level libraries in RedHawk Linux provide interfaces to some of the real-time features present in the Red Hat Linux kernel. Descriptions of the real-time features in RedHawk Linux are provided in the section “Real-Time Features in RedHawk Linux” later in this section.

Except for the kernel, all Red Hat components operate in their standard fashion. These include Linux utilities, libraries, compilers, tools and installer unmodified from Red Hat.

Linux conforms to many of the interface standards that are defined by POSIX. However, Linux does not completely conform to these standards. Red Hat Linux has the same level of POSIX conformance as other Linux distributions that are based on the 2.6 series of kernels. Linux on the Intel x86 architecture has defined a binary standard of its own which allows shrink-wrapped applications that are designed to run on the Linux/Intel x86 platform to run on Concurrent’s iHawk platform.

7.2 RedHawk Linux Kernels

There are three flavors of RedHawk Linux kernels. The system administrator can select which version of the RedHawk Linux kernel is loaded via the GRUB boot loader. The three flavors of pre-built RedHawk kernels are:

Kernel Name	Kernel Description
vmlinuz-2.6.0-test11-bk3-RedHawk-2.0-trace	Kernel with trace points but no debug checks (Default kernel)
vmlinuz-2.6.0-test11-bk3-RedHawk-2.0-debug	Kernel with both debug checks and kernel trace points
vmlinuz-2.6.0-test11-bk3-RedHawk-2.0	Kernel with no trace points and no debug checks

The default RedHawk Linux trace kernel has been built with kernel trace points enabled. The kernel trace points allow the NightTrace™ tool to trace kernel activity.

The debug kernel has been built with both debugging checks and kernel trace points enabled. The debugging checks are extra sanity checks that allow kernel problems to be detected earlier than they might otherwise be detected. However, these checks do produce extra overhead. If you are measuring performance metrics on your system, this activity would be best performed using a non-debug version of the kernel.

The RedHawk installation script installs the kernels appropriate to the architecture on which it is being installed (i386 or x86_64).

7.3 Real-Time Features in RedHawk Linux

Table 7-1 contains a brief description of the real-time features included in the RedHawk Linux operating system.

Table 7-1 Real-Time Features in RedHawk Linux

Feature	Description
Processor Shielding	Enhancements to the scheduler and the <code>/proc</code> file system allow the user to dynamically shield a set of processors from interrupts and other background activity that can interfere with determinism in an application's execution time. Shielded processors provide the best possible environment for guaranteeing that a real-time application can meet its execution time deadlines. See the <i>RedHawk Linux User's Guide</i> and the <code>shield(5)</code> man page for complete details.
Processor Affinity	In a real-time application where multiple processes execute on multiple CPUs, it is desirable to have explicit control over the CPU assignments of all processes in the system. This capability is supplied by Concurrent through the <code>mpadvise(3)</code> library routine and the <code>run(1)</code> command. See the <i>RedHawk Linux User's Guide</i> and the man pages for more details.
Real-Time Scheduler	The real-time scheduler provides fixed-length context switch times regardless of how many processes are active in the system. It also provides a true real-time scheduling class that operates on a symmetric multiprocessor.
Low Latency Enhancements	In order to protect shared data structures used by the kernel, the kernel protects code paths that access these shared data structures with spin locks and semaphores. The locking of a spin lock requires that preemption, and sometimes interrupts, be disabled while the spin lock is held. A study was made which identified the worst-case preemption off times. The low latency enhancements applied to RedHawk Linux modify the algorithms in the identified worst-case preemption off scenarios to provide better interrupt response times.

Table 7-1 Real-Time Features in RedHawk Linux (Continued)

Feature	Description
Kernel Preemption	The ability for a high priority process to preempt a lower priority process that is currently executing inside the kernel is provided in RedHawk Linux. Under standard Linux the lower priority process would continue running until it exited from the kernel, creating longer worst case process dispatch latency. Data structure protection mechanisms are built into the kernel to support symmetric multiprocessing.
User-level Preemption Control	When an application has multiple processes that can run on multiple CPUs and those processes operate on data shared between them, access to the shared data must be protected to prevent corruption from simultaneous access by more than one process. The most efficient mechanism for protecting shared data is a spin lock; however, spin locks cannot be effectively used by an application if there is a possibility that the application can be preempted while holding the spin lock. To remain effective, RedHawk provides a mechanism that allows the application to quickly disable preemption. See the <i>RedHawk Linux User's Guide</i> and the <code>resched_cntl(2)</code> man page for more information about user-level preemption control.
User-level Spin Locks	RedHawk Linux busy-wait mutual exclusion tools include a low-overhead busy-wait mutual exclusion variable (a spin lock) and a corresponding set of macros that allow you to initialize, lock, unlock and query spin locks. To be effective, user-level spin locks must be used with user-level preemption control. Refer to the <i>RedHawk Linux User's Guide</i> for details.
Fast Block/Wake Services	Many real-time applications are composed of multiple cooperating processes. These applications require efficient means for doing inter-process synchronization. The fast block/wake services developed by Concurrent allow a process to quickly suspend itself awaiting a wakeup notification from another cooperating process. See the <i>RedHawk Linux User's Guide</i> and the <code>postwait(2)</code> and <code>server_block(3)</code> man pages for more details.
Hyper-threading	Hyper-threading is a feature of the Intel Pentium Xeon processor on iHawk systems that allows for a single physical processor to appear to the operating system as two logical processors. Two program counters run simultaneously within each CPU chip so that in effect, each chip is a dual-CPU SMP. With hyper-threading, physical CPUs can run multiple tasks "in parallel" by utilizing fast hardware-based context-switching between the two register sets upon things like cache-misses or special instructions. Refer to the <i>RedHawk Linux User's Guide</i> for more information on how to effectively use this feature in a real-time environment on iHawk systems.

Table 7-1 Real-Time Features in RedHawk Linux (Continued)

Feature	Description
usermap(3) and /proc mmap(2)	The usermap(3) library routine has been added to RedHawk Linux, which lets a process map portions of another process' address space into its own address space utilizing the /proc filesystem mmap(2) . Thus, monitoring and modifying other executing programs becomes simple CPU reads and writes within the application's own address space, without incurring the overhead of /proc filesystem read(2) and write(2) system service calls. Refer to the <i>RedHawk Linux User's Guide</i> for more information.
POSIX Timers (high-resolution)	Support for high-resolution POSIX clocks and timers is included in RedHawk Linux. There are four system-wide POSIX clocks that can be used for such purposes as time stamping or measuring the length of code segments. POSIX timers allow applications to use relative or absolute time based on a high resolution clock and to schedule events on a one-shot or periodic basis. Applications can create multiple timers for each process. In addition, high-resolution sleep mechanisms are provided which can be used to put a process to sleep for a very short time quantum and specify which clock should be used for measuring the duration of the sleep. See the <i>RedHawk Linux User's Guide</i> for more information on POSIX timers.
POSIX Message Queues	POSIX message passing facilities based on IEEE Standard 1003.1b-1993 are included in RedHawk Linux, implemented as a file system. POSIX message queue library routines allow a process to create, open, query and destroy a message queue, send and receive messages from a message queue, associate a priority with a message to be sent, and request asynchronous notification when a message arrives. POSIX message queues operate independently of System V IPC messaging, which is also available under RedHawk Linux. See the <i>RedHawk Linux User's Guide</i> for more information.
XFS Journaling File System	Support for the XFS journaling file system from SGI is included in RedHawk Linux. Journaling file systems use a journal (log) to record transactions. This drastically cuts the complexity of a file system check, reducing recovery time. XFS is a multithreaded, 64-bit file system capable of handling large files and large filesystems, extended attributes, variable block sizes, is extent based and makes extensive use of Btrees (directories, extents, free space) to aid both performance and scalability. Both user and group quotas are supported. See the <i>RedHawk Linux User's Guide</i> for more information.
High Resolution Process Accounting	In the standard Linux kernel, the system accounts for a process' CPU execution times using a very coarse-grained mechanism. This means that the amount of CPU time charged to a particular process can be very inaccurate. The high resolution process accounting facility provides a mechanism for very accurate CPU execution time accounting, allowing better performance monitoring of applications. This facility is used by standard Linux CPU accounting services and the Performance Monitor. See the <i>RedHawk Linux User's Guide</i> for more details.

Table 7-1 Real-Time Features in RedHawk Linux (Continued)

Feature	Description
/ proc Modifications	Modifications were made to the process address space support in /proc to allow a privileged process to read or write the values in another process' address space. This is for support of the NightProbe data monitoring tool and for the NightView debugger.
RCIM Driver	<p>A driver was added for support of the Real-Time Clock and Interrupt Module (RCIM). This multi-purpose PCI card has the following functionality:</p> <ul style="list-style-type: none"> - allows the connection of four external device interrupts - four real time clocks that can interrupt the system - four programmable interrupt generators which allow generation of an interrupt from an application program <p>These functions can all generate local interrupts on the system where the RCIM card is installed. Multiple RedHawk Linux systems can be chained together, allowing up to eight of the local interrupts to be distributed to other RCIM-connected systems. This allows one timer, or one external interrupt or one application program to interrupt multiple RedHawk systems almost simultaneously to create synchronized actions. In addition, the RCIM contains a synchronized high-resolution clock so that multiple systems can share a common time base. See the <i>Real-Time Clock and Interrupt Module (RCIM) PCI Form Factor User's Guide</i> for additional information.</p>
Frequency-Based Scheduler	The Frequency-Based Scheduler (FBS) provides a mechanism for scheduling applications that run according to a predetermined cyclic execution pattern. The FBS also provides a very fast mechanism for waking a process when it is time for that process to execute. In addition, the FBS allows the performance of cyclical applications to be tracked, with various options available to the programmer when deadlines are not being met. The FBS is the kernel mechanism that underlies the NightSim GUI for scheduling cyclical applications. See the <i>NightSim User's Guide</i> and the <i>Frequency-Based Scheduler (FBS) User's Guide</i> for additional information.
Capabilities Support	Pluggable Authentication Modules (PAM) provides a mechanism to assign privilege to users. This allows a non-root user to run applications that require privileges that only root would normally be allowed. For example, the ability to lock pages in memory is provided by one predefined privilege that can be assigned to individual users or groups. See the <i>RedHawk Linux User's Guide</i> and the pam_capability(8) and capability.conf(5) man pages for more information about capabilities support.
Kernel Trace Facility	Support was added to RedHawk Linux to allow kernel activity to be traced. This includes mechanisms for inserting and enabling kernel trace points, reading trace memory buffers from the kernel, and managing trace buffers. The kernel trace facility is used by the NightTrace tool.

Table 7-1 Real-Time Features in RedHawk Linux (Continued)

Feature	Description
ptrace Extensions	<p>The ptrace debugging interface has been extended to support the capabilities of the NightView debugger. Features added include:</p> <ul style="list-style-type: none"> - the capability for a debugger process to read and write memory in a process not currently in the stopped state - the capability for a debugger to trace only a subset of the signals in a process being debugged - the capability for a debugger to efficiently resume execution at a new address within a process being debugged - the capability for a debugger process to automatically attach to all children of a process being debugged
64-bit Capabilities	<p>RedHawk Linux operates on the 64-bit AMD Opteron processor. It is a full 64-bit operating system that executes both 32-bit and 64-bit applications in native mode.</p>
Kernel Debuggers	<p>The kgdb debugger is supported through an open source patch. It allows the RedHawk Linux kernel to be debugged with gdb as if it were a user application. gdb runs on a separate system, communicating with the kernel being debugged through the console's serial port. For complete information about kgdb, see http://kgdb.sourceforge.net. Currently, kgdb is supported on iHawk systems only.</p>
Kernel Core Dumps/ Crash Analysis	<p>This open source patch provides the support for dumping physical memory contents to a file as well as support for utilities that do a post mortem analysis of a kernel core dump. See the lcrash(1) man page and the section entitled "Using the Linux Kernel Crash Dump (LKCD)" on page 30 of these notes or the <i>RedHawk Linux User's Guide</i> for additional information on LKCD. Currently, LKCD is supported on iHawk systems only.</p>

7.4 NightStar Development Tools

On iHawk systems, users can schedule, monitor, debug and analyze run time behavior of their real-time applications and the RedHawk Linux operating system kernel using the following NightStar tools:

- NightTrace Event Analyzer
- NightSim Periodic Scheduler
- NightProbe Data Monitor
- NightView Debugger

NightStar's non-intrusive operation ensures real-time determinism.

Refer to the *RedHawk Linux NightStar Tools Release Notes*, Pubs No. 0898008-*x.x* (where *x.x* = latest version) for additional information.

8.0. Changes in this Release

This section describes enhancements and other changes introduced in Version 2.0.

8.1 Enhancements

8.1.1 Support for 64-bit Processing

RedHawk Linux Version 2.0 includes support for the 64-bit AMD Opteron processor. RedHawk Linux is a full 64-bit operating system that executes both 32-bit and 64-bit applications in native mode on the Opteron processor. Existing 32-bit applications run natively with no “emulation mode” to degrade performance.

Complete information about the AMD Opteron processor and 64-bit processing can be found at www.amd.com.

8.1.2 Dynamic Creation of Custom Kernel Trace Events

RedHawk Linux now provides the ability to dynamically create kernel trace events. Previously, there existed only a single “CUSTOM” event that could be used to arbitrarily add events to device drivers, applications and kernel debug code. Now a range of event IDs can be used for dynamic custom kernel events. In the code to be traced, requests for allocation of a unique ID and a name for the tracepoint can be provided. The NightTrace analyzer tool includes these dynamic events in its display.

Creating dynamic kernel trace events is described in the “Device Drivers” chapter of the *RedHawk Linux User’s Guide*.

8.2 Modifications

8.2.1 High Resolution Process Accounting

Previously, per-process time measurements using the high resolution time stamp counter (TSC) were performed by the High Resolution Process Timing Facility. This facility has been redesigned and renamed to “high resolution process accounting.”

This facility provides measurements of system, user, interrupted system and interrupted user time for each process. On kernels configured with this facility, services such as **ps (1)**, **top (1)**, **times (2)** and **getrusage (2)** use it for standard CPU accounting instead of the local timer interrupt. On kernels not so configured, the local timer interrupt is used for CPU accounting as was the case previously. Configuration is accomplished through the CONFIG_HRACCT kernel tunable accessible through the **General Setup** selection of the Linux Kernel Configuration menu.

There are library routines and **/proc** files available for obtaining process execution time measurements and TSC frequencies. Library routines are also provided for converting measurements to different formats. Refer to the **hracct (3)** and **hracct (7)** man pages or the *RedHawk Linux User’s Guide* for more information about high resolution process accounting.

8.2.2 Frequency-Based Scheduler and Performance Monitor

Along with the changes made for high resolution process accounting (described above), some changes have been made to the Frequency-Based Scheduler (FBS) and the Performance Monitor (PM).

The timing mode, which determines whether PM includes or excludes time spent servicing interrupts in its measurements, is no longer tied to how CPU accounting is performed on the system.

The kernel tunables used to configure these utilities have changed; they are still accessible through the Frequency-Based Scheduler selection of the Linux Kernel Configuration menu. Note that both FBS and high resolution process accounting must be configured in order to use PM.

For information about FBS and PM, see the *RedHawk Linux Frequency-Based Scheduler (FBS) User's Guide*.

8.2.3 Real-Time Clock and Interrupt Module (RCIM)

The following additions have been made to support the Real-Time Clock and Interrupt Module (RCIM):

- You can now modify configuration options dynamically by using `echo (1)` to write a suitable configuration string to `/proc/driver/rcim/config`. This requires write permission to that file and should only be done when the RCIM is not in use.
- An `ioctl (2)` command (IOCTLSIGATTACH) has been provided that allows a process to register with each possible RCIM subdevice for a signal to be delivered on every occurrence of an interrupt generated by that subdevice.
- The RCIM driver now permits distributed interrupts to be used on isolated RCIMs.
- Output to `/proc/driver/rcim/status` has been expanded to include the board IRQ and firmware revision level and the driver major number.
- `lspci (8)` now correctly identifies the RCIM board.
- The RCIM registers and calculations to determine maximum synchronization cable lengths have been added as appendices to the RCIM User's Guide.

For more information about the RCIM, refer to the *Real-Time Clock and Interrupt Module (RCIM) PCI Form Factor User's Guide*.

8.2.4 Rescheduling Variables

If a rescheduling variable is locked when a parent process makes a call to `fork (2)`, `vfork (2)` or `clone (2)`, the rescheduling variable will now abort with `RESCHED_EFORK`. Previously, the rescheduling variable was locked in the child process created by the system call.

8.2.5 POSIX Message Queue Facility

The POSIX message queue facility has been upgraded in RedHawk Linux to version 4.10 of the open source patch. In this patch, `ioctls` have been replaced with a system call implementation. In addition, read/write files under

`/proc/sys/fs/mqueue` are included for setting or determining message queue limits. The maximum size limit for all queues on the system (MQ_MAXSYSSIZE) has been removed.

As in previous releases of RedHawk, all applications that call the message queue library routines must link to Concurrent's real-time library (libccur_rt). On the iHawk system, linking dynamically retains compatibility with earlier releases of RedHawk.

Complete information about using POSIX message queues can be found in the *RedHawk Linux User's Guide*.

9.0. Known Issues

Special consideration should be given to the following areas:

1. There are sporadic problems with password files that are distributed to multiple systems via NIS. This can prevent valid users from logging into the system when their accounts are made known to that system via NIS.
2. NFS mounts of large partitions will sporadically have problems in seeing all of the files that should be mounted in the partition.
3. When running the Linux Test Project test suite LTP-2002(*), the user must remove the **Fork07** test case prior to building the test suite to prevent a system hang on RedHawk systems with greater than 1GB of memory. This is a known bug in the official Linux kernel. For more information, refer to the Linux Test Project web site at <http://ltp.sourceforge.net/>
4. The “Console Redirection” BIOS feature, when enabled, has been observed to interfere with the proper operation of the integrated VGA video and the XFree86 X server with some iHawk platforms, such as the Dell PowerEdge™ 6650.
5. 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.
6. Using USB keyboards with the **kdb** kernel debugger is not supported.
7. Note that in order to use the **pam_capability(8)** feature with **ssh**, the **/etc/ssh/sshd_config** file must have the following option set:

```
UsePrivilegeSeparation no
```

8. Currently, commands run with **rexec(1)** on a RedHawk target fail to inherit capabilities even when the appropriate **pam_capability.so** entry is made in the **/etc/pam.d/rexec** file. Use the **rsh** facility instead.
9. When a system configuration change is made that affects the number of CPUs (for example when booting with a uniprocessor kernel or with hyper-threading), 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.

10. Kernel crash dumps and hyper-threading are not currently supported on a model 868 8-way iHawk system. Hyper-threading should be disabled in the BIOS before booting RedHawk Linux on this system.
11. Under earlier versions of RedHawk Linux, output from **lspci(8)** displays the RCIM peripheral as “PLX Technology, Inc.: Unknown device.” To correct this so that the RCIM is correctly identified, update

`/usr/share/hwdata/pci.ids` with the version of this file supplied in the RedHawk kernel tree. Contact Concurrent Software Support (see page 35 for details) if you require assistance.

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

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

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

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

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

13. Note that ide-scsi emulation has been deprecated in the 2.6.0 kernel series (and therefore the RedHawk 2.X series). This feature may cause system instability and its use is discouraged.

10.0. Using the Linux Kernel Crash Dump (LKCD)

Note: Currently, LKCD is supported only on iHawk 860 systems with the exception of model 868 (8-way).

LKCD contains kernel and user level code designed to:

- save the kernel memory image when the system dies due to a software failure
- recover the kernel memory image when the system is rebooted
- analyze the memory image to determine what happened when the failure occurred

The memory image is stored into a dump device, which is represented by one of the disk partitions on the system.

10.1 Documentation

After LKCD has been installed on your system, the documents listed in Table 10-1 are available at `/usr/share/doc/lkcd`. The user should be familiar with the contents of these documents.

Table 10-1 LKCD and Lcrash Documents

File Name	Title of Document
lkcd_tutorial.pdf	LKCD Installation and Configuration
lcrash.pdf and lcrash.htm	Lcrash HOWTO

Additional information about LKCD and the LKCD project can be found at:

<http://lkcd.sourceforge.net/>

10.2 Installation/Configuration Details

The `lkcdutils` RPM is automatically installed as part of the RedHawk Linux installation. The default RedHawk kernel configurations include the LKCD kernel patch. Concurrent has added scripts to automatically patch the `/etc/rc.d/rc.sysinit` file to configure LKCD to take dumps and to save dumps at boot time. LKCD will automatically self-configure to use the swap partition as the dump device. This is done by creating the file `/dev/vmdump` as a symbolic link to the swap partition.

The dump is saved to a disk partition and then copied to the dump directory on reboot. The system administrator must set up the disk partitions so there is enough space in `/var/log/dump` and the swap partition.

The `/etc/sysconfig/dump` configuration file contains configuration details. This file may be edited by the user, for example to select compression types (GZIP compressed dumps are the default) or to change the directory where dumps are saved.

10.3 Forcing a Crash Dump on a Hung System

LKCD contains a Magic System Request (SysRq) option to force a dump. By default, Concurrent has configured the kernel with this option. To use SysRq, it must be enabled through the `/proc` filesystem as follows:

```
echo 1 >/proc/sys/kernel/sysrq
```

The configuration file `/etc/sysctl.conf` sets the default value.

Two methods to force a dump are described below. Use the method applicable to your configuration:

Using PC keyboard:	Use the Alt+SysRq+C
Using serial console:	Break followed by C

An example of how you can send a Break using a serial console is illustrated below:

Using minicom as terminal emulator	Ctrl+A+F to send a Break, followed by C
------------------------------------	-----------------------------------------

Refer to the `lcrash` documentation for information on how to run `lcrash`. A dump is a set of files, for example:

```
# cd /var/log/dump
# ls
0 1 bounds

# cd 0
# ls
analysis.0  dump.0  kerntypes.0  lcrash.0  map.0

# lcrash -n 0

Please wait...
  Initializing vmdump access... Done.
  Loading system map... Done.
  Loading type info (Kerntypes)... Done.
  Initializing arch specific data... Done.
  Loading ksyms from dump..... Done.

>> stat | less
```

This should display information about when the dump was taken, the panic string and the kernel messages.

11.0. Useful Information

11.1 Man Pages

Man pages can be found online. Type **man** followed by a command or man page name to display information about the command.

11.2 RCIM Connection Modes

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

Isolated mode	Neither the input nor output connector is used. There is no connection between this RCIM and other RCIMs.
Master mode	The RCIM is connected to one or more other RCIMs, and the RCIM is at the head of the chain. There is an output cable connection from a master RCIM but there is no input cable connection into it. The RCIM master is unique in that it controls the synchronized clocks. See Chapter 3 of the <i>Real-Time Clock and Interrupt Module (RCIM) PCI Form Factor User's Guide</i> for a discussion of synchronized clocks.
Pass-through Slave mode	The RCIM is connected to one or more other RCIMs but not at the head of the chain. There is an input cable connection into a pass-through slave RCIM and an output cable connection to another slave.
Final Slave mode	The RCIM is connected to one or more other RCIMs at the end of the chain. There is an input cable connection into a final slave RCIM but no output cable connection.

11.3 File System Swap Size Guidelines

The table below is intended to help you determine the partition size of the swap file system. The swap partition size is based on the size of the system's main memory plus an additional 5 percent. This swap space allocation should be adequate to accommodate system crash dumps.

Table 11-1 (next page) supplies the recommended swap partition sizes for various sizes of main memory.

Table 11-1 Swap Partition Size Guidelines

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


12.0. Software Updates and Support

12.1 General

Available updates to RedHawk Linux software products can be downloaded from: <http://redhawk.ccur.com>. This site also includes the most recent documentation available.

The site login windows are shown below.

Customer login ID and password can be found in the shipping documents accompanying the system.



12.2 How to Download and Install Updates

Upon login, you will see a list of packages that you have not yet downloaded. Click on the packages to download them and save them in a suitable location (e.g. `/root/updates`).

To install the newly downloaded packages, log in as root in single-user mode. Change directory to the location of the updates and issue the following command:

```
rpm -Fvh * .rpm
```

The command should complete successfully, though it may take several minutes depending on the number of updates being installed. When complete, exit single-user mode (^D).

If you have problems, please contact the Concurrent Software Support Center (see “Direct Software Support” on page 35 for further information).

13.0. Direct Software Support

For answers to basic RedHawk Linux questions, check out the FAQ posted on Concurrent's web site: <http://www.ccur.com>. Select Integrated Solutions Division -> Support -> FAQ.

If you need assistance or information about your system, please contact the Concurrent Software Support Center at our toll free number (1-800-245-6453). Our customers outside the continental United States can contact us directly at 1-954-283-1822. The Software Support Center operates Monday through Friday from 8 a.m. to 7 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.

Concurrent provides a Software Action Request (SAR) form which our customers can fill out and submit to their local field analyst or the Software Support Center. This procedure ensures that your request is entered into our SAR database for follow-up and action.

To obtain copies of SAR forms, call the Software Support Center and request form number CSD1833B.

