

 REDHAWK *Architect™ User's Guide*

Copyright 2019 by Concurrent Real-Time, Inc. All rights reserved. This publication or any part thereof is intended for use with Concurrent Real-Time products by Concurrent Real-Time personnel, customers, and end-users. It may not be reproduced in any form without the written permission of the publisher.

The information contained in this document is believed to be correct at the time of publication. It is subject to change without notice. Concurrent Real-Time makes no warranties, expressed or implied, concerning the information contained in this document.

To report an error or comment on a specific portion of the manual, photocopy the page in question and mark the correction or comment on the copy. Mail the copy (and any additional comments) to Concurrent Real-Time, 2881 Gateway Drive, Pompano Beach, FL 33069. Mark the envelope “**Attention: Publications Department.**” This publication may not be reproduced for any other reason in any form without written permission of the publisher.

Concurrent Real-Time 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.

Printed in U. S. A.

Revision History:	Level:	Effective With:
November 2008	000	RedHawk Linux 5.1
January 2009	100	RedHawk Linux 5.2
February 2009	200	RedHawk Linux 5.2
July 2009	300	RedHawk Linux 5.2
October 2009	400	RedHawk Linux 5.2
July 2010	600	RedHawk Linux 5.4
October 2011	700	RedHawk Linux 6.0
April 2012	720	RedHawk Linux 6.0
December 2012	800	RedHawk Linux 6.3
July 2013	900	RedHawk Linux 6.3
September 2013	920	RedHawk Linux 6.3
February 2014	930	RedHawk Linux 6.3
August 2014	940	RedHawk Linux 6.5
September 2014	950	RedHawk Linux 6.5
October 2014	960	RedHawk Linux 6.5
May 2015	7.0	RedHawk Linux 7.0
August 2015	7.0-1	RedHawk Linux 7.0
June 2016	7.2	RedHawk Linux 7.2
December 2016	7.2-1	RedHawk Linux 7.2
August 2017	7.2-2	RedHawk Linux 7.2
October 2017	7.3	RedHawk Linux 7.3
April 2018	7.3-1	RedHawk Linux 7.3
September 2018	7.5	RedHawk Linux 7.5
March 2019	7.5-1	RedHawk Linux 7.5

Scope of Manual

This manual describes the RedHawk Architect™, an easy-to-use GUI interface for creating and maintaining a runtime and development environment for a target computer.

Structure of Manual

This manual consists of:

- Chapter 1, which introduces you to RedHawk Architect and guides you through its use.
- Chapter 2, which explains how to use the security extension of the Advanced Security Edition of Architect.
- Chapter 3, which explains Importing ISO Images to avoid repetitive manual optical media insertion.
- Chapter 4, which explains PXE Management.
- Appendix A explains Manual DHCP configuration for Architect PXE targets.
- An Index containing an alphabetical reference to key terms and concepts and the pages where they occur in the text.

Syntax Notation

The following notation is used throughout this manual:

<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.
	Operating system and program output such as prompts, messages and listings of files and programs appears in list type.
[]	Brackets enclose command options and arguments that are optional. You do not type the brackets if you choose to specify these options or arguments.
hypertext links	When viewing this document online, clicking on chapter, section, figure, table and page number references will display the corresponding text. Clicking on Internet URLs provided in blue type

will launch your web browser and display the web site. Clicking on publication names and numbers in **red** type will display the corresponding manual PDF, if accessible.

Related Publications

The following table lists Concurrent Real-Time documentation for RedHawk Architect and the components that can be installed using RedHawk Architect. Depending upon the document, they are available online on RedHawk Linux systems or from Concurrent Real-Time's documentation web site at <http://redhawk.concurrent-rt.com/docs>.

RedHawk Architect	Pub. Number
<i>RedHawk Architect Release Notes</i>	0898600
<i>RedHawk Architect User's Guide</i>	0898601
RedHawk Linux	
<i>RedHawk Linux Release Notes</i>	0898003
<i>RedHawk Linux User's Guide</i>	0898004
<i>RedHawk Linux Cluster Manager User's Guide</i>	0898016
<i>RedHawk Linux FAQ</i>	N/A
NightStar RT Development Tools	
<i>NightView User's Guide</i>	0898395
<i>NightTrace User's Guide</i>	0898398
<i>NightProbe User's Guide</i>	0898465
<i>NightTune User's Guide</i>	0898515

Contents

Preface	iii
Chapter 1 Using RedHawk Architect	
Introducing Architect	1-1
Creating a root File System for Target Systems	1-1
Running Architect	1-2
Creating a New Session	1-4
Selecting Software to Install in the Image	1-4
Selecting Base Distribution Linux Packages	1-5
Using the Base Environments View	1-6
Using the Categories and Groups View	1-7
Using the All Packages View	1-7
Using the Selected Packages View	1-8
Selecting RedHawk OS Options	1-9
Selecting NightStar Options	1-11
Configuring an Image	1-11
Configuring General Settings	1-12
Configuring a Console	1-13
Configuring Networking	1-15
Configuring File Systems	1-17
Simple Disk Partitioning	1-18
Advanced Disk Partitioning	1-19
Building an Image	1-23
Customizing an Image	1-27
Software Updates	1-27
System Services	1-29
Kernel Manager	1-30
Configure Custom Kernel	1-31
Import Kernel Configuration	1-32
Export Kernel Configuration	1-33
Compile Custom Kernel	1-33
Remove Custom Kernel	1-34
Additional RPMs	1-35
Installing Board Support Packages	1-35
File Manager	1-36
Chroot Shell	1-36
Image Cleanup	1-38
Deploying an Image	1-38
Deploy to USB Device	1-39
Install via USB drive	1-42
Install via DVD media	1-43
Installing via PXE over a Network	1-45
Bootting Diskless via PXE over a Network	1-47
Deploy to Virtual Machine	1-54
Editing an Existing Session	1-55

Chapter 2 Security Extensions

SELinux	2-1
Configuring SELinux	2-1
Security Content Automation Protocol (SCAP)	2-4
Introduction to SCAP	2-4
Overview of SCAP Workflow	2-4
Understanding SCAP Evaluation and Remediation Scans	2-5
Configuring SCAP	2-6
Building the SCAP-configured Image	2-8
Customizing the SCAP-configured Image	2-9
Deploying the SCAP-configured Image	2-10
Running SCAP Scans on Deployed Target Systems	2-10
Customizing SCAP Content Using SCAP Workbench	2-12
Known SCAP Issues	2-14

Chapter 3 Importing ISO Images

Importing ISO Images	3-1
Importing ISO Images From Optical Media	3-2
Copying ISO Images From Existing ISO Images	3-3
Linking To Existing ISO Images	3-3
Deleting Imported ISO Images	3-4

Chapter 4 PXE Management

Enabling PXE on Targets	4-1
Initializing PXE Services	4-1
Managing PXE Images	4-3
PXE Installer Images	4-4
PXE Diskless Images	4-5
Managing PXE Targets	4-7
Adding Targets	4-7
Adding Single Targets	4-7
Adding Multiple Targets	4-9
Removing Targets	4-11
Editing Targets	4-11

Appendix A Manual DHCP Configuration

Overview	A-1
Installing DHCP Configuration	A-2

Index

Index-1

Using RedHawk Architect

This chapter introduces you to RedHawk Architect and provides instructions for its use.

Introducing Architect

RedHawk Architect is an easy-to-use GUI interface for configuring, building and flashing embedded solutions.

RedHawk Architect greatly simplifies the following tasks to create and maintain a runtime and development environment for single board computers (SBCs):

- installing custom configurations of the CentOS® or Red Hat® Enterprise Linux distribution
- installing and configuring the RedHawk™ operating system
- installing SBC-specific board support packages (BSPs)
- installing NightStar™ RT application development tools
- installing RedHawk and NightStar software updates
- maintaining and reconfiguring an SBC's root file system
- deploying root file system images onto target systems or virtual machines

With Architect, you choose the Linux and application modules that will be installed with the RedHawk kernel. For example, you can select all packages in categories such as Virtualization or a subset of the packages, for a minimal configuration. Architect allows the Linux file system to be customized and minimized for embedded applications using flash memories as small as 1 GB.

Architect creates and processes a configuration file based upon your selections and performs actual RPM package installations. It prompts you to insert the necessary RedHawk, CentOS, Red Hat, and NightStar media depending upon the features selected.

Creating a root File System for Target Systems

To create a target file-system image that can be used on a single board computer, use RedHawk Architect on a supported host system to perform the following steps:

1. Select the software to install in the image.
2. Configure the image.
3. Customize the image for your embedded application.

4. Deploy the image on your target boards or to virtual machines.

These procedures are described in the sections that follow. The steps may be repeated to change the image and/or deploy it any number of times.

Running Architect

For instructions on installing RedHawk Architect, refer to the *RedHawk Architect Release Notes*.

Architect must be run as the root user.

To run Architect, type **architect** at a shell prompt:

```
# architect
```

Note that the **sudo (8)** command can alternatively be used to run Architect if that method is preferred to using a root shell.

The very first time Architect is invoked after it is installed, a dialog appears presenting you with the Concurrent Real-Time End User Agreement.



Figure 1-1 Architect End User Agreement

When Architect starts, a dialog appears presenting you with the option to start a new session or load an existing session.

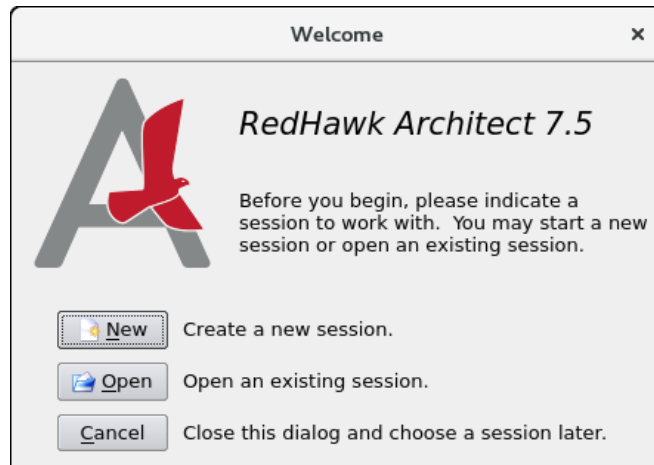


Figure 1-2 Opening RedHawk Architect Dialog

To start a new session, click on the **New** button. See “Creating a New Session” on page 1-4 for details.

A session can be saved at any time and loaded later to continue work on the file system image. To edit an existing session, click on the **Open** button. See “Editing an Existing Session” on page 1-55 for details.

When the **Cancel** button is clicked, the RedHawk Architect main window appears, as shown in the following figure.

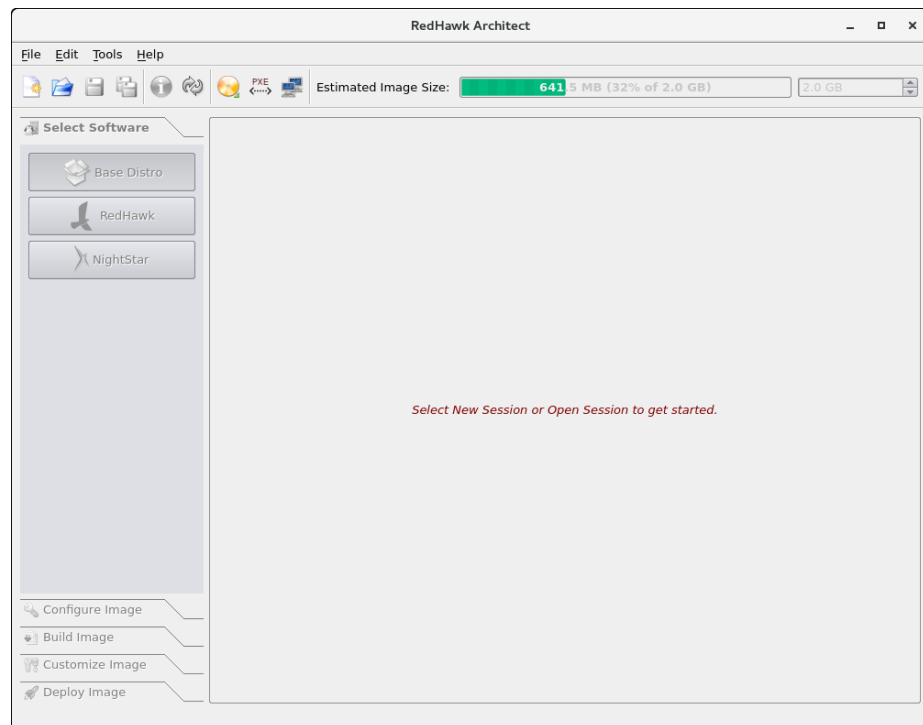




Figure 1-3 RedHawk Architect Main Window


From this window, the most common tasks performed are:

- start a new session by clicking on the New Session icon  or selecting **New Session** from the **File** menu along the top of the window. See “Creating a New Session” on page 1-4 for details.
- edit an existing session by clicking on the Open Session icon  or selecting **Open Session** in the **File** menu. See “Editing an Existing Session” on page 1-55 for details.

Creating a New Session

An Architect session describes all decisions made about a particular target file-system image, including:

- the target SBC
- which software should be installed
- how the software should be configured

When you select the **New** button from the opening Architect dialog, or the **New Session** icon  or **New Session** from the **File** menu along the top of the RedHawk Architect main window, the **New Session** dialog, shown below, displays.




Figure 1-4 New Session Dialog

This dialog enables you to specify the version of RedHawk to be used for creating the target file-system image. Be sure that you have the correct version of the RedHawk, CentOS, Red Hat Enterprise Linux media or ISO files necessary to create a target file-system image of the specified RedHawk release.

Selecting Software to Install in the Image


To select the software to install in the file system image, click on **Select Software** from the toolbox on the left side of the RedHawk Architect main window. This allows you to select software from the following three groups:

- Base Distribution Linux packages

- RedHawk Linux operating system
- NightStar tools

Some initial selections are made for you by default; e.g. the core RedHawk OS. These packages appear with a gray check mark and cannot be deselected.

The Estimated Image Size gauge at the top of the RedHawk Architect main window indicates the approximate size the file system image will be when built. It also indicates the percentage of the target board's root device that will be consumed by the image.

Once an image is built, you may click on the  Refresh Image Size button to calculate the actual image size as it is stored on disk. Alternatively, you may select Refresh Image Size from the Tools menu. The Estimated Image Size gauge will be updated to reflect the *current* actual size.

The spin control box to the right of the Estimated Image Size gauge may be used to change the desired maximum size of the image. This value cannot exceed the known size of the root device but it can be made smaller. If the size of the root device is unknown the maximum value allowed is 1 terabyte.

The Undo button can be used to reverse the last package select or package deselect operation. This can be used repeatedly to reverse several operations if desired, which is useful for experimenting with package sets to see the effect on the estimated image size.

To get more information about a package, right-click to display a menu of options. Multiple packages can be processed as a group by highlighting the packages and then right-clicking to display the menu of options. When choosing the select or deselect menu options, software dependent on the highlighted packages will be automatically selected or deselected. The figure below shows three packages selected as a group.

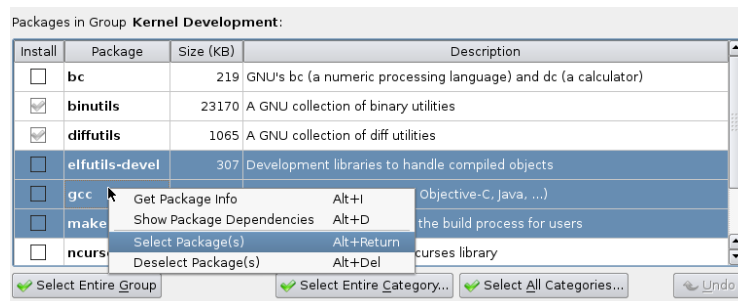


Figure 1-5 Selecting Multiple Packages

Selecting Base Distribution Linux Packages

To select CentOS or Red Hat packages for the file system image, click on the Base Distro selection from the Select Software toolbox.

CentOS or Red Hat packages may be navigated by way of several “Package Views”. Select the desired package view from the Package Views drop-down menu. The following views are available and are described in the subsections that follow.

- Base Environments

- Categories and Groups
- All Packages
- Selected Packages

Note that the Base Environments view will be initially selected; a base package environment must be chosen before the other package views will become available.

Using the Base Environments View

The Base Environments view requires the user to choose a high-level task-based characterization of the CentOS or Red Hat packages that will be initially selected. This view should be very familiar to users that have previously performed a native CentOS or Red Hat installation; it is shown in the following figure.

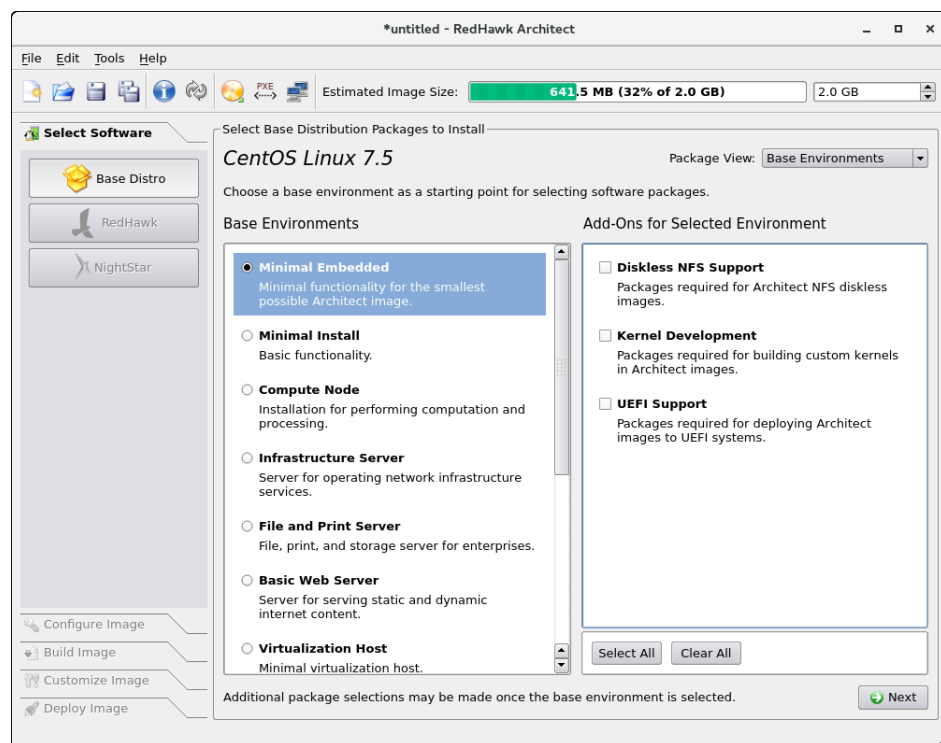


Figure 1-6 Choosing Target Characterization, Base Environments View

For example, if the target will primarily be used to run a web server, choose the **Basic Web Server** environment. Note that the set of base environments available may be different depending on the current session's distribution type and revision.

To see more information about a particular environment, click on the **Get Environment Info** button that is displayed when you place the cursor over the environment and right-click.

Once a Base Environment has been chosen, a list of corresponding optional package groups will be displayed in the **Add-Ons for Selected Environment** area. You can choose these package groups individually or press the **Select All** and **Clear All** buttons to affect all optional package groups at once.

After you have chosen the desired base environment and associated optional packages, press the **Next** button at the lower right to add all of the corresponding packages to the session and enable the other package views for further package customization.

Using the Categories and Groups View

The **Categories and Groups** view provides a view of CentOS or Red Hat packages organized in a hierarchy of groups. The package group hierarchy is the standard CentOS and Red Hat package group hierarchy. This view is shown in the following figure.

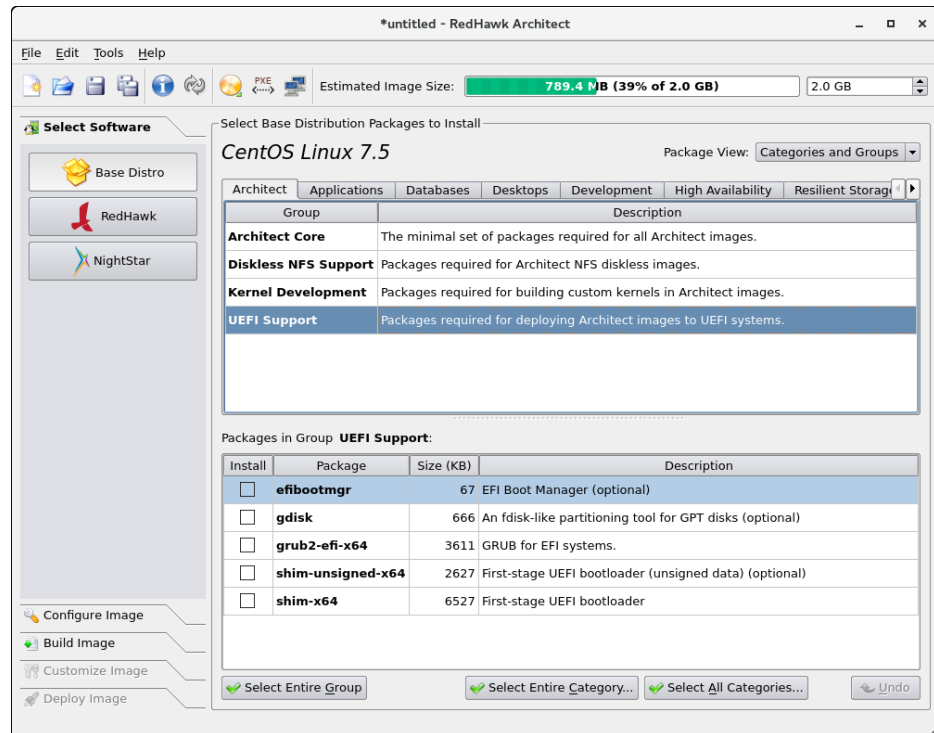


Figure 1-7 Selecting Base Distro Packages, Categories and Groups View

Packages may be selected or deselected by choosing a package group in the upper pane and then checking or un-checking packages in that group in the lower pane. All the packages in a chosen group may be selected with the **Select Entire Group** button.

All the packages that are in all the groups of the currently chosen package category may be selected with the **Select Entire Category** button. Also, the **Select All Categories** button may be used to select all of the packages in all of the groups of all of the categories.

Using the All Packages View

The **All Packages** view, as shown in the following figure, provides a sorted linear list of all CentOS or Red Hat packages.

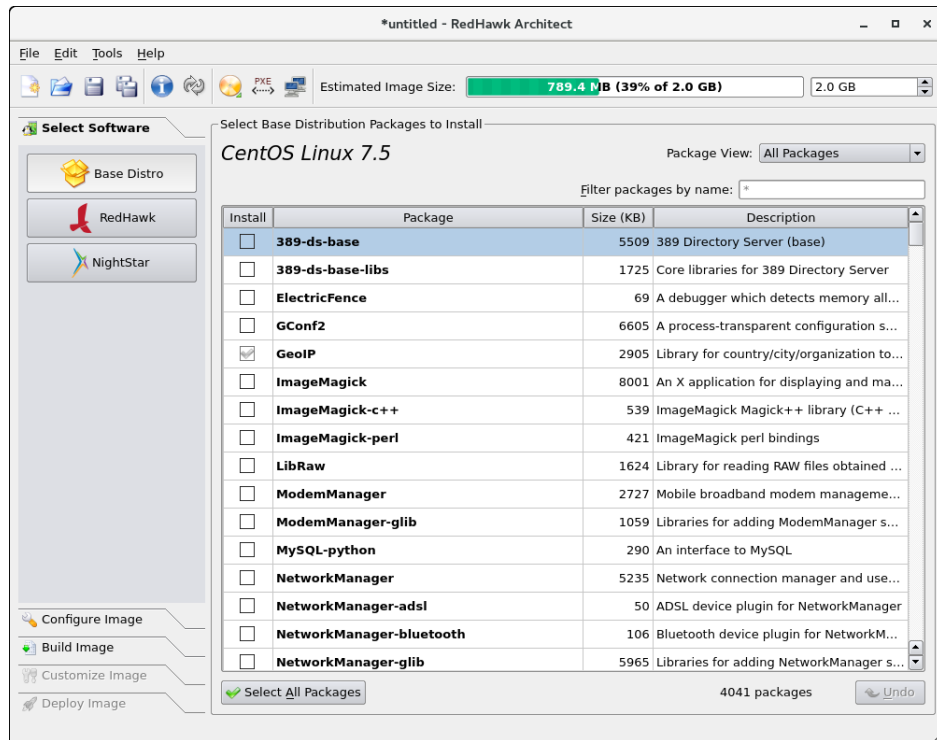


Figure 1-8 Selecting Base Distro Packages, All Packages View

Packages may be selected or deselected from this list. The Filter packages by name box allows you to search for packages by name.

All packages can be selected by clicking on the Select All Packages button.

Using the Selected Packages View

The Selected Packages view, as shown in the following figure, provides a sorted linear list of all CentOS or Red Hat packages that are currently selected for installation.

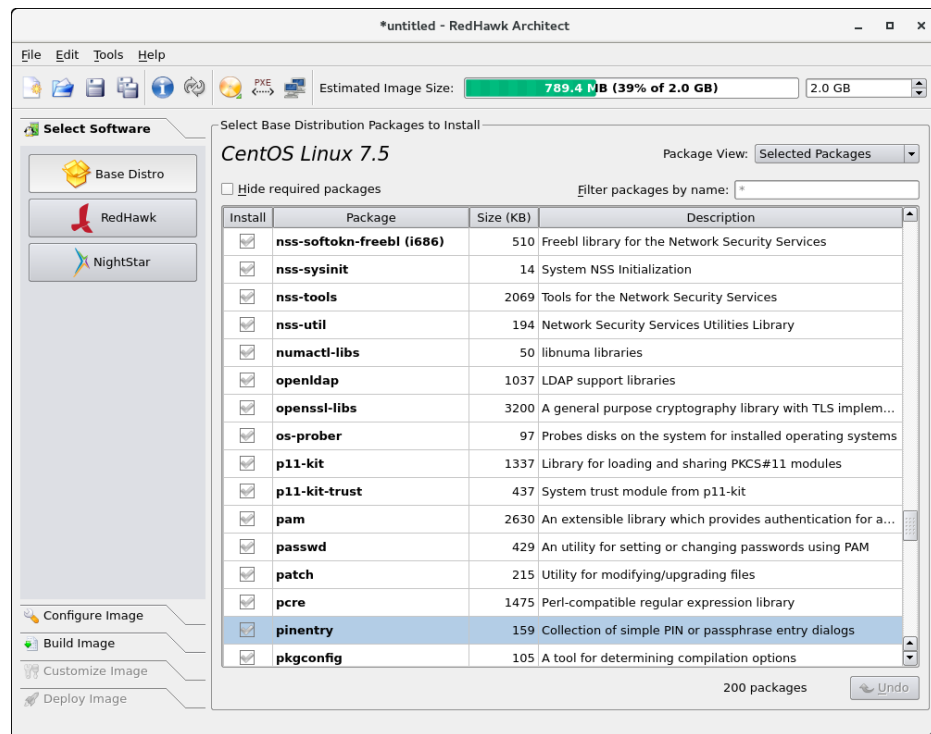


Figure 1-9 Selecting Base Distro Packages, Selected Packages View

Packages may be deselected from this list. The Filter packages by name box allows you to search for packages by name.

To exclude the required packages from the list, check the Hide required packages check box. When this box is checked, only the optional packages are shown.

Selecting RedHawk OS Options

To select RedHawk Linux OS and kernels for the file system image, click on the RedHawk selection from the Select Software toolbox. The RedHawk page, shown in the following figure, displays.

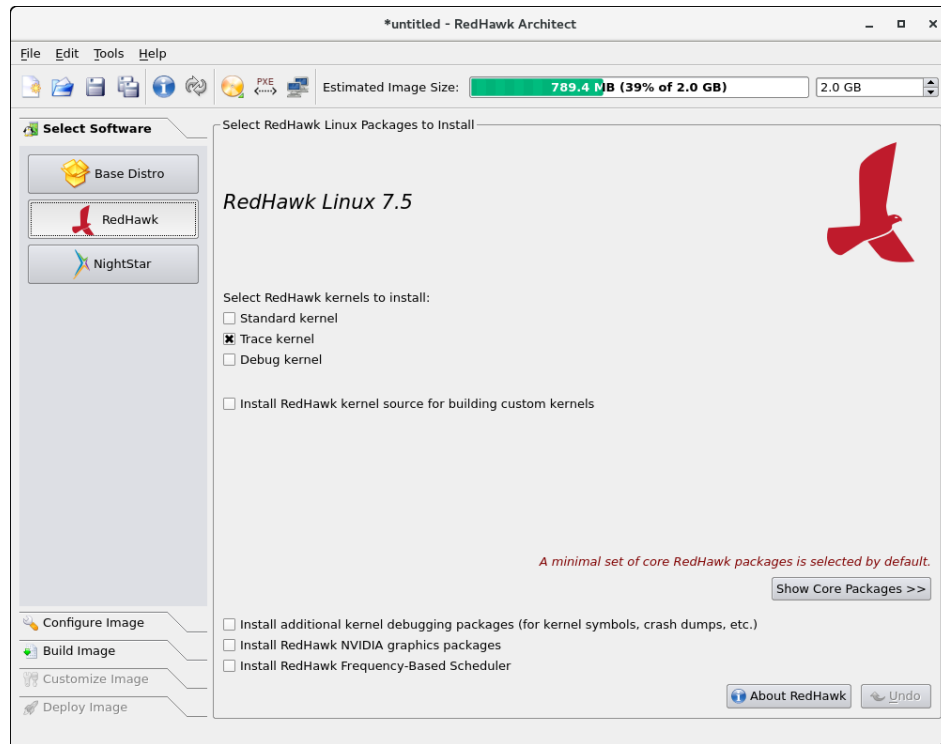


Figure 1-10 Selecting RedHawk Options

Select which RedHawk kernel(s) to install by checking the appropriate check box(es): Standard, Trace and/or Debug. The standard kernel does not have tracing or debugging capabilities and it is the smallest sized kernel option. The trace kernel does offer tracing capabilities but it does not have debugging capabilities. The debug kernel offers both debugging and tracing capabilities. Note that at least one kernel *must* be selected; the GUI enforces this by ensuring that a sole remaining selection cannot be deselected.

Select Install RedHawk kernel source for building custom kernels to ensure that the complete kernel source code will be installed in the image. The kernel source is only required for building custom kernels and loadable kernel drivers.

Advanced users may wish to press the Show Core Packages >> button to refine exactly which RedHawk packages they wish to install from the complete set of RedHawk packages that are available on the media. Normally this is not necessary, but the option exists to facilitate minimizing the image size for very small flash devices.

Select Install additional kernel debugging packages to install extra support for live kernel debugging. This option is also required for RedHawk to be able to create crash dumps. See the *RedHawk User's Guide* for more information.

If the target system has an NVIDIA graphics or CUDA card you may want to select the Install NVIDIA graphics packages radio button. Note that you should only select this option if the target system actually has NVIDIA hardware.

Select Install Frequency-Based Scheduler if you have previously purchased and wish to install the RedHawk FBS software into the target image.

Selecting NightStar Options

To select NightStar tools for the image, click on the NightStar selection from the Select Products toolbox. The NightStar RT page, shown in the following figure, displays.

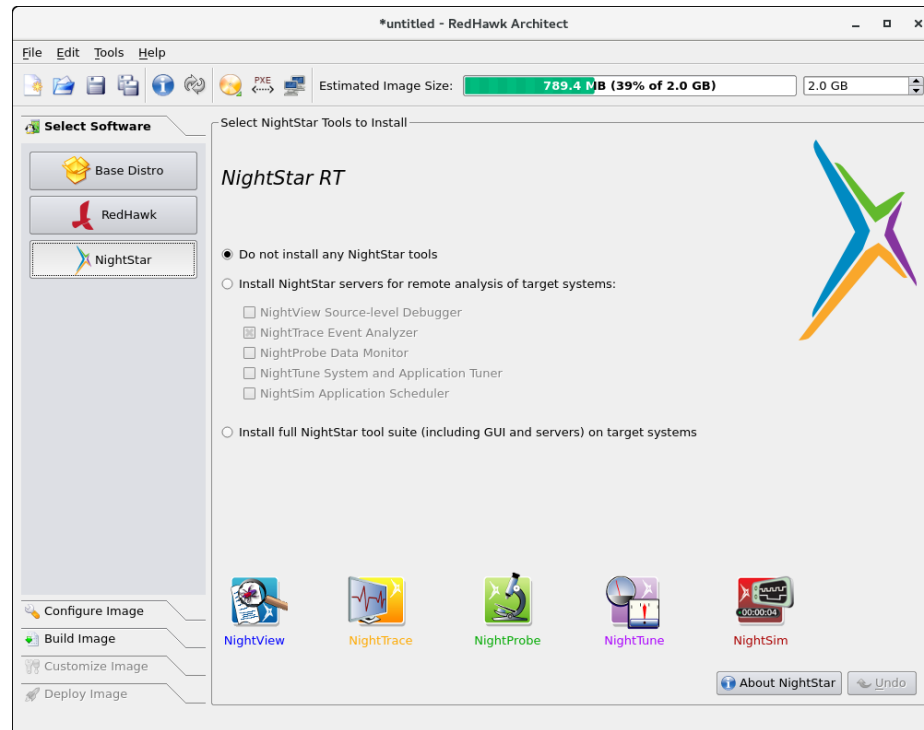


Figure 1-11 Selecting NightStar Tools

By default no NightStar tools will be installed in the target image. Choose the **Install select NightStar servers only** radio button if you wish to only install NightStar remote support for various tools. You may select individual servers from the list by clicking on the check boxes for each tool. The remote support allows NightStar tools running on a host system to connect to and control the target remotely.

Choose the **Install all NightStar clients and servers** radio button to indicate that all NightStar servers and clients are to be installed in the image. This allows the target to run all NightStar tools locally, in addition to providing the remote support described above.

Click on the **About NightStar** button to see a detailed description of each of the individual NightStar tools that are available for installation.

Configuring an Image

It is possible to configure a target file-system image before or after the image has been built by selecting **Configure Image** from the toolbox on the left side of the

RedHawk Architect main window. This selection is available before and after an image is built, however note that there are additional Apply buttons present on the pages *after* an image has been built. It is necessary to apply any changes made after the image has been built in order for the changes to be reflected in the on-disk file-system image.

To configure the file system image, select **Configure Image** from the toolbox on the left side of the RedHawk Architect main window. This allows you to configure the following four groups:

- General Settings
- Console
- Networking
- File Systems

Some initial selections are made for you by default.

Configuring General Settings

To configure time zone, root password and default system run level for the file system image, click on **General Settings** from the **Configure Image** toolbox. The **General Settings** configuration page appears, as shown in the following figure.

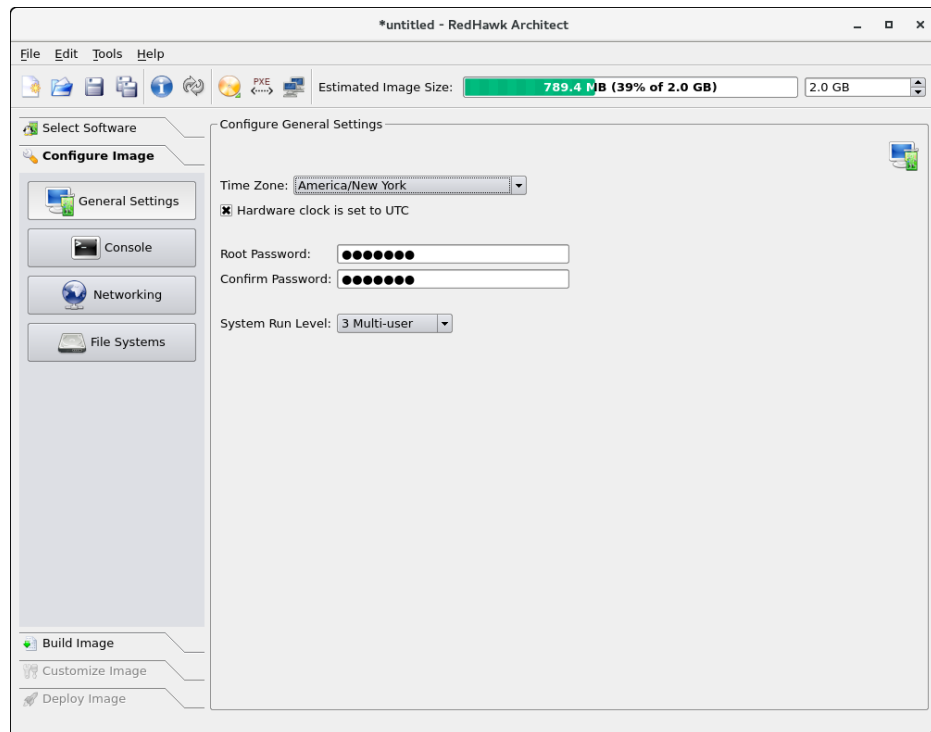


Figure 1-12 General Settings Configuration Page

In the **Time Zone** section, select the proper time zone for your location from the drop-down menu. Click in the check box to indicate if your system clock uses UTC.

NOTE

By default the Hardware clock is set to UTC check box is selected, so be sure to set the target system's BIOS clock in Coordinated Universal Time. If you do not select this, set the BIOS clock according to the selected time zone.

In the Root Password section, enter the root password in the Password field. Reenter it in the Confirm Password field.

NOTE

The default root password is redhawk (all lowercase letters and only one word with no spaces).

In the Run Level section, select the desired default run level from the drop-down menu.

If a change is made to the general settings after the target file-system image has been built, an *Out-of-Sync Notice* will appear at the bottom of the page, as shown in the following figure:

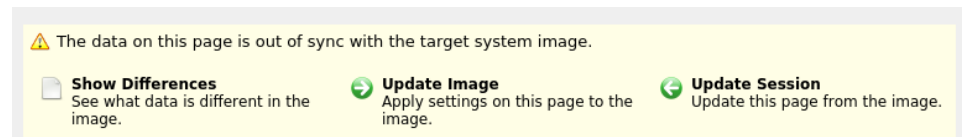


Figure 1-13 General Settings and Image Out-of-Sync Notice

The appearance of an Out-of-Sync Notice on any page indicates that the settings currently displayed in the session do not match the state of the associated target image. Click on **Show Differences** to see which settings are currently out-of-sync. To resolve the issue, it is necessary to either click on **Update Image** or **Update Session**.

The **Update Image** button will apply the currently displayed settings to the target image, whereas the **Update Session** button will change the currently displayed settings to match the state of the target image. The Out-of-Sync Notice will disappear once an update direction has been selected.

Configuring a Console

To configure a serial console for the file system image, click on **Console** from the **Configure Image** toolbox. The **Configure Console** page appears, as shown in the following figure.

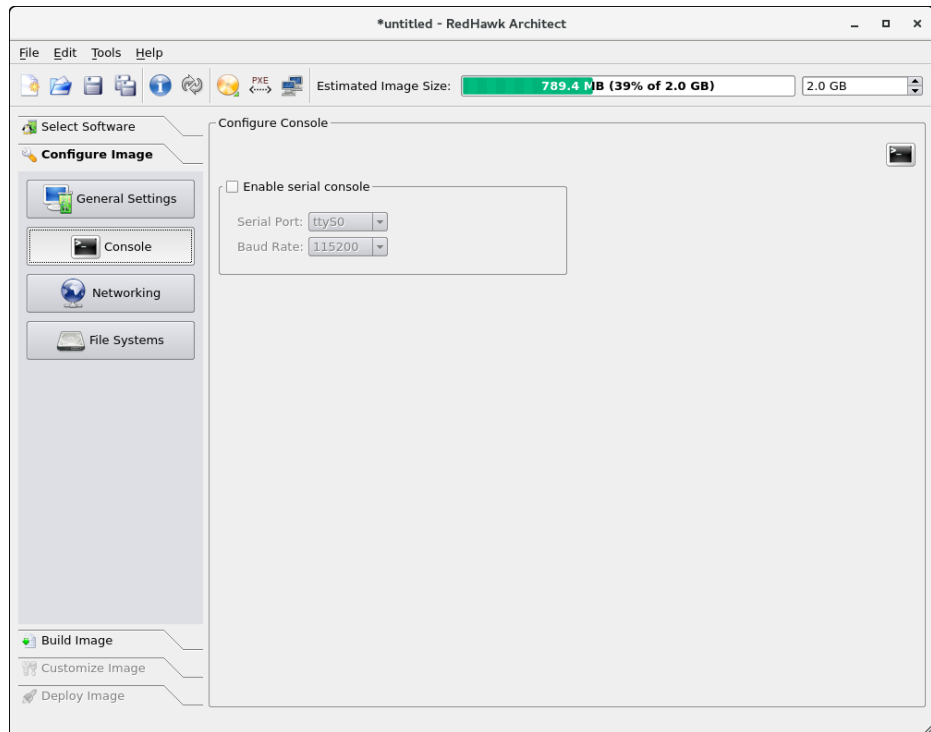


Figure 1-14 Console Configuration Page

Click on the **Enable serial console** check box to activate the fields that define the port and baud rate for the console.

Select a port from the **Serial Port** drop-down menu.

Select a baud rate from the **Baud Rate** drop-down menu.

If a change is made to the console settings after the target file-system image has been built, an *Out-of-Sync Notice* will appear at the bottom of the page, as shown in the following figure:

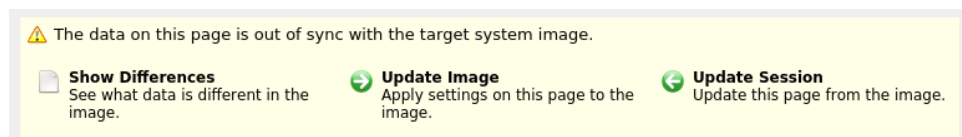


Figure 1-15 Console Settings and Image Out-of-Sync Notice

The appearance of an *Out-of-Sync Notice* on any page indicates that the settings currently displayed in the session do not match the state of the associated target image. Click on **Show Differences** to see which settings are currently out-of-sync. To resolve the issue, it is necessary to either click on **Update Image** or **Update Session**.

The **Update Image** button will apply the currently displayed settings to the target image, whereas the **Update Session** button will change the currently displayed settings to match the state of the target image. The *Out-of-Sync Notice* will disappear once an update direction has been selected.

NOTE

If your target system does not have a serial port do not configure a serial console on this page.

Configuring Networking

To configure networking for the file system image, click on **Networking** from the **Configure Image** toolbox. The **Configure Networking** page appears, as shown in the following figure.

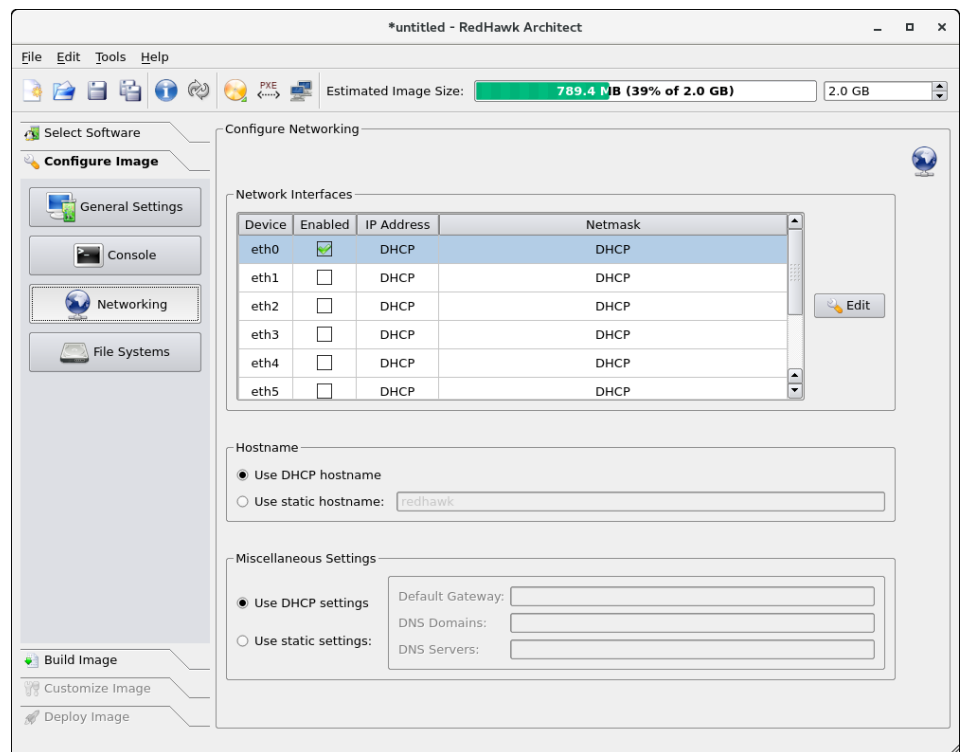


Figure 1-16 Network Configuration Page

All available network interfaces are listed in the **Network Interfaces** section. There may be more or less interfaces shown depending on the target board selected.

To configure a particular network interface, click on the interface to select it, then click on the **Edit** button. The **Configure Network Interface** dialog shown in the following figure displays.

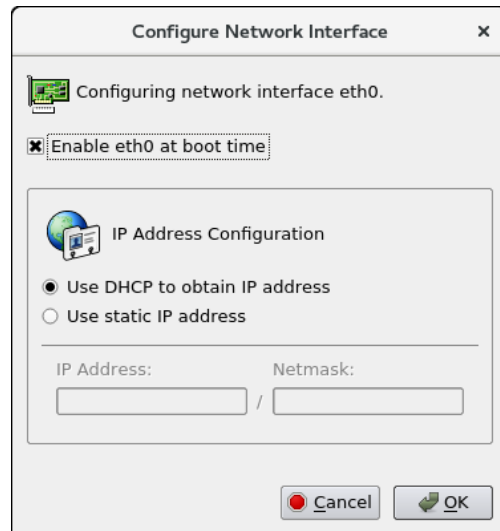


Figure 1-17 Configure Network Interface Dialog

The selected network interface is displayed at the top of the dialog.

Click on the **Enable eth0 at boot time** check box to enable/disable the interface automatically on boot.

Choose the **Use DHCP to obtain IP address** radio button to enable dynamic address configuration, or choose the **Use static IP address** radio button to enable manual address configuration. For manual configurations, enter the IP address and netmask in the appropriate fields.

Click on **OK** to apply the settings to the image and close the dialog. Click **Cancel** to cancel changes.

On the **Configure Networking** dialog under the **Hostname** and **Miscellaneous Settings** areas, either choose to use DHCP or supply the hostname, default gateway, domains, and DNS server addresses in the appropriate fields. Note that multiple DNS domains and DNS servers may be specified by separating multiple entries with either spaces or commas. Be sure to choose to use DHCP appropriately if a DHCP server will be providing some or all of the network parameters dynamically.

If a change is made to the network settings after the target file-system image has been built, an *Out-of-Sync Notice* will appear at the bottom of the page, as shown in the following figure:

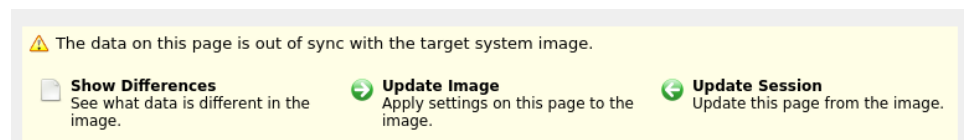


Figure 1-18 Network Settings and Image Out-of-Sync Notice

The appearance of an Out-of-Sync Notice on any page indicates that the settings currently displayed in the session do not match the state of the associated target image. Click on **Show Differences** to see which settings are currently out-of-sync. To resolve the issue, it is necessary to either click on **Update Image** or **Update Session**.

The **Update Image** button will apply the currently displayed settings to the target image, whereas the **Update Session** button will change the currently displayed settings to match the state of the target image. The Out-of-Sync Notice will disappear once an update direction has been selected.

Configuring File Systems

To configure file system options for the file system image, click on **File Systems** from the **Configure Image** toolbox. The **Configure File System** page appears, as shown in the following figure.

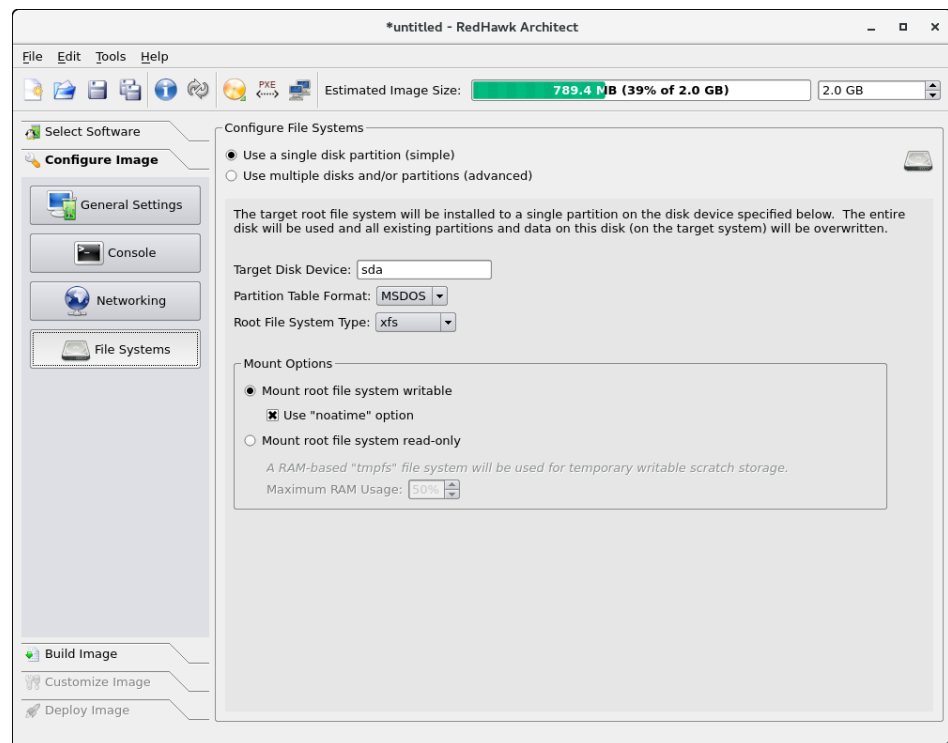


Figure 1-19 File System Configuration Page

There are two different partitioning modes supported: **Use a single disk partition (simple)** and **Use multiple disks and/or partition (advanced)**. This page defaults to the simple disk partitioning mode.

NOTE

When using the PXE Diskless deployment method, all file system configuration settings are ignored; any local drive media that is present on the target will be untouched and ignored. See “Booting Diskless via PXE over a Network” on page 1-47 for more information.

Simple Disk Partitioning

Simple disk partitioning is the traditional partitioning that was offered by early versions of RedHawk Architect. In this mode only a single partition will be created on the specified disk device.

Indicate the desired root device in the **Target Disk Device** field. Use the **Root File System Type** pull-down menu to select the file-system type that you wish to use for the file system that will be created on the target.

Select the desired **Partition Table Format** to be used when initializing the root device; both the MSDOS and GPT partition table formats are supported.

Select the desired **Root File System Type** to be used for the file system on the disk partition; currently the XFS, EXT4, EXT3 and EXT2 file system types are supported.

By default **Mount writable** is selected and the root file system will be mounted with both read and write permissions.

Check the **Use “noatime”** option box to mount the root file system with the *noatime* option. This helps to minimize the number of writes to the root device when root is *not* mounted read-only.

Select **Mount root file system read-only** to mount the root file system *read-only*. Mounting the root file system read-only offers improved security and it will also help preserve the life of root flash devices. When mounting the root file system read-only, a RAM-based file system is then allocated for temporary storage. The **Maximum RAM Usage** for this file system is set, by default, to 50 percent of RAM. The default can be changed by clicking on the up and down arrows of the spin control box.

If a change is made to the file-system settings after the target file-system image has been built, an *Out-of-Sync Notice* will appear at the bottom of the page, as shown in the following figure:

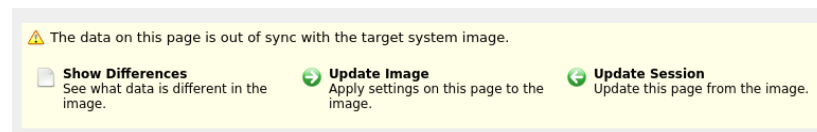


Figure 1-20 File-System Settings and Image Out-of-Sync Notice

The appearance of an Out-of-Sync Notice on any page indicates that the settings currently displayed in the session do not match the state of the associated target image. Click on **Show Differences** to see which settings are currently out-of-sync. To resolve the issue, it is necessary to either click on **Update Image** or **Update Session**.

The Update Image button will apply the currently displayed settings to the target image, whereas the Update Session button will change the currently displayed settings to match the state of the target image. The Out-of-Sync Notice will disappear once an update direction has been selected.

Advanced Disk Partitioning

The advanced disk partition mode provides a more modern and flexible disk partitioning scheme. In this mode you can configure multiple partitions and even multiple disks using the Disk File Systems tab; and special file systems like **tmpfs** and **bind** via the Special File Systems tab. The All File Systems tab will list all the file systems to be configured on the target system.

Use the Disk File Systems tab to configure multiple partitions and/or multiple disks. The default page is shown in the figure below.

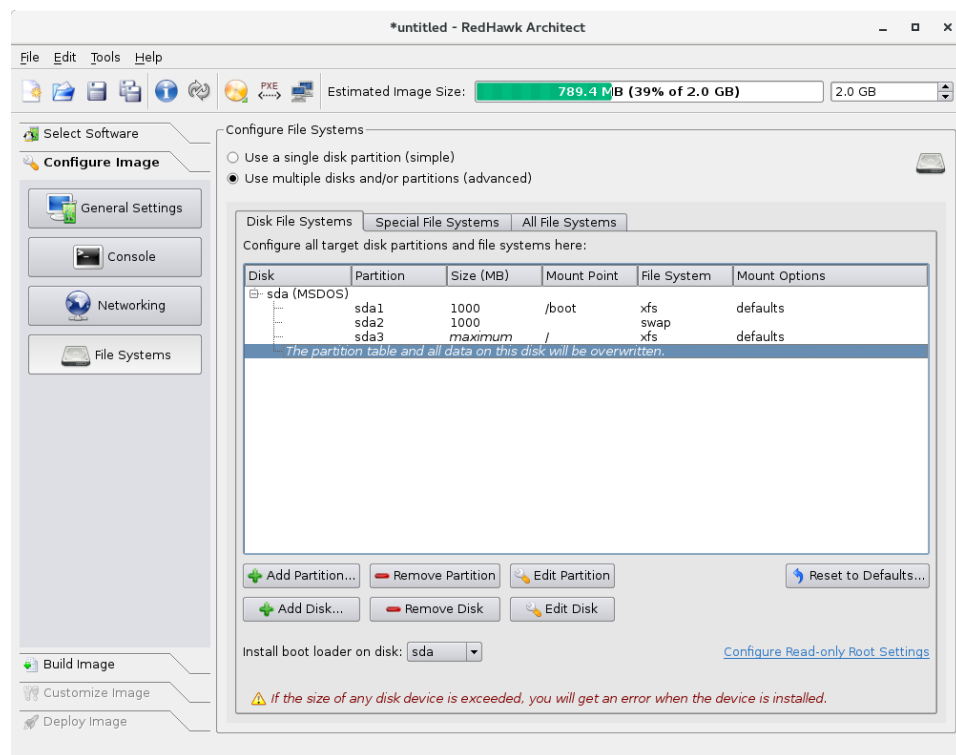


Figure 1-21 Advanced Disk Partitioning Disk File Systems

Press Add Partition... to add new partitions to the currently selected disk.

Press Remove Partition to remove the currently selected partition.

Press Edit Partition to edit attributes of the currently selected partition.

Press Add Disk to add a new disk to the set of currently available disks.

Press Remove Disk to remove the currently selected disk from the set of available disks.

Press **Edit Disk** to edit attributes of the currently selected disk.

Use the **Install boot-loader on MBR of disk** pull-down menu to choose the disk that you will be booting from if multiple disks have been defined.

NOTE

You must ensure that any additional disks defined using advanced partitioning in fact exist on the target for installation to succeed.

NOTE

Multiple disks cannot be partitioned with the **USB Device** deployment tool. In order to use multiple disks you must deploy with one of the **Installer** methods (via DVD, USB or PXE).

To configure the root file system read-only, click on the **Configure Read-only Root Settings** link on the lower right hand of the **Configure File Systems** page. This will bring up a dialog that will instruct you on the steps to take and on implementation choices. The first step is to configure the root file system as read-only via the **Edit Partition** button. Then, click again on the **Configure Read-only Root Settings** link for information on the next steps.

Temporary storage is required when root is configured as read-only. This can be achieved via a RAM-based file system (the default) or by creating a writable file system named `/var/lib/stateless/writable` using the **Add Partition** button. Note that the RAM-based file system size is configurable as a percentage of RAM space. While these two options are writable they are not persistent over boots. Optionally, a persistent file system can be created with a mount point of `/var/lib/stateless/state`. When finished partitioning the disk, verify the configuration by clicking on the **Configure Read-only Root Settings** link one last time.

The figure below shows an example of a root read-only partition scheme with the required scratch storage for root configured as a disk partition (`/var/lib/stateless/writable`) and the optional persistent disk partition (`/var/lib/stateless/state`).

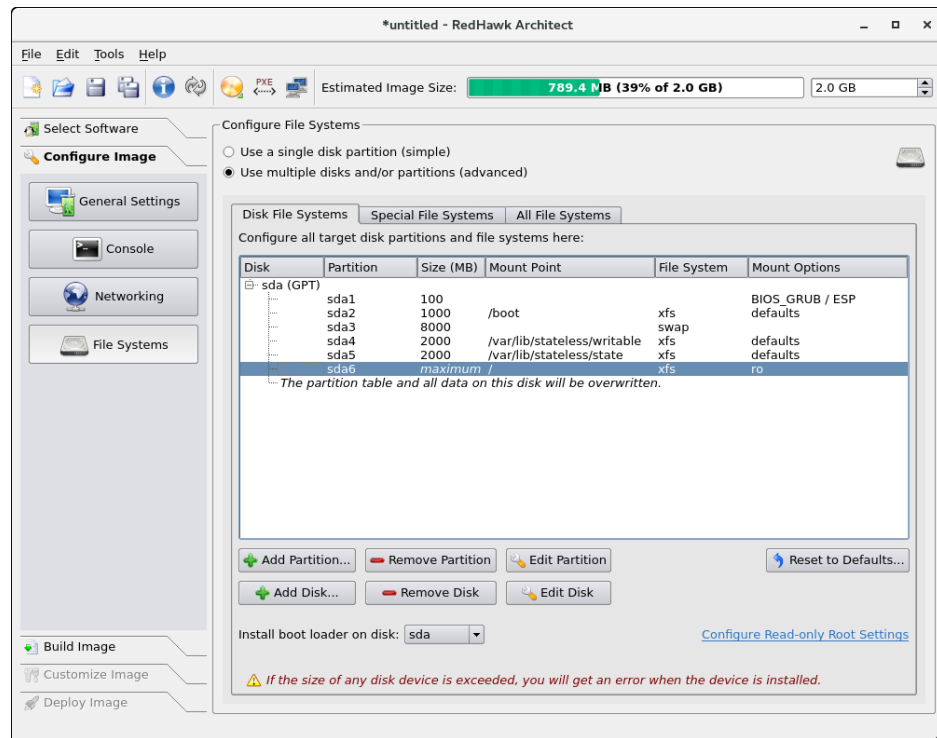
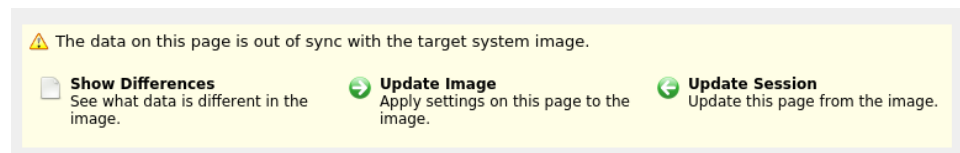


Figure 1-22 Example root read-only partitioning scheme

Note that if a change is made to the file-system settings after the target file-system image has been built, an *Out-of-Sync Notice* will appear at the bottom of the page, as shown in the figure below.



The Special File Systems tab is used to configure special (non-disk) file systems. Initially this page is blank but in the following figure below two example entries have been added; one of type **tmpfs** and one of type **bind**. See the **mount(8)** man page for more information on these special file system.

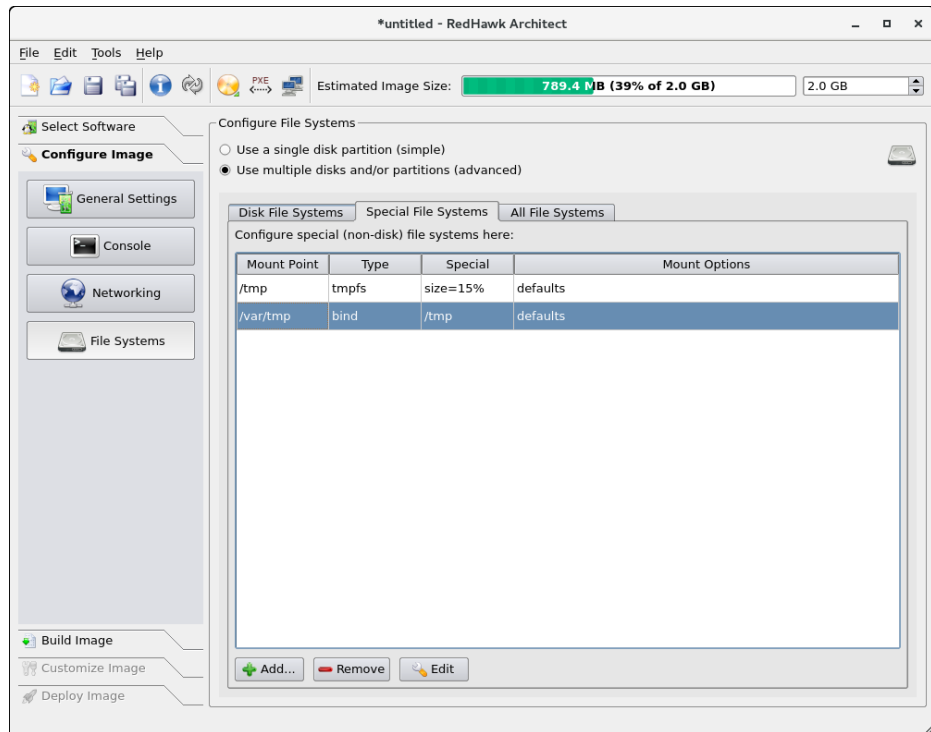


Figure 1-23 Example Special File Systems page entries

The All File Systems tab is used to view all the file systems to be configured on the target system. Both disk and special file system entries are listed. The following figure shows entries corresponding to the examples above.

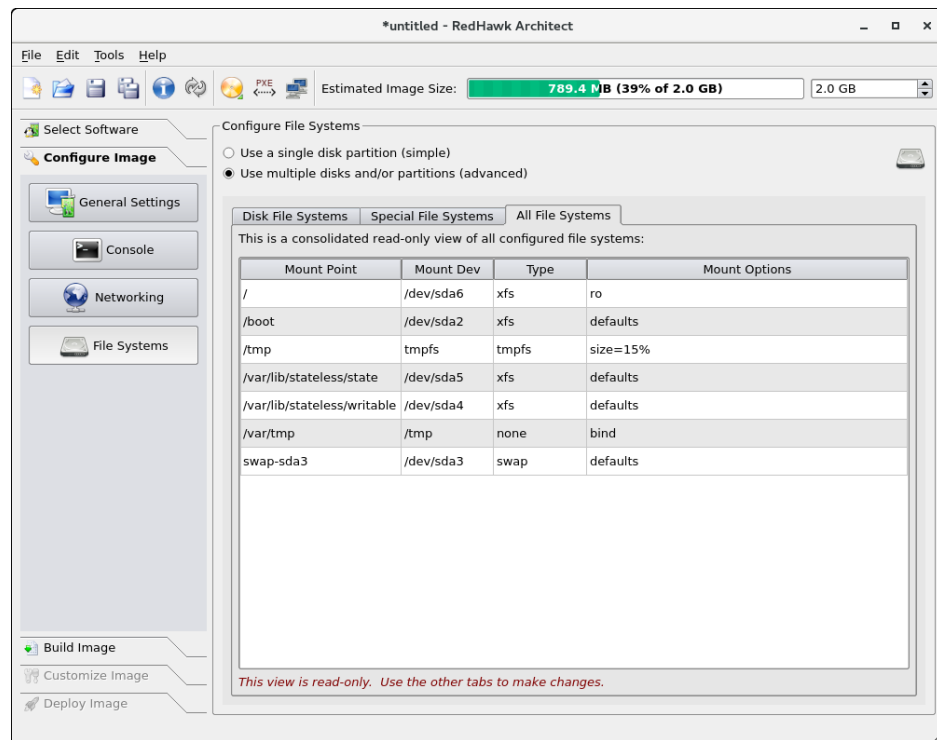


Figure 1-24 Example All File Systems list

Building an Image

To build the file system image by installing the selected software, select **Build Image** from the toolbox on the left side of the RedHawk Architect main window. The Build Image page shown in the following figure displays.

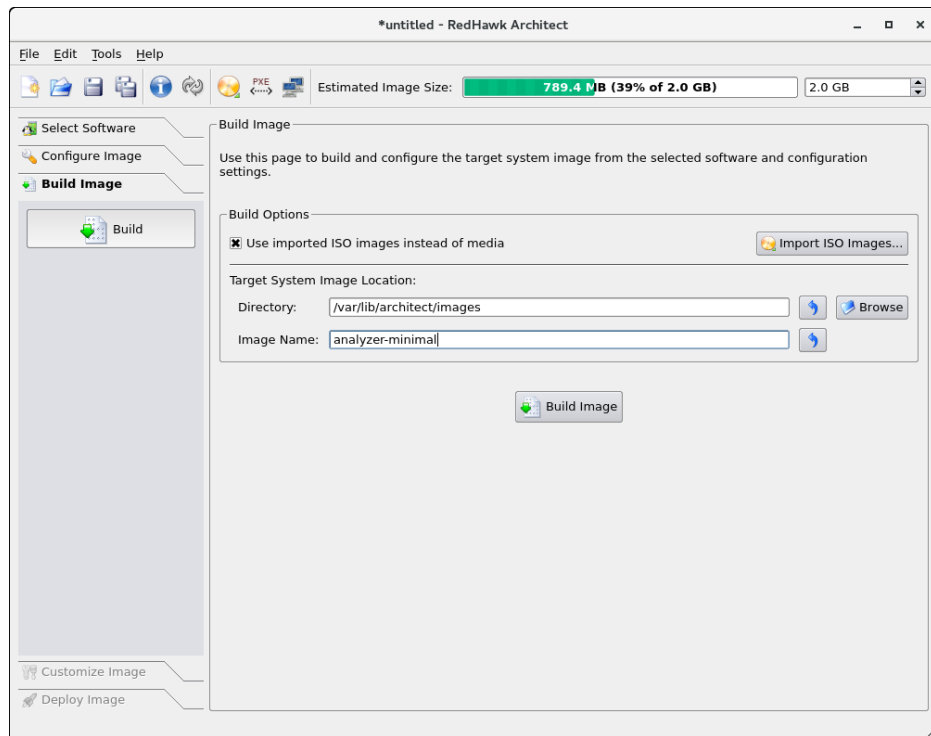


Figure 1-25 Build Image Page

Choose a directory in which to build the file system image and enter it in the **Directory** field, or click on the **BROWSE** button to display a file browser from which to choose.

NOTE

Do not use **/tmp** as the target directory. Packages like “tmp-watch” might remove files that have not been accessed in a certain number of days, thereby sabotaging the image directory.

Choose a name for the file system image and enter it in the **Image Name** field.

NOTE

Make sure that the directory you specify has enough free disk space to hold one or more file system images, each of which can be several gigabytes in size.

Click on the **Build Image** button to begin the build process. The rest of this section assumes that you have *not* previously imported the ISOs from their respective media by clicking on the **Import ISO Images...** button or selecting **Media ISO Manager** in the **Tools** menu. Advanced users may wish to do that to avoid inserting

DVD or CD media repeatedly. See “Chapter 3: Importing ISO Images” for more information.

Dialogs are presented to guide you through the process of installing the software into the image. For example, you will be prompted to insert various DVD or CD media, as shown in the following figure. Follow the directions to load the media, then click **OK** to begin.



Figure 1-26 Build Prompt to Insert CentOS Updates Media

When **OK** is selected, the CentOS installation begins. The **Build Image** screen overlays the RedHawk Architect main window and tracks the progress, as shown in the following figure.

Clicking **Abort** at any time in the build process aborts the build. A confirmation message then displays and you must click on the **Close** button to close the message box and reactivate the RedHawk Architect main window.

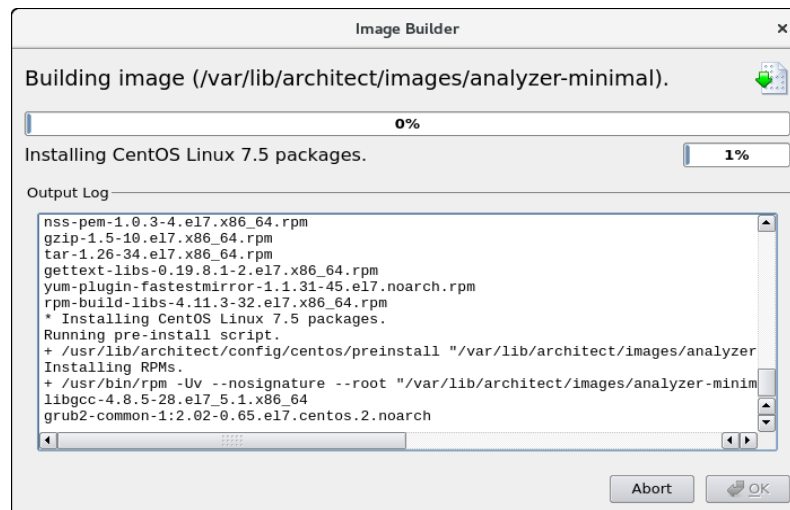


Figure 1-27 Status of CentOS Installation

An overall progress bar at the top of the **Image Builder** screen shows the progress of the entire build; the entire build will be complete once this progress bar is full.

The current stage of the build is listed immediately underneath the overall progress bar, along with a smaller stage-specific progress bar; the current stage of the build will be complete once this stage-specific progress bar is full, and it will reset for the next stage.

An **Output Log** status area in the lower half of the dialog shows the detailed output that is generated during the entire build, including any error messages generated by the

build process. Note that critical error messages result in pop-up error dialogs with which you can interact.

Abort

Click on this button to abort the build. A confirmation dialog displays allowing you to confirm or decline aborting the build process.

OK

Once the build is completed or aborted, click on the OK button to close the Build Image screen and reactivate the RedHawk Architect main window.

When the CentOS installation is complete, the dialog shown in the following figure appears.




Figure 1-28 Build Prompt to Insert RedHawk Media

Load the RedHawk Linux media, then click OK. The RedHawk installation begins and the Image Builder screen tracks the progress, as shown in the following figure.

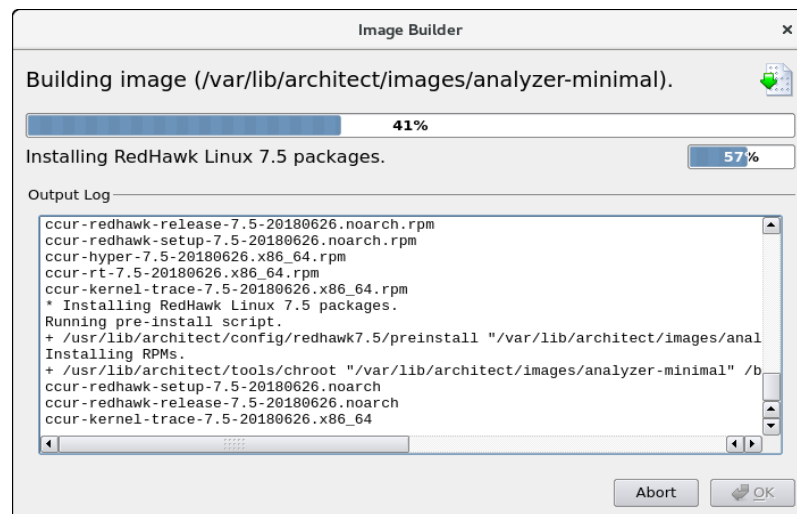


Figure 1-29 Status of RedHawk Installation

The same steps as above will be repeated for any optional software selected during the **Select Software** step; a prompt will ask the user to insert the product disc and the software will be installed in the target's build image.

Customizing an Image

To further customize the file system image, select **Customize Image** from the toolbox on the left side of the **RedHawk Architect** main window. This allows you to customize the following groups:

- Software Updates
- System Services
- Kernel Manager
- Additional RPMs
- File Manager
- Chroot Shell
- Image Cleanup

Each of these customizations will be fully described in the following sections.

NOTE

Image customizations are *not* saved in the session, and will not be automatically re-applied to future images built from the session.

Software Updates

To install RedHawk and NightStar updates into the file system image, click on **Software Updates** in the **Customize Image** toolbox. The **Software Updates** page appears, as shown in the following figure.

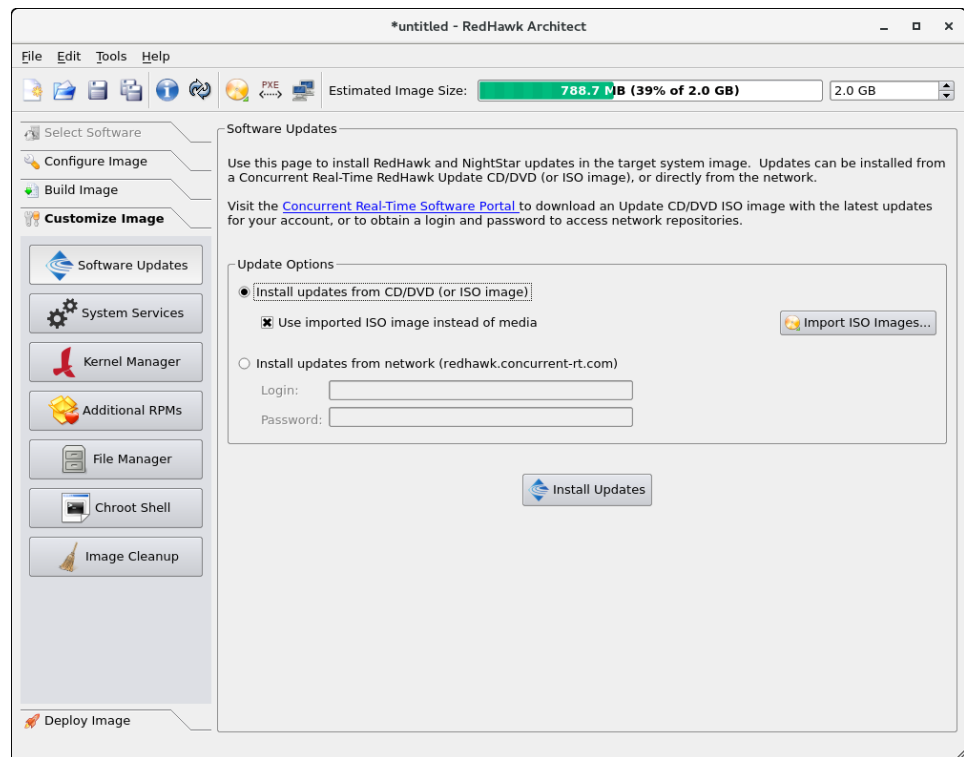


Figure 1-30 Software Updates Page

Installing updates can be done from local media (DVDs or ISO images) or directly over the network if the host system is connected to the Internet.

Select **Install updates from CD/DVD (or ISO image)** if you wish to use local media; then press the **Install Updates** button and you will be prompted to insert the media.

Select **Install updates from network instead of media** if you wish to download updates via the Internet. You will need to enter your site's assigned login and password to be granted access to the RedHawk Updates repositories, and you will also need an active maintenance subscription.

Follow the instructions as they are presented. You should see something similar to the following dialog displayed once all updates have been successfully installed.

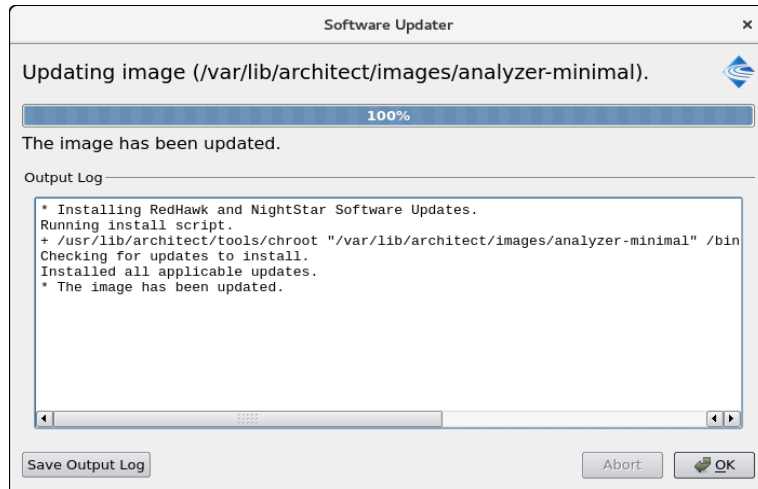


Figure 1-31 Software Updater Dialog

System Services

To customize the settings of the system services that are present in the target image, click on System Services in the Customize Image toolbox. The System Services page appears, as shown in the following figure. Note that the actual list of system services shown depends on the set of packages installed in the target image.

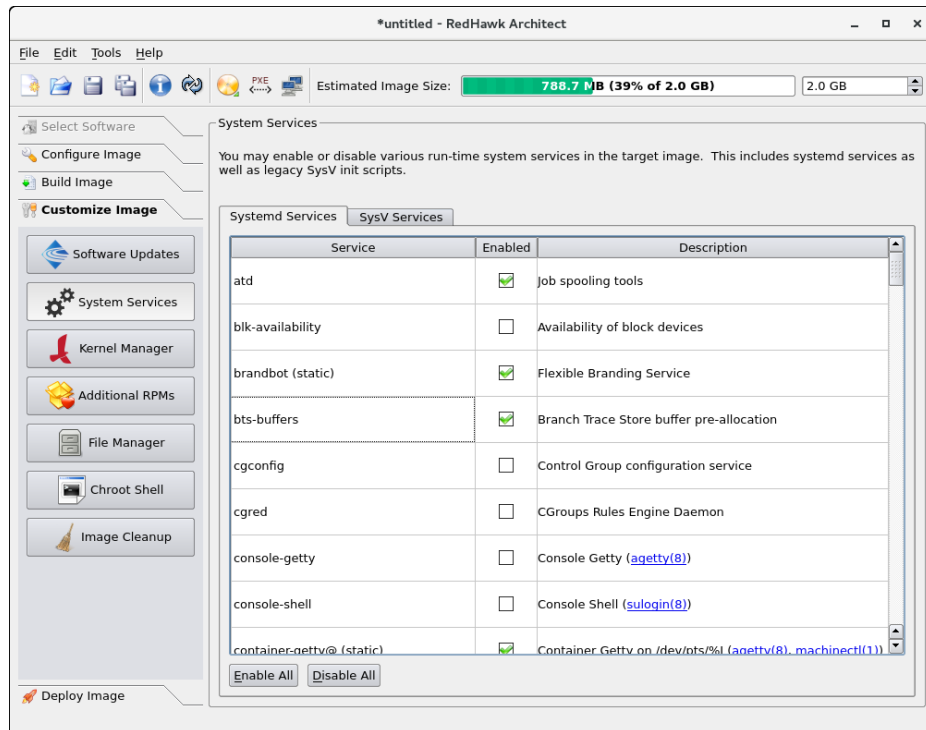


Figure 1-32 System Services Page

There are tabs for both modern `Systemd Services` and also legacy `SysV Services`. Only the system services that are actually present in the built target image are available for customization on the `System Services` page.

Note that any changes made on the `System Services` page take effect in the target image immediately.

Kernel Manager

By default you can choose to boot a standard RedHawk kernel in your target image. However you may also wish to customize the kernel to include additional components or possibly to exclude existing components. In order to customize the kernel, you must select the RedHawk kernel source software when configuring the target.

To customize kernel settings for the target image, click on `Kernel Manager` from the `Customize Image` toolbox. The `Kernel Manager` page appears, as shown in the following figure.

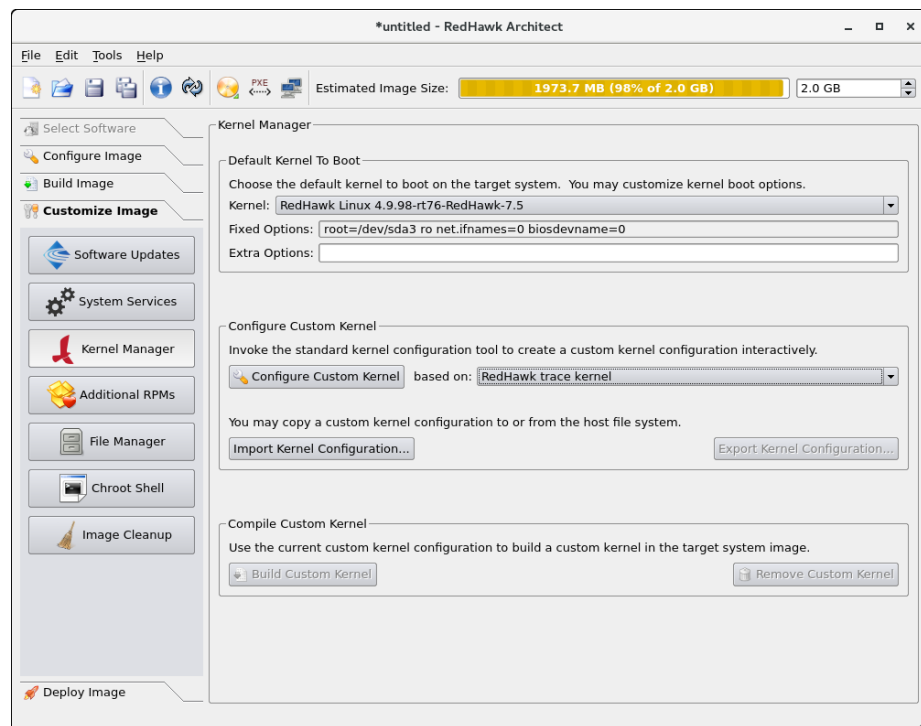


Figure 1-33 Kernel Manager Page

The `Kernel Manager` allows you to perform different functions with the kernel configuration in the target image.

The `Kernel` pull-down menu allows you to choose which installed kernel should be the default kernel that boots in the target image. Any change made to this setting is customized in the target image immediately.

The Fixed Options text area displays the required kernel boot options for the selected kernel; these kernel boot options are fixed and may not be changed by the user.

The Extra Options text field displays optional kernel boot options for the selected kernel; these kernel boot options are fully customizable by the user.

The Configure Custom Kernel area contains functions related to configuring custom kernels for the target image. The Compile Custom Kernel area contains functions related to building custom kernels for the target image. These functions will be described in the following sections.

Note that only one custom kernel configuration, and therefore one custom kernel, can be associated with a specific target image at any given time.

Configure Custom Kernel

The Configure Custom Kernel button begins the process of creating a custom kernel configuration. The custom kernel configuration is based upon the kernel configuration that is selected in the drop-down menu that is immediately to the right of the Configure Custom Kernel button.

The choices in the drop-down menu are: RedHawk standard kernel, RedHawk trace kernel, RedHawk debug kernel and Custom kernel (available once a custom kernel configuration has been imported or configured). The first three create new configurations based on the configurations of the standard RedHawk kernels.

The Custom kernel choice bases the new configuration on the current custom kernel configuration that is associated with the image; thus, the Custom kernel choice can be used to further customize a configuration that you have already customized or imported.

Pressing the Configure Custom Kernel button will bring up two different dialog windows. The first dialog window displays overall configuration progress status, as shown in the following figure.

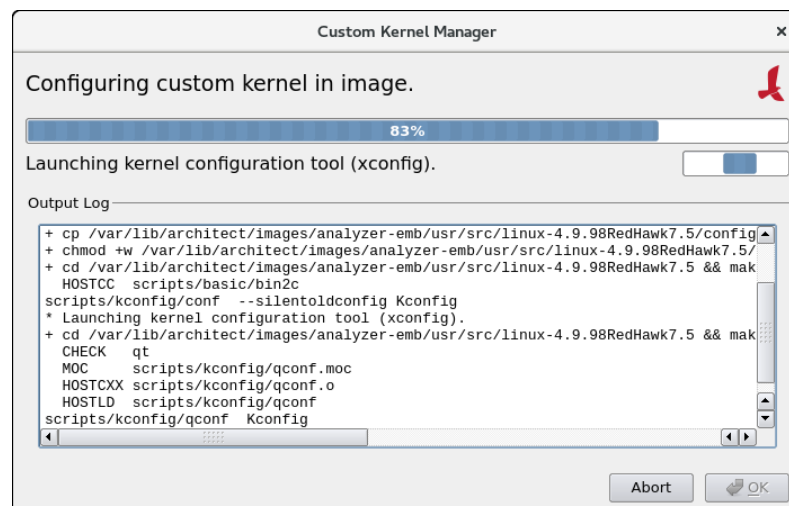


Figure 1-34 Custom Kernel Dialog

This window shows the status of running the `ccur-config` command in the target image kernel source directory. The `ccur-config` command will eventually bring up the Linux Kernel Configuration window to customize the kernel, as shown in the following figure.

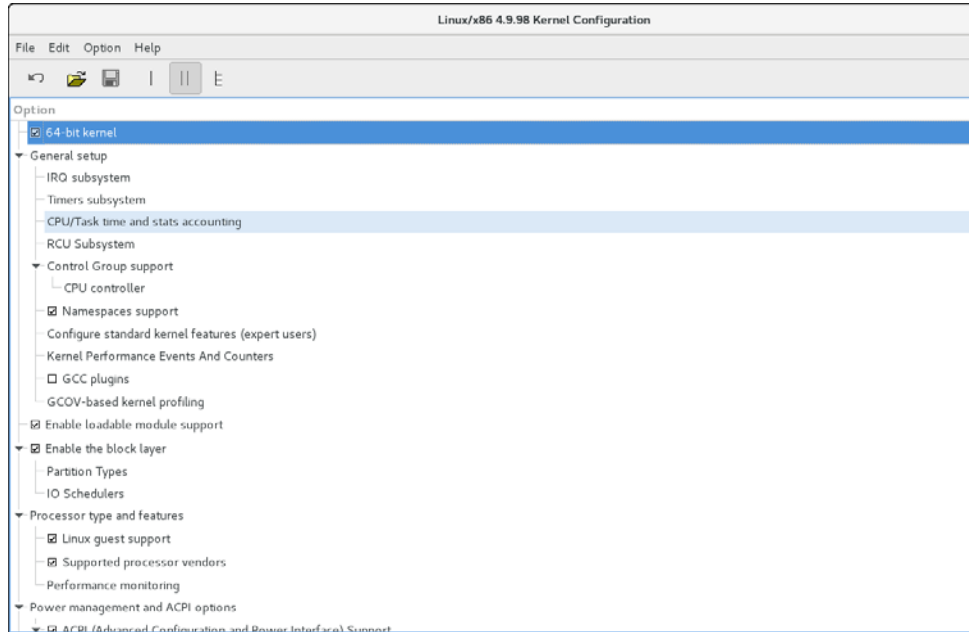


Figure 1-35 Linux Kernel Configuration Dialog

This window allows you to customize almost any aspect of the custom kernel configuration. It is expected that users who are performing this step have a thorough understanding of Linux kernel configuration.

Note that you must **Save** the kernel configuration before you exit the Linux Kernel Configuration window. Failure to **Save** the configuration will result in an error being displayed in the Custom Kernel Manager dialog window and no changes to the custom kernel configuration will be made.

NOTE

Certain compilation related RPMs must be installed on the host system in order to successfully configure and build a custom kernel (e.g. `make`, `gcc`). If any of these RPMs are missing you will be presented with a dialog detailing which RPMs must first be installed on the host system before you can proceed.

Import Kernel Configuration

The Import Kernel Configuration button allows you to choose a Linux kernel configuration file on the host system and import it to become the custom kernel configuration in the target image.

Note that once a custom kernel configuration has been imported you can further customize it by using the **Configure Custom Kernel** button and selecting the Custom kernel to base the configuration on.

Export Kernel Configuration

The **Export Kernel Configuration** button allows you to copy the target's current custom kernel configuration to the host system.

Compile Custom Kernel

The **Build Custom Kernel** button allows you to build and install a complete custom kernel in the target image. You must first have created a custom kernel configuration, either by using the **Configure Custom Kernel** button or by using the **Import Kernel Configuration** button.

Building a custom kernel compiles each file that comprises the Linux kernel and this process can take quite a bit of time to complete. Once you start the process, you will see the **Custom Kernel Manager** dialog appear, as shown in the following figure and it will describe the entire process.

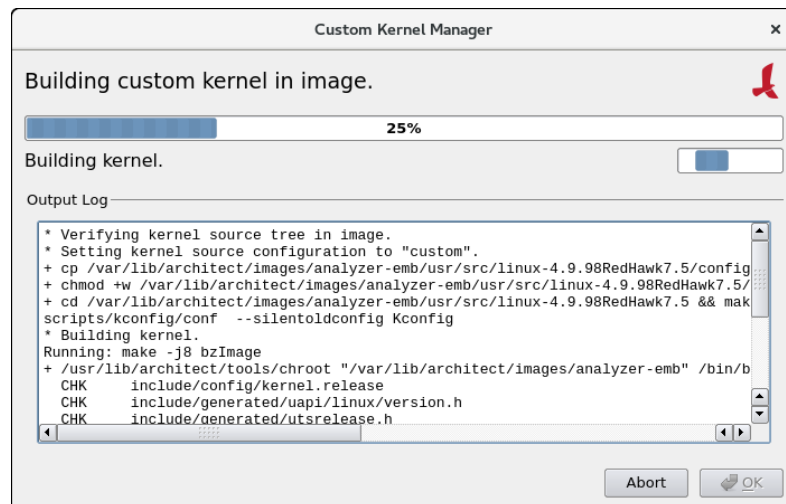


Figure 1-36 Initial Build Progress

Initially **ccur-config** will be invoked and once that completes the kernel build stages will begin, as shown in the following figure.

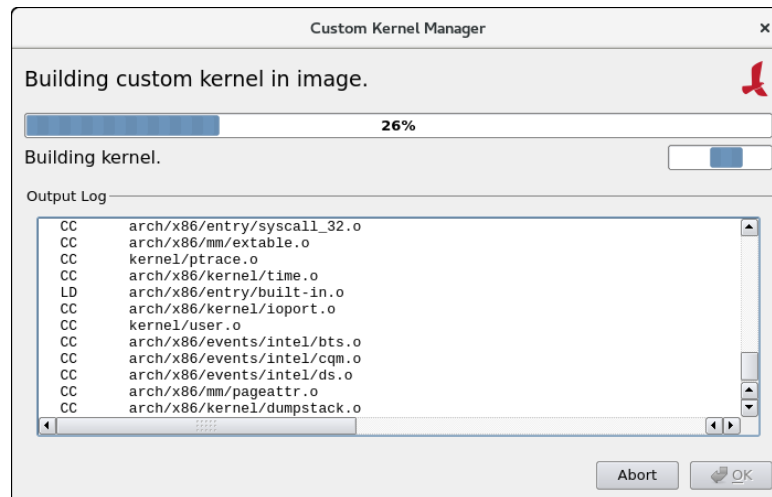


Figure 1-37 Kernel Build Stages

Finally, once the entire build and install process is complete, the kernel source tree will be cleaned to free up the temporary space used to build the kernel. At this point the entire build process will be complete, as shown in the following figure.




Figure 1-38 Kernel Build Complete

The custom kernel automatically becomes the default kernel to boot. If this choice is not desired, change the kernel to boot using the **Kernel To Boot** area as described above.

Remove Custom Kernel

The **Remove Custom Kernel** button allows you to remove the current custom kernel from the target image. This will remove the entry in **grub.conf** as well as all of the associated kernel files in the image.

Note that the custom kernel configuration itself is not removed. Thus, it is still possible to build a custom kernel based on the current custom kernel configuration that still remains in the target image.

Additional RPMs

To install additional RPMs into the target file system image manually, click on **Additional RPMs** from the **Customize Image** toolbox. The **Install Additional RPMs** page appears, as shown in the following figure.

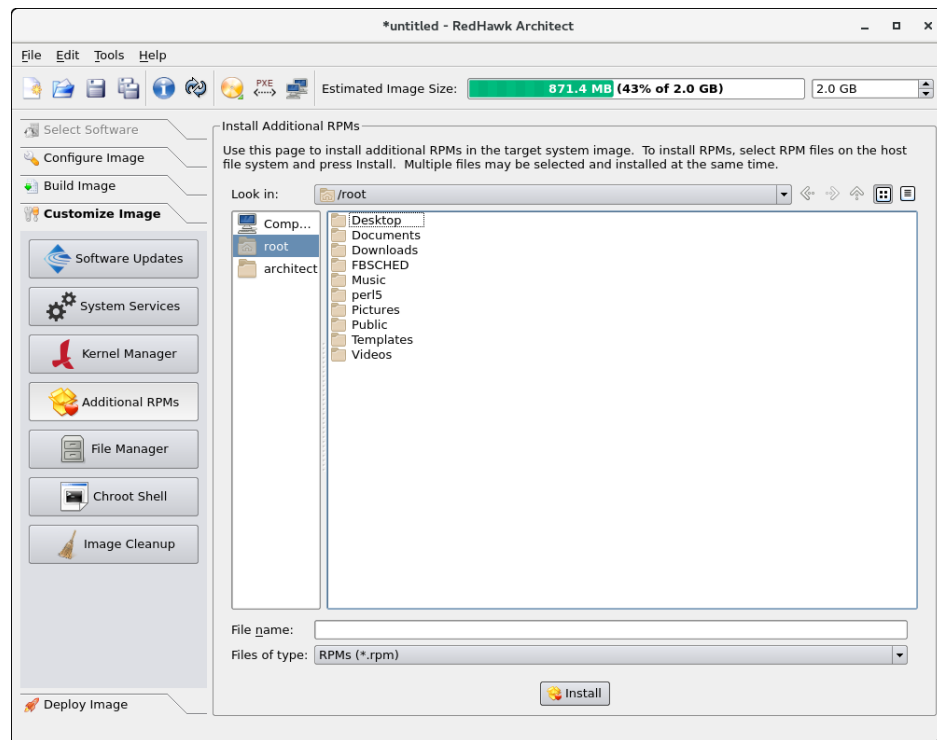


Figure 1-39 Install Additional RPMs Page

The **Install Additional RPMs** page can be used to locate RPM files on the host system and then easily install them into the target file system image. Note that the interface supports multiple selection; if you have a set of RPMs that have dependencies upon each other you will need to select all of the RPMs simultaneously to have them properly installed together into the target file system image.

Installing Board Support Packages

Concurrent Real-Time provides *Board Support Packages* (BSPs) for several supported SBCs. These BSPs are distributed as RPMs that may be installed in an image using the **Additional RPMs** page as described above. Contact Concurrent Real-Time (support@concurrent-rt.com or 1-800-245-6453) for information on how to obtain BSPs for a particular SBC.

File Manager

To copy various files into the target file system image manually, click on File Manager from the Customize Image toolbox. The File Manager page appears, as shown in the following figure.

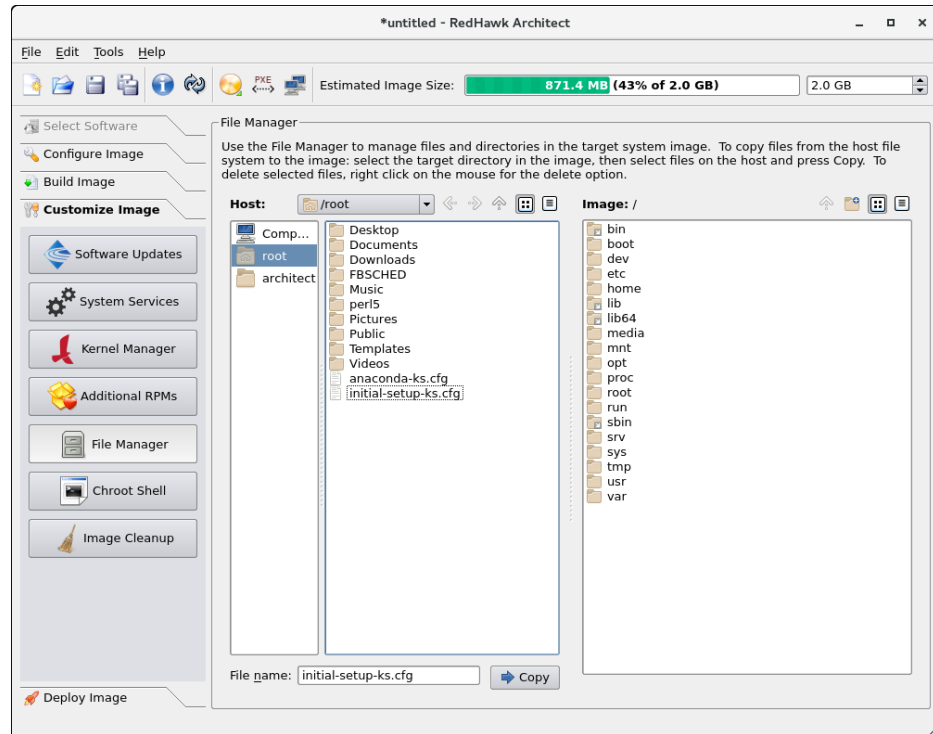


Figure 1-40 File Manager Page

The File Manager page supports many features including multiple selection, the ability to create new directories in the target file system image, and the ability to delete files in the target system image.

Chroot Shell

To customize the target file system image manually, click on the Chroot Shell from the Customize Image toolbox. The Chroot Shell page appears, as shown in the following figure

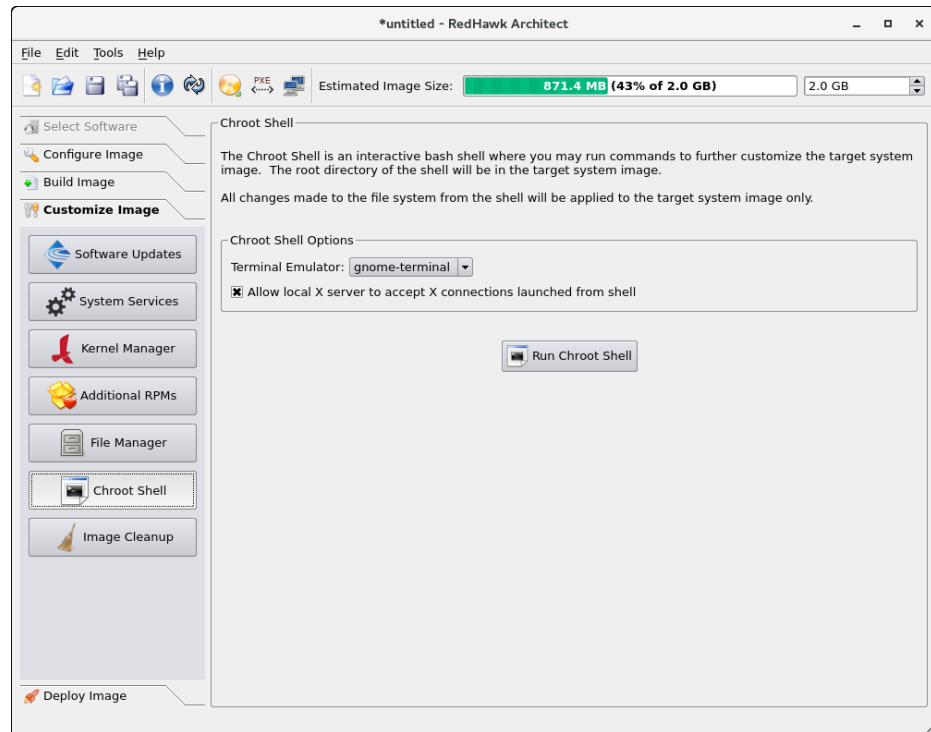


Figure 1-41 Chroot Shell Page

From this toolbox you can open a “chroot” shell in a terminal window. Select the type of terminal from the dropdown and click on the Run chroot shell button. A terminal screen opens, as shown in the following figure.

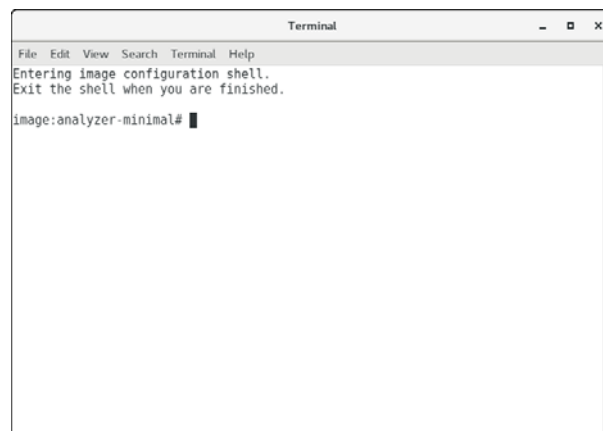


Figure 1-42 chroot Shell

This provides a shell with the root directory being the file system image directory. All changes made to system files (including software installed or removed) will be done in the file system image directory only. The host's root file system will not be affected.

Exit the shell when changes are complete.

Image Cleanup

You may reduce the size of the file system image by removing various types of files that may be unnecessary for the image. To remove unnecessary files from the image, click on Image Cleanup in the Customize Image toolbox. The Image Cleanup page appears, as shown in the following figure.

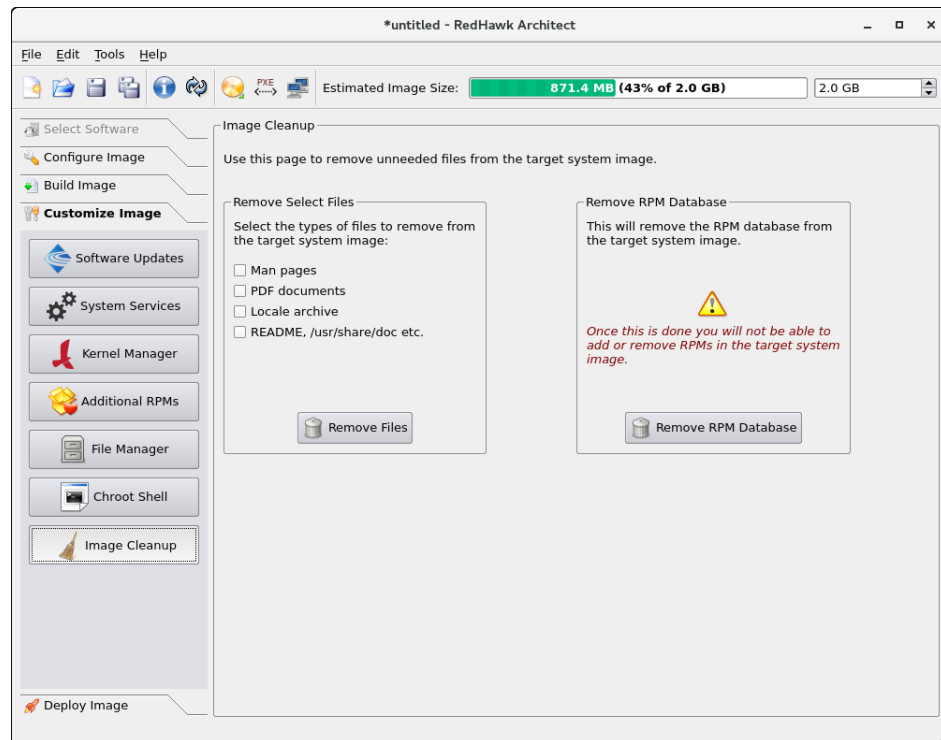


Figure 1-43 Image Cleanup Page

Select the types of files to remove from the file system image and click the **Remove Files** button.

To remove the RPM database from the file system click the **Remove RPM Database** button. Once this is done you will lose all ability to manage RPMs in the image. This cannot be undone. Only do this once you are sure you do not have to add or update any more RPMs in the image.

Deploying an Image

Target root file system images can be deployed onto target boards in several different ways with RedHawk Architect.

- USB devices can be directly flashed with the root file system image. This includes USB drives and also CompactFlash cards in CompactFlash-to-

USB adapters. These devices can then be inserted into a target board and the board will boot the image upon a restart. See “Deploy to USB Device” on page 1-39 for more information.

- A USB drive installer can be created with the root file system on it. Architect creates a bootable installation USB drive that will boot on the target and install the root file system onto the target board's local media. Once complete, the USB drive is removed and the board will boot the image upon a restart. See “Install via USB drive” on page 1-42 for more information.
- DVD media installers can be created with the root file system image on it. Architect creates a bootable installation DVD that will boot on the target and install the root file system onto the target board's local media. Once complete, the DVD is removed and the board will boot the image upon a restart. See “Install via DVD media” on page 1-43 for more information.
- RedHawk Architect can deploy the root file system image over the network. It can deploy an installer that will install the root file system image onto the target board's local media, or it can deploy the root file system image via NFS for fully diskless booting. See “Installing via PXE over a Network” on page 1-45 for more information on network installation, and see “Booting Diskless via PXE over a Network” on page 1-47 for more information on diskless booting.
- RedHawk Architect can deploy the root file system image directly to a virtual machine image that can be booted via QEMU. See “Deploy to Virtual Machine” on page 1-54 for more information.

In addition to deploying root file system images to target boards, Architect also supports deploying root file system images to *virtual machine images*, which can be booted in virtual machines running directly on the host. Using this feature, it is possible to test target system images without the use of any target hardware.

The UEFI firmware target configuration is currently supported by all deployment methods with the exception of the DVD Installer method. In the USB Device and Virtual Machine deployment methods, the Configure for UEFI firmware box must be set if the intended target system utilizes UEFI firmware. In the other deployment methods (PXE and USB Installer), there is no UEFI configuration box as those methods work with either UEFI or BIOS systems.

Deploy to USB Device

To copy a target root file system image to a USB device, select **Deploy Image** from the toolbox on the left side of the RedHawk Architect main window and click on the **USB Device** button. This will display the **Deploy to USB Device** page, as shown in the following figure.

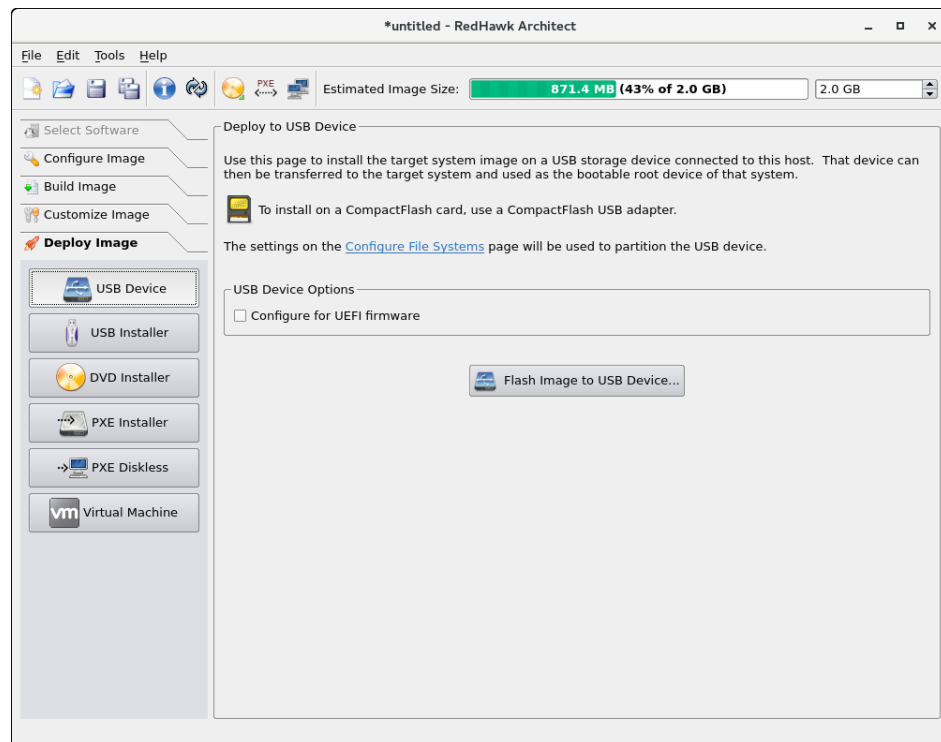


Figure 1-44 Deploy to USB Device Page

The Flash Image to USB Device... button allows you to copy a root file system image onto a USB flash device (e.g. a standard USB Flash Drive or a CompactFlash that is connected directly to the host machine via a USB-to-CompactFlash adapter). Note that IDE/SATA CompactFlash adapters are not supported at this time.

Make sure to select the Configure for UEFI firmware check box if the intended target system utilizes UEFI firmware.

NOTE

CompactFlash devices and USB drives can be bought inexpensively at many retail stores that sell computer accessories. Note that the duration of the flashing process depends upon the performance rating of the specific CompactFlash device or USB drive. It is recommended to use CompactFlash devices or USB drives that have a minimum of a 40MB/s read/write performance rating.

Pressing the Flash Image to USB Device... button will begin copying the target root file system onto the USB device. The host system will be scanned for attached USB flash storage devices. If multiple devices are found a choice will be presented to the user, otherwise the sole device found will be selected by default. Once a device is found or chosen, a confirmation dialog similar to the following will appear:



Figure 1-45 Flash Device Confirmation

Press OK to confirm the operation and then the copy will begin, as shown in Figure 1-46.

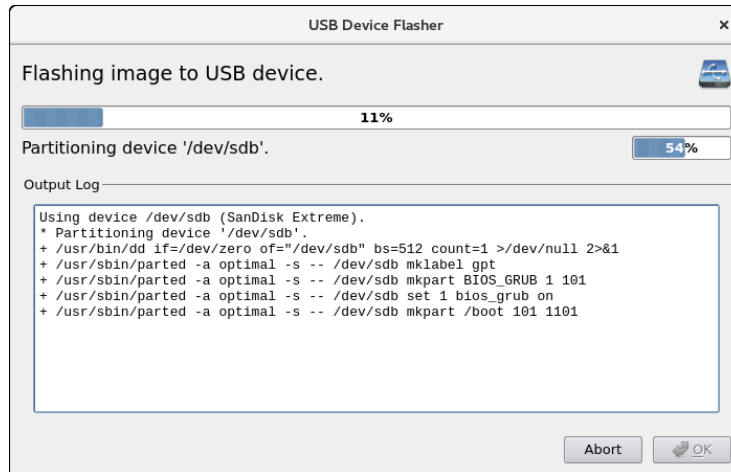


Figure 1-46 Flash Copy In Progress

Note that no initial check is made to determine whether the image will fit onto the size of the selected USB device. If the copy fails because of insufficient space, an error message will be displayed, as shown in the following figure.




Figure 1-47 Flash Dialog Error

If the USB device is large enough to hold the image, and no other error occurs during the copy, a success dialog will be presented, as shown in Figure 1-49.




Figure 1-48 Device Removal Notification

Remove the USB device if desired and then click **OK** to continue. You will now be presented with a final dialog indicating that the transfer is complete.



Figure 1-49 Flash Copy Completed

Once the copy has completed successfully, the USB device can then be placed onto the intended target board and the board can be reset to boot into a fresh RedHawk installation.

Install via USB drive

To create a bootable USB drive that will install the target root file system image into a target system, select **USB Installer** from the toolbox on the left side of the RedHawk

Architect main window. This will display the Deploy via Installation USB Drive page, as shown in the following figure.

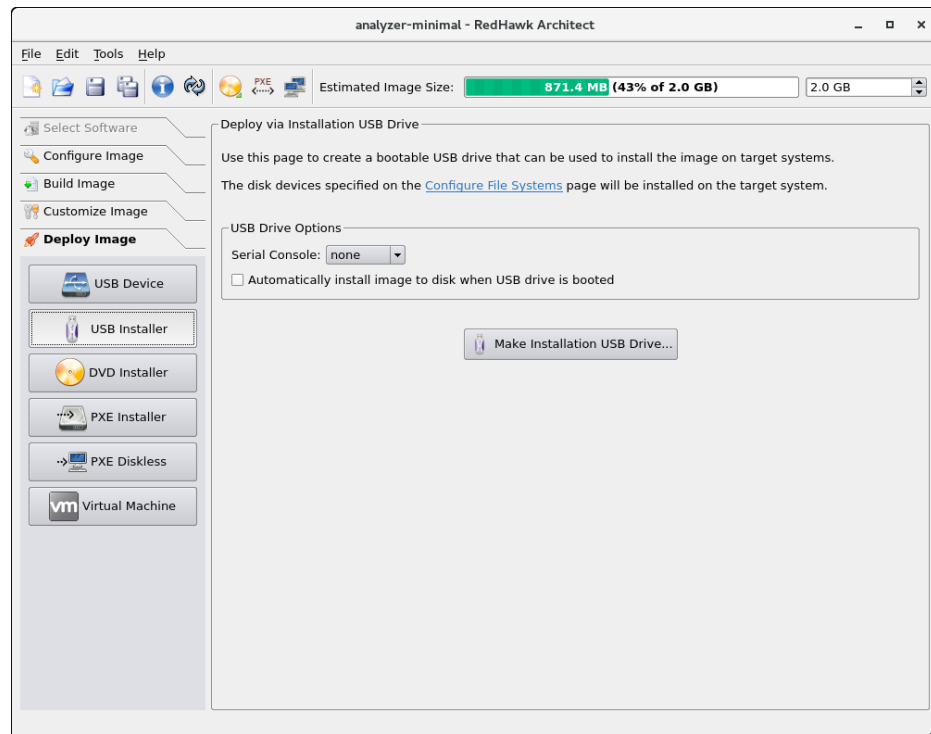


Figure 1-50 Deploy via Installation USB Drive Page

Press the **Make Installation USB Drive...** button to write a bootable installer image to an attached USB drive.

Choose the **Serial Console** setting that the target system will use for communicating with the host. If set to **none** the target will default the console to the VGA display.

Choose **Automatically install image to disk when USB drive is booted** to create a USB drive that will install the target root file system image onto a target board's local media without any prompting or user interaction.

NOTE

This will destroy any data on the target system's local media whenever booted into a target system and therefore this should be used carefully; however, it is useful on systems without an attached console display or configured and connected serial console.

Install via DVD media

To create a bootable DVD media that will install the target root file system image into a target system, select **DVD Installer** from the toolbox on the left side of the RedHawk

Architect main window. This will display the Deploy via Installation DVD page, as shown in the following figure.

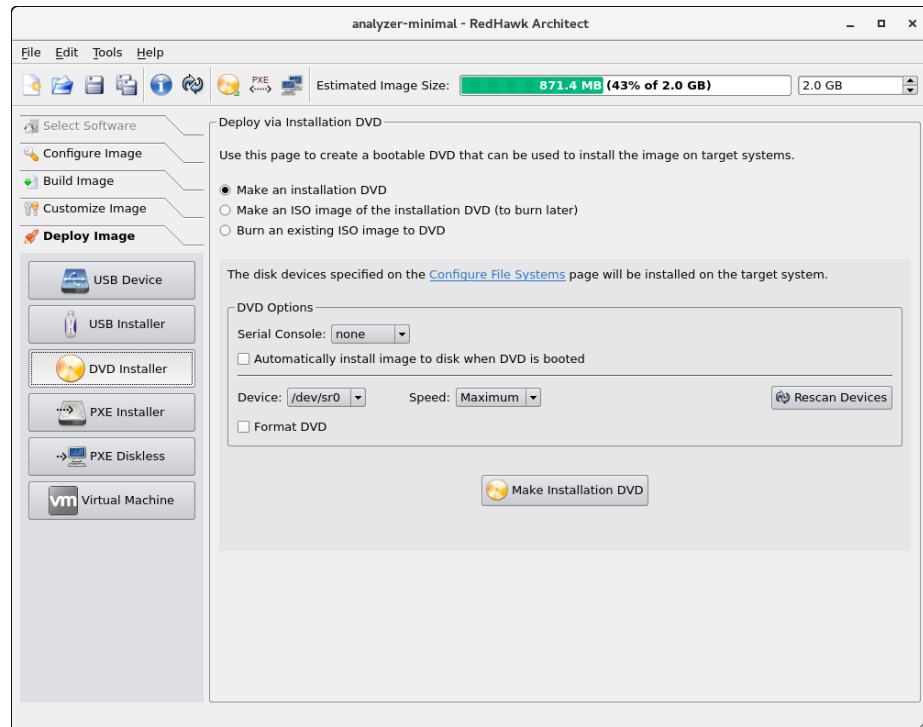


Figure 1-51 Deploy via Installation DVD Page

Choose **Make one installation DVD** to directly burn a DVD that will install the target root file system image onto DVD media. No ISO image will be saved on disk in this mode.

Choose **Make an ISO image for the installation DVD** to create an ISO file that contains an installer image. This ISO image can be later burned to DVD, or it may be useful with other tools or for long term storage.

Choose **Burn an existing ISO image to DVD** to burn a previously created ISO image to DVD media.

Depending on which operation mode is chosen, various options will be available for selection. Choose the options and settings that are appropriate for your specific needs.

NOTE

Targets running UEFI firmware are currently not supported via the DVD deployment method.

Installing via PXE over a Network

RedHawk Architect can deploy a root file system image to a target system over an Ethernet network connecting the host machine to the target machine. Installation of the root file system is performed by first creating a PXE-bootable installation image. A target machine can boot this installation image via PXE, which will then remotely copy the root file system image into the target's local drive media.

This deployment method does not require the preparation of any removable installation media and it is often the fastest installation deployment method, however it does require some initial networking configuration on both the host and target systems.

NOTE

Various host system networking services must be properly configured before attempting to deploy a PXE-bootable installation image for the first time. If you have not configured the host networking services yet, you will need to invoke the PXE Target Manager and choose to Initialize PXE Services. See "Managing PXE Targets" on page 4-7 for more information.

To create a PXE-bootable installation image that will install the target root file system image over a network into a target system, select PXE Installer from the toolbox on the left side of the RedHawk Architect main window. The Deploy via PXE Installation page will then appear, as shown in the following figure.

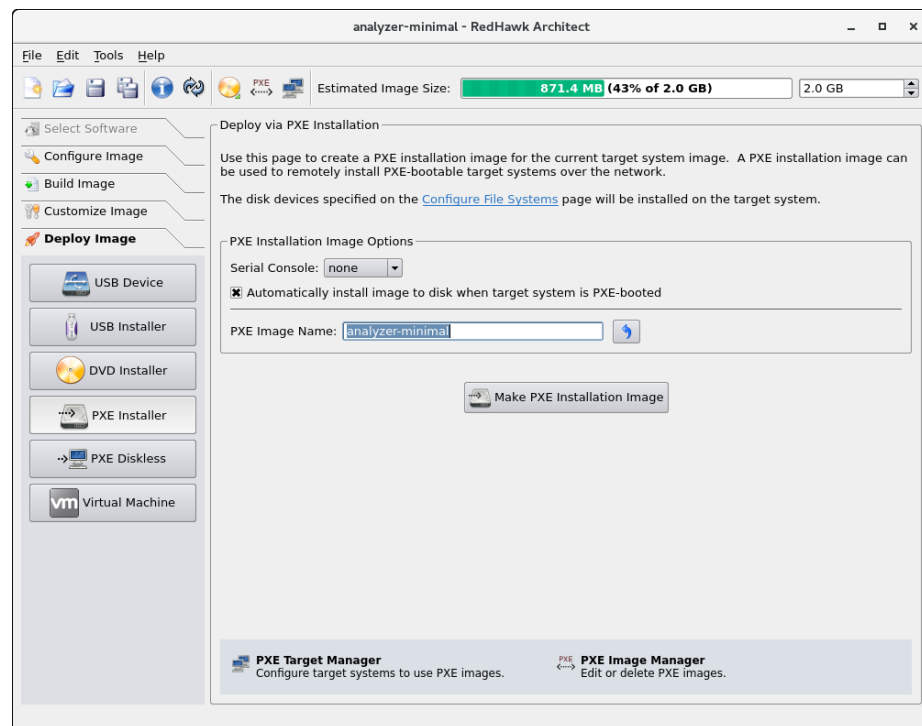


Figure 1-52 Deploy via PXE Installation Page

Choose the **Serial Console** setting that the target system will use for communicating with the host. If set to **none** the target will default the console to the VGA display.

Verify that the **Automatically install image to disk when target system is PXE-booted** checkbox is checked to have the target perform installation non-interactively. If this checkbox is not checked then the target will first display a menu and a user must press a key before the installation will begin.

Enter a **PXE Image Name** for the installation image that will be created. Each installation image must have a unique name to identify it, though the names can be arbitrarily chosen by the user. Multiple images may be created and shared between targets. See “Managing PXE Images” on page 4-3 for more information.

Press **Make PXE Installation Image** to begin building the named PXE installation image. The PXE Installation Image Builder dialog will appear as shown in the following figure.

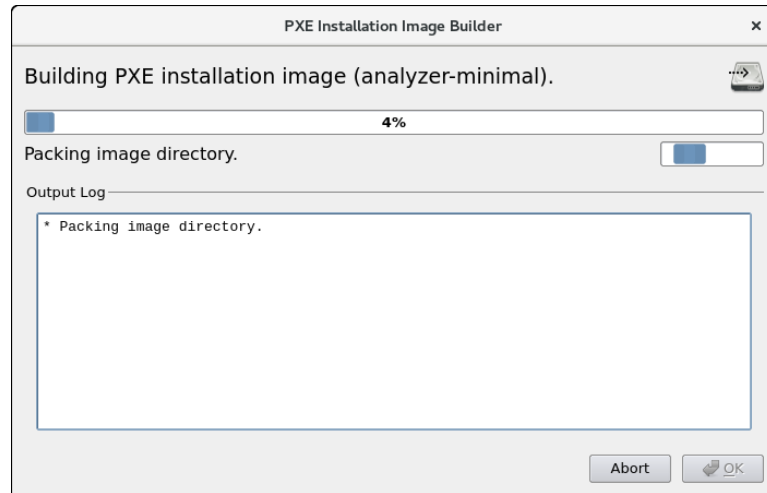


Figure 1-53 PXE Installation Image Builder Dialog

NOTE

The PXE installation images are placed under a directory named **architect** which must reside under the system's **tftpboot** directory. The **tftpboot** directory defaults to **/var/lib/tftpboot**. While this directory is configurable, at this time the Architect tool only supports the default location.

Packing the PXE installation image will take several minutes. Once complete the PXE Installation Image Builder dialog will appear as shown in the following figure.

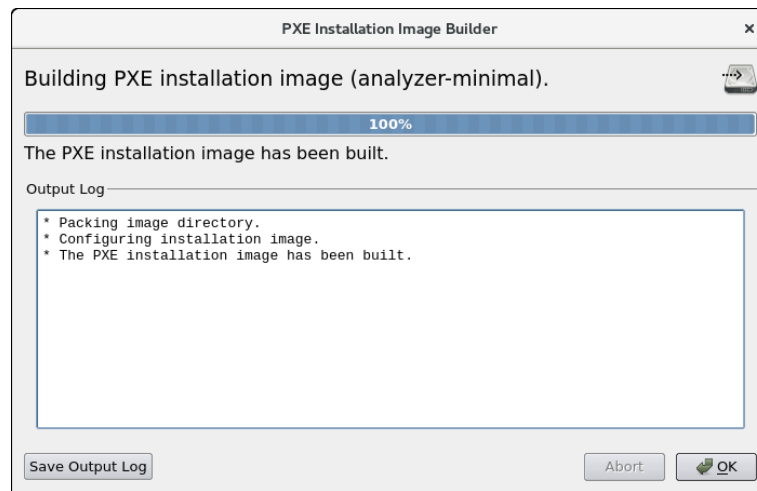


Figure 1-54 PXE Installation Image Building Complete

Press **Okay** to dismiss the dialog.

Once the PXE installation image is successfully built you can use the **PXE Target Manager** to schedule installation of the image for specific targets, and you can use the **PXE Image Manager** to edit or delete PXE installation images. See “Managing PXE Targets” on page 4-7 for more information.

Booting Diskless via PXE over a Network

RedHawk Architect can create and then deploy a PXE-bootable diskless image to a diskless target system. Host and target connect over an Ethernet network. This deployment method does not require any local drive media to be present on the target system; any local drive media that is present on the target will be untouched and ignored. This method also requires some initial networking configuration on both the host and target systems. Note that the file system configuration for this deployment method is custom and ignores the settings in the **File Systems** configuration page.

There are two different implementations for booting diskless. The first option uses NFS, the second a Live RAMDISK. With NFS, the target machine boots a diskless image via PXE, which will then mount the root file system image via NFS. In a Live RAMDISK boot, the entire root file system is downloaded to the target's RAM.

NFS versus Live RAMDISK considerations:

- **Persistent Storage:** with the NFS option, the kernel mounts the root file system read-only over NFS but the user may optionally configure persistent storage via the **Configure Read-only Root Settings** link explained below. With the Live RAMDISK option, the entire root file system is writeable but volatile.
- **Network Connection:** with the NFS option, the host and target must maintain an Ethernet working connection for the duration of the time the target is booted. With the Live RAMDISK, the connection is required only for booting.

- **Boot time and RAM Allocations:** with the NFS option, the read-only root file system is accessed via NFS although some system directories that require to be writeable are RAM-based and volatile. With the Live RAMDISK option the entire SquashFS root file system is downloaded and copied to RAM during the boot.

To create a PXE-bootable diskless image that will mount the target root file system image via NFS on a target system, select **PXE Diskless** from the toolbox on the left side of the RedHawk Architect main window and then select on the **Make an NFS diskless image** radio button at the top of the page. The **Deploy to Diskless Systems** page will then appear as shown in the figure below.

NOTE

Various host system networking services must be properly configured before attempting to make an NFS diskless installation image and before booting a RAMDISK diskless installation image. If you have not configured the host networking services prior, you will instead be presented with a page allowing you to **Initialize PXE Services**. See “Managing PXE Targets” on page 4-7 for more information. Once PXE Services have been initialized you will be allowed to continue.

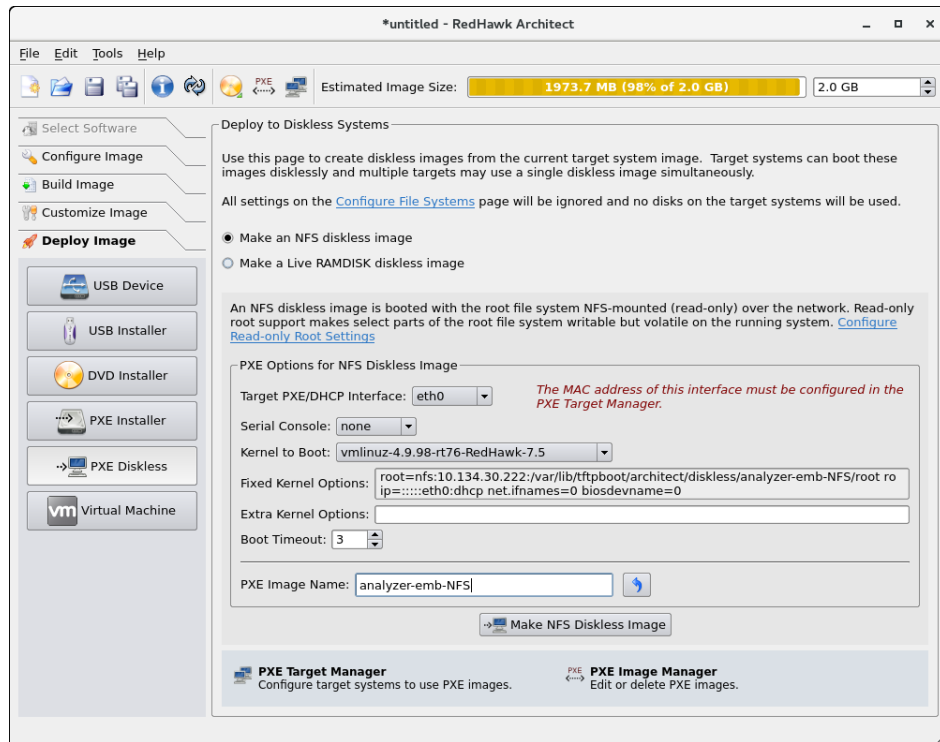


Figure 1-55 Initial PXE NFS Diskless Deployment Page

When the Make a Live RAMDISK diskless image radio button is selected, a different but similar page will appear as shown in the figure below.

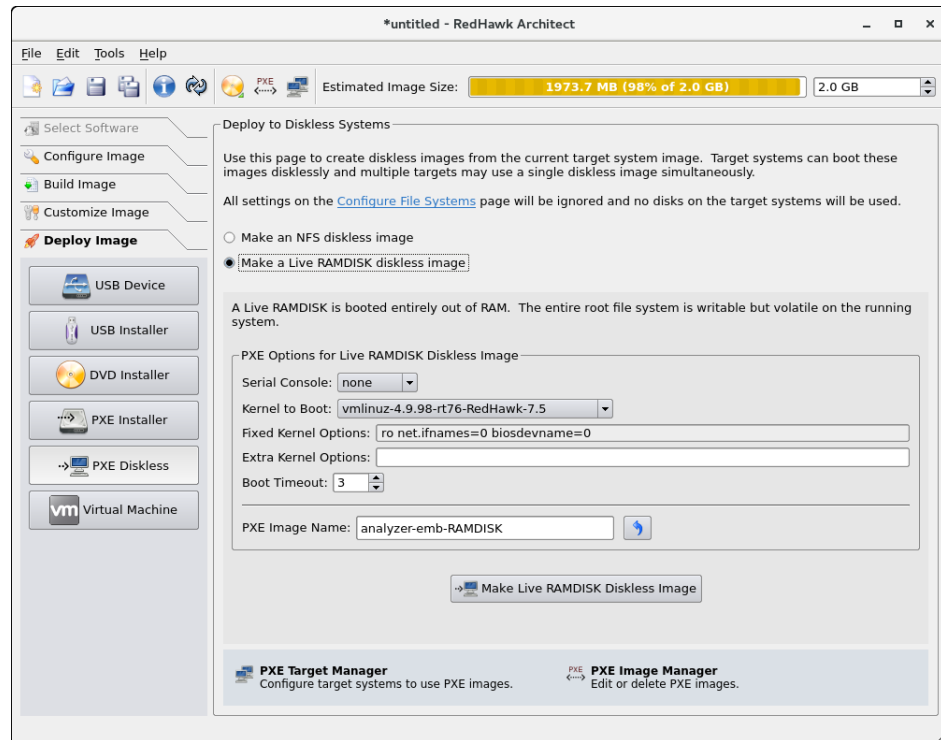


Figure 1-56 Initial PXE RAMDISK Diskless Deployment Page

The following settings are common to both creating an NFS and a RAMDISK bootable diskless image, with the exception of the **Configure Read-only Root Settings** which pertains only to making NFS diskless images.

Choose the **Target PXE/DHCP Interface** network interface that the target system will use for communicating with the host. The target hardware must be configured to perform a PXE broadcast on this network interface at boot time.

Choose the **Serial Console** setting that the target system will use for communicating with the host. If set to **none** the target will default the console to the VGA display.

Choose the **Kernel to Boot** for the target. This will default to the kernel that has already been chosen as the default in the **Kernel Manager**, however a diskless image may specify a different default if desired.

The **Fixed Kernel Options** text area displays the required kernel boot options for the selected kernel; these kernel boot options are fixed and may not be changed by the user.

Enter any **Extra Kernel Options** that you would like to use for the diskless image. All kernel parameters specified here will be appended to the kernel's boot-time options. See the **kernel-parameters.txt** file in the kernel source **Documentation** directory for a complete list of the standard boot options.

Modify the `Boot Timeout` count to change the number of seconds the boot menu will be displayed before the diskless image will start booting. Increase the timeout if you wish to have more time to interrupt the boot menu to choose different kernels or boot options.

The `Configure Read-only Root Settings` pertains only when making NFS diskless images but not RAMDISK diskless images. Click on this link to adjust the size of RAM space allotted for temporary storage space. Use the up and down arrows to change the default size. Persistent storage, private to each target, can be accessed on the target under `/var/lib/stateless/state` and on the nfs server under `/var/lib/tftpboot/clientstate/<target-system>`.

Enter a `PXE Image Name` for the diskless image that will be created. Each diskless image must have a unique name to identify it, though the names can be arbitrarily chosen by the user. Multiple images may be created and shared between targets. See “Managing PXE Images” on page 3-4 for more information.

When choosing to create an NFS Diskless Image, Press `Make NFS Diskless Image` to begin building the named PXE diskless image. The PXE Diskless Image Builder dialog will appear as shown in the following figure.




Figure 1-57 PXE NFS Diskless Image Builder Dialog

Creating the PXE NFS diskless image should take several minutes. Once complete the PXE Diskless Image Builder dialog will appear as shown in the following figure.

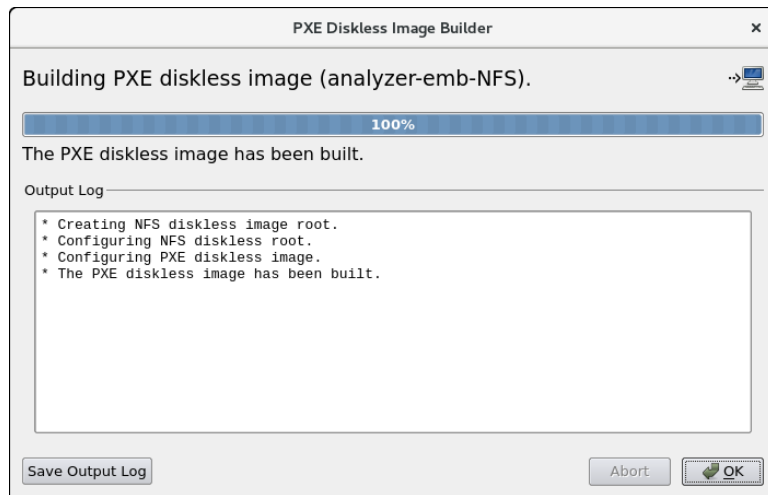


Figure 1-58 PXE NFS Diskless Image Building Complete

Press **OK** to dismiss the dialog.

When choosing to create a **RAMDISK** Diskless image, Press **Make Live RAMDISK Diskless Image** to begin building the named PXE diskless image. The PXE Diskless Image Builder dialog will appear as shown in the following figure.

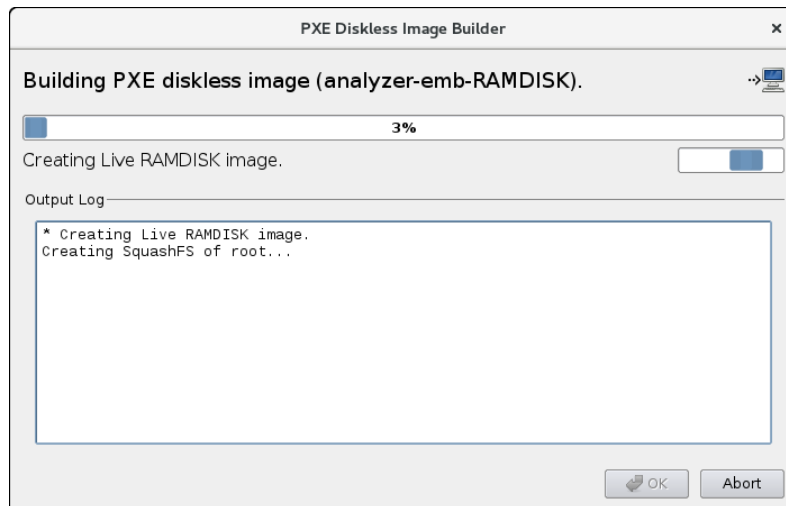


Figure 1-59 PXE RAMDISK Diskless Image Builder Dialog

Creating the PXE RAMDISK Diskless image should take several minutes. Once complete the PXE Diskless Image Builder dialog will appear as shown in the following figure.

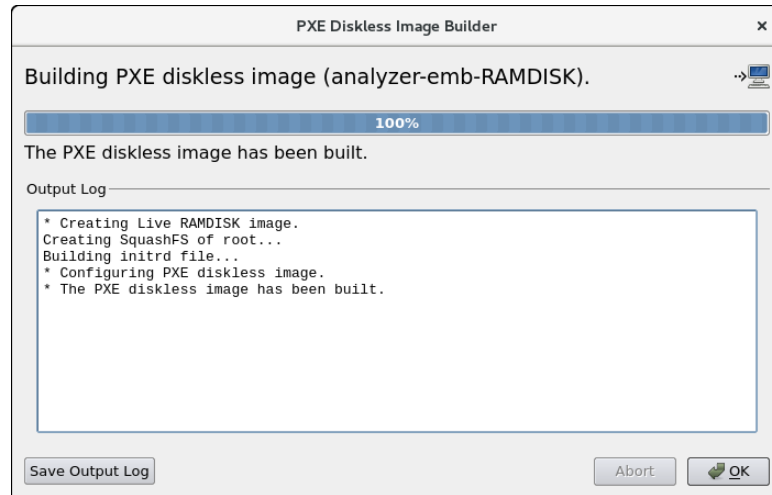


Figure 1-60 PXE RAMDISK Diskless Image Building Complete

Press **Okay** to dismiss the dialog.

NOTE

The PXE diskless images are placed under a directory named `architect` which must reside under the system's `tftpboot` directory. The `tftpboot` directory defaults to `/var/lib/tftpboot`. While this directory is configurable, at this time the Architect tool only supports the default location.

Once the PXE diskless image is successfully built you can use the PXE Target Manager to configure diskless booting of the image for specific targets, and you can use the PXE Image Manager to edit or delete PXE diskless images. See "Managing PXE Targets" on page 4-7 for more information.

Deploy to Virtual Machine

To deploy a target root file system image to a virtual machine image which can be booted in a virtual machine, click on the **Virtual Machine** button. This will display the Deploy to Virtual Machine page, as shown in the following figure.

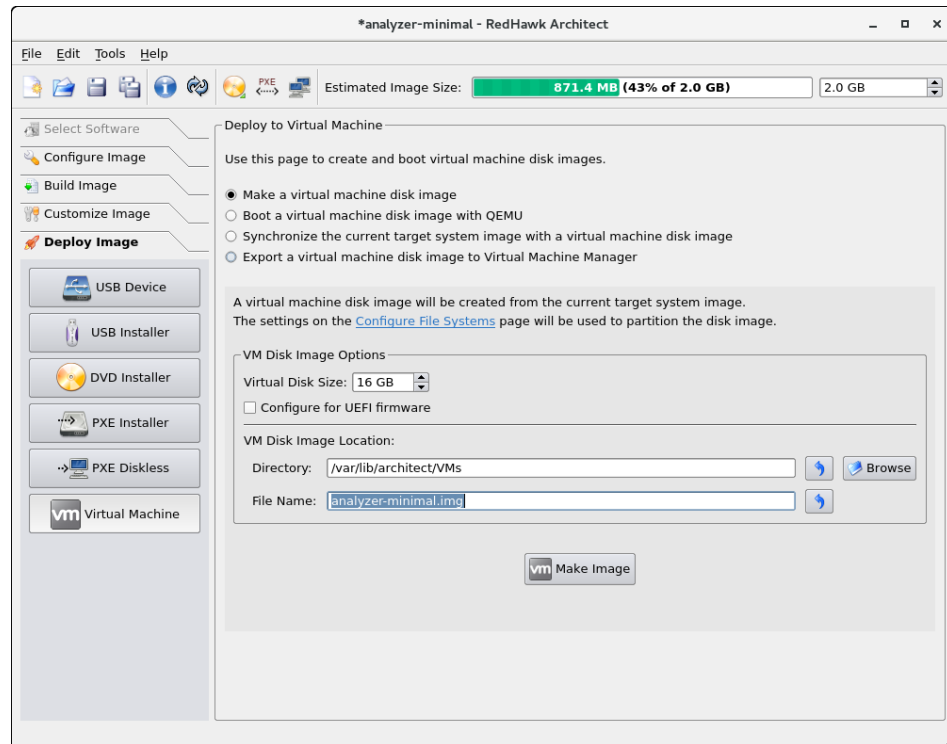


Figure 1-61 Deploy to Virtual Machine Page

The Deploy to Virtual Machine Page defaults to the **Make a virtual machine disk image** radio button selection. In this mode, pressing the **Make Image** button will simply create a virtual machine image file from the session's root file system image. The name and location of the virtual machine image file created can be customized using the **Directory** and **File Name** text fields, and the **directory Browse** button.

Make sure to select the **Configure for UEFI firmware** check box in the **VM Disk Image Options** section if the intended target system utilizes UEFI firmware.

Selecting the **Boot a virtual machine disk image in a QEMU virtual machine** radio button allows you to boot a previously created virtual machine image directly on the host using the **QEMU PC System Emulator**. Choose the virtual machine image using the **VM Disk Image to Boot** text field or the **file Browse** button.

Selecting the **Synchronize the current target system image with a virtual machine disk image** radio button allows you to perform file synchronization between the target system image and the virtual machine disk image in both directions:


- Choose **Update files** in the current target system image (to match VM disk image) and press the **Sync Image** button to import into the target system image all file changes that have been made inside the booted virtual machine image.
- Choose **Update files** in the VM disk image (to match the current target system image) and press the **Sync Image** button to export all file changes that have been made in the target system image into the virtual machine disk image. The exported changes will be visible in the virtual machine disk image the next time it is booted using QEMU.


These two synchronization features provide additional flexibility for customizing a target system image. Image customization can also be accomplished inside a booted virtual machine, and this customization is very natural as the environment closely resembles the final booted environment that will be available on the actual target hardware.


Selecting **Export a virtual machine disk image to Virtual Machine Manager** allows you to utilize a previously created virtual machine image with very flexible and powerful virtual machine management tools provided on the host. Once the image has been exported, the graphical VMM tools can boot and manipulate the image completely independently of Architect. See the **virt-manager(1)** man page for more information.

Editing an Existing Session

A session can be saved at any time and loaded later to continue work on a file system image.

To save the current session click on the **Save Session** icon  or on **Save Session** in the **File** menu. Selecting **Save Session As** in the **File** menu displays a file selection dialog and allows you to save the current session using a different name.

To make a copy of the current session click on the **Duplicate Session** icon  or on **Duplicate Session** in the **File** menu. Duplicating a session makes a copy of the current session and optionally copies an existing image to go with it.

To load an existing session, click on the **Open Session** icon  or on **Open Session** in the **File** menu. You may also click on the **Open** button from the opening dialog when Architect first starts.

Security Extensions

The Advanced Security Edition of Architect includes a security extension which adds support for configuring, creating, and deploying target system images with SELinux and SCAP security policy enabled to enhance the security of deployed target systems.

This chapter explains how to use these security extensions.

Note

The security extensions of the Advanced Security Edition of Architect are provided by an optional package named **ccur-architect-security**. If this package is not installed on the system, the security extensions will not be available.

SELinux

The Advanced Security Edition of Architect includes a security extension which adds support for NSA Security-enhanced Linux (SELinux).

Configuring SELinux

The instructions for creating, building and deploying a target configured with SELinux support are the same as detailed in Chapter 1 Using RedHawk Architect with the addition of the instructions to configure SELinux that follow.

Before building the target system image, configure SELinux in the image via the [Configure SELinux](#) page. Click on the SELinux button of the [Configure Image](#) toolbox to display the [Configure SELinux](#) page. By default, SELinux is disabled in RedHawk but it can be configured in the permissive or enforcing mode in this page. See figure below.

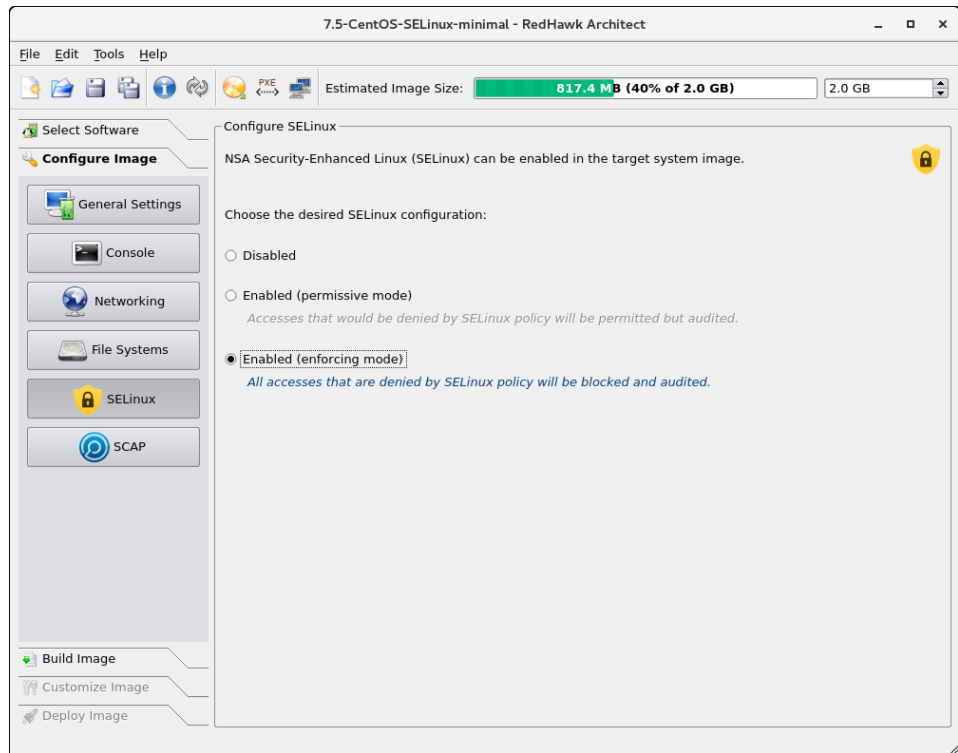


Figure 2-1 SELinux enabled in the enforcing mode

After SELinux is configured, you may proceed to follow the general instructions to build a target system image. See "Building an Image" on page 1-23.

Alternatively, SELinux can be configured after the target's image is built. After the target's image is created and following a change in the SELinux configuration, the user will be prompted to update the target system image. Click on the button labeled **Update Image** that appears at the bottom of the page to update the target's image. See figure below.

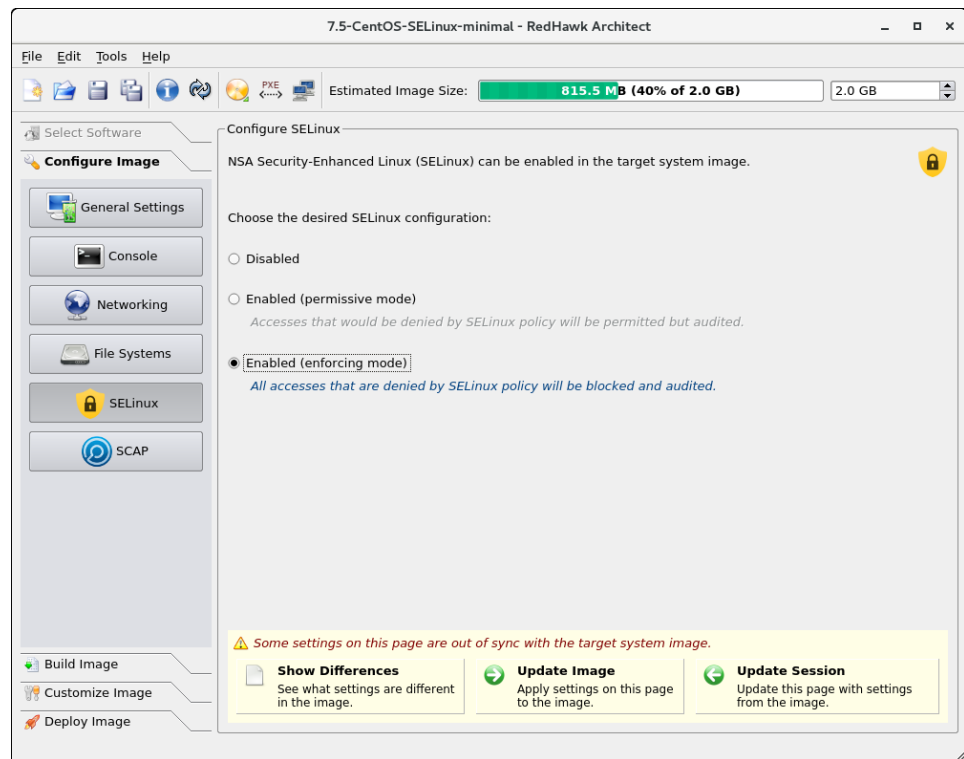


Figure 2-2 SELinux enabled after the target system image built

Note

SELinux support is not possible over NFS, therefore, the diskless deployment method that uses NFS is not supported.

Note

Optionally after the target system image is built, the target's software can be automatically updated via the Software Updates button of the Customize Image toolbox. If the `ccur-redhawk-setup` package is updated, the SELinux configuration will be reset to disabled as this is the default for RedHawk systems. This package is rarely updated but in such a case, reconfigure SELinux after the software update.

Security Content Automation Protocol (SCAP)

Introduction to SCAP

The Secure Content Automation Protocol (SCAP) was developed by the U.S. government's NIST organization to create security-oriented operating system configuration checklists.

The SCAP Security Guide implements security guidances recommended by respected authorities, namely PCI DSS, STIG, and USGCB. The SCAP Security Guide transforms these security guidances into a machine readable format referred as content files which can be used to audit your system in an automated way. The SCAP content files are provided in the **scap-security-guide** package and are installed in the directory **/usr/share/xml/scap/ssg/content/**. There are files for every platform available in the forms XCCDF (Extensible Configuration Checklist Description Format), OVAL (Open Vulnerability Assessment Language) or datastream documents. In most cases, the datastream is used, which are the file names ending with **-ds.xml**.

The SCAP Security Guide builds multiple security benchmarks and corresponding profiles from a single SCAP content. Profiles provide a set of rules to be applied. The DISA STIG RHEL 7 rules provides required settings for US Department of Defense systems. As an example, following are the corresponding SCAP settings when building a RedHawk 7.5 (RHEL Server) target system image and selecting the DISA STIG profile:

content file: **/usr/share/xml/scap/ssg/content/ssg-rhel7-ds.xml**
benchmark: Guide to the Secure Configuration of Red Hat Enterprise Linux
profile: DISA STIG for Red Hat Enterprise 7

Custom profiles can also be derived from existing profiles using the SCAP Workbench graphical tool. This is often referred to as SCAP content tailoring.

Overview of SCAP Workflow

The following steps are typically performed to configure, build, and deploy targets that adhere to some SCAP security policy:

1. Configure the desired SCAP security policy. The desired SCAP security policy is chosen prior to building a target system image. Image software and configuration settings are made to comply with the chosen security policy before the target system image is built.
2. Build target system image. A post-build SCAP remediation scan is done automatically in the target system image when the image is built.
3. Optionally run additional post-build SCAP scans. Additional SCAP scans can be run in the chroot of the target system image prior to image deployment to the target systems. These scans may be evaluation scans or remediation scans. If desired, manual remediation can also be performed at this time.

4. Run additional SCAP scans on deployed targets. These scans may be evaluation scans or remediation scans. If desired, manual remediation can also be performed at this time.

Note

It is common for some auto-remediation rules to fail when run in the chroot of a target system image; therefore an additional remediation scan and/or manual remediation is often required to be performed after target deployment. It may be possible to avoid these extra steps by doing appropriate manual remediation in the chroot of the target system image and/or using a custom SCAP tailoring file.

Understanding SCAP Evaluation and Remediation Scans

SCAP scans are performed with the `oscap(8)` tool installed by the `openscap-scanner` package. Two types of scans may be performed: evaluation and remediation. Evaluation scans report on the current security status of the system, according to a selected security profile. Remediation scans attempt to fix security discrepancies found on the system, then report on the status of this process. This process is also called "auto-remediation".

Additionally, the SCAP Workbench GUI tool can be used to perform SCAP scans. This tool has the ability to scan a remote system over an SSH connection.

Auto-remediation is not perfect. Rarely are all SCAP rules fixed by auto-remediation, therefore manual remediation or custom SCAP tailoring is often required to achieve the desired level of security policy compliance.

Manual remediation can be done by editing user-level configuration files. This can either be done in a chroot of a target system image or on target systems once deployed.

Custom SCAP tailoring can be used to modify existing SCAP policy. This is useful to change or exclude certain security rules in a profile. Custom tailoring also provides a way to codify manual remediation steps into auto-remediation scripts contained within a custom SCAP profile. Most of the tools that accept a SCAP content file as input, will optionally also allow both a SCAP tailoring file and its corresponding SCAP content file to be specified. This includes `oscap(8)`, `scap-workbench(8)`, and the Architect GUI.

Note

Some SCAP rules may be broken and do not pass evaluation no matter what you do. Always be sure to use the most up-to-date SCAP content files provided by your host distribution. Before creating target system images, update the SCAP content files on the host by using the command: `yum update scap-security-guide`. If the package is not installed, install it with the command: `yum install scap-security-guide`.

Configuring SCAP

The instructions for creating, building and deploying a target configured with SCAP are the same as the general instructions detailed in Chapter 1 Using RedHawk Architect. The configuration instructions here are additional and specific to SCAP configurations.

SCAP must be configured before the target system image is built. If an image has already been created, it must first be removed via the **Delete Image** button of the **Build Image** toolbox.

To configure SCAP, click on the **SCAP** button of the **Configure Image** toolbox. The **Configure SCAP Security Policy** page will display the default SCAP content file loaded and the corresponding benchmark(s) and profile(s) available for the corresponding RedHawk release specified when the target's session was created.

In the figure below the following default SCAP settings are observed for a session building a RedHawk 7.5 CentOS target system image. The various profiles to choose from are listed in the top window.

content file: `/usr/share/xml/scap/ssg/content/ssg-centos7-s.xml`

benchmark: Guide to Secure Configuration of Red Hat Enterprise Linux

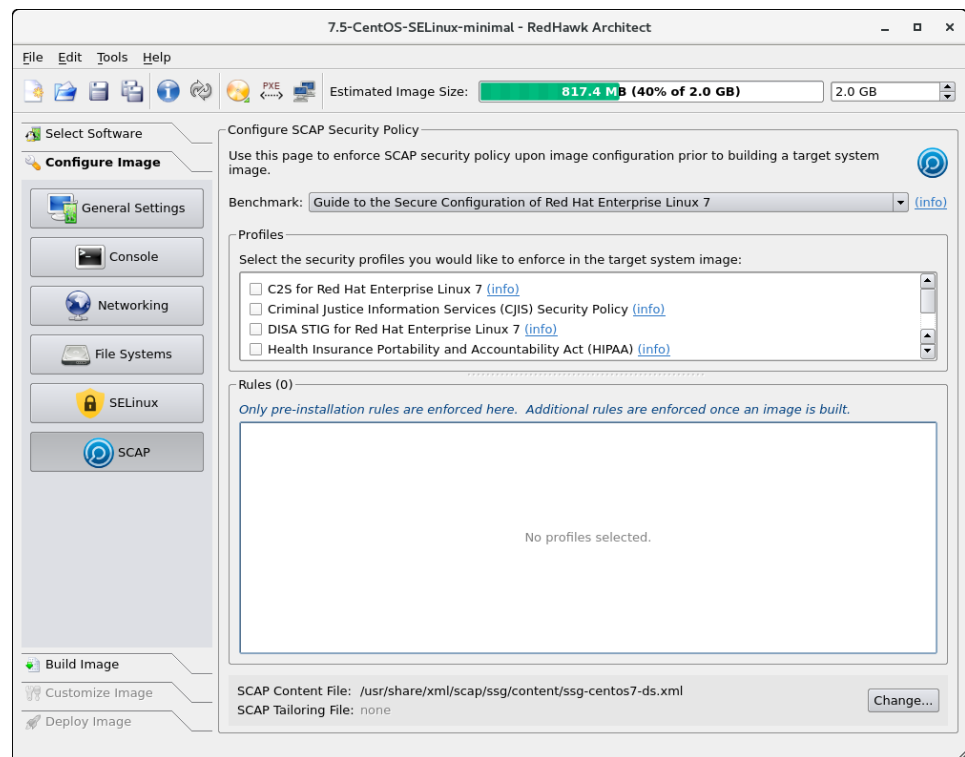


Figure 2-3 Configure SCAP Security Policy page

The default content file in the above figure provides a choice of benchmarks and profiles. The default content file used is listed in the left bottom corner of the page. Use the **Change** button on the right hand corner of the page to load a content file different from the default.

Note

To execute pre-installation SCAP rules, Architect relies on special formatting that is unique to Red Hat content files in the scap-security-guide package. If more recent SCAP content is required, additional scans can be done using newer content after the initial target system image has been built.

Click on the (info) links in the Configure SCAP Security Policy page to obtain information about the benchmark and the profiles listed. Click on the box by a profile to select that profile. Note that one or more profiles may be selected.

Once one or more profiles are selected, a list of pre-installation rules will appear in the bottom window of the Configure SCAP Security Policy page as shown in the figure below. Note that each one has an (info) link also. For each pre-install rule to be applied, there are two buttons with choices to either Fix or Ignore the rule. The target system image cannot be built until each pre-installation rule is either applied or ignored.

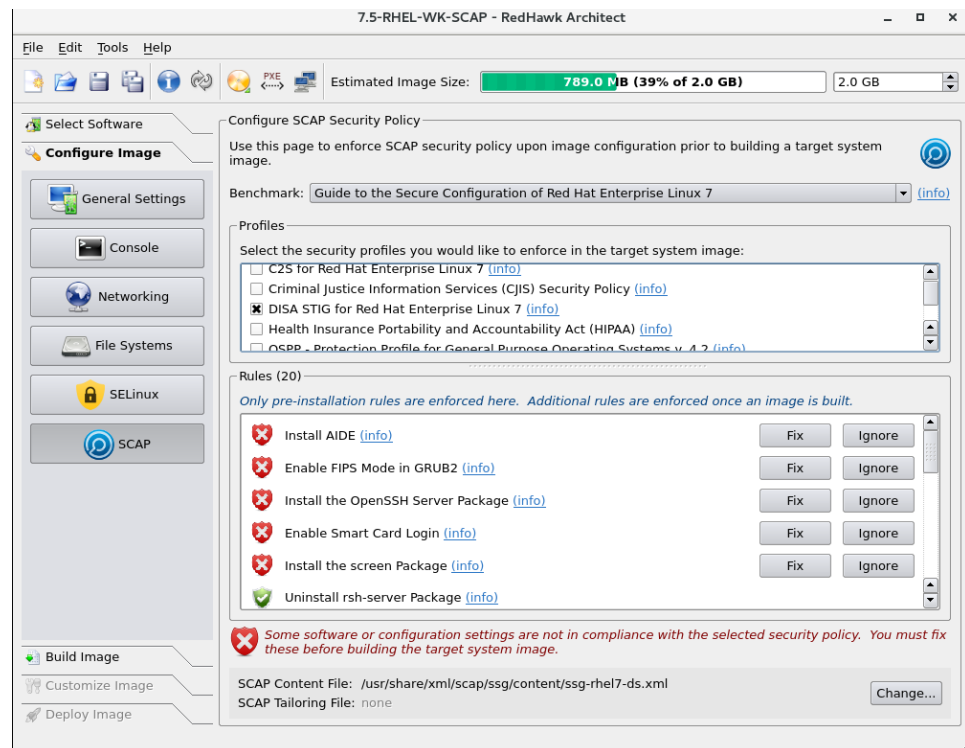


Figure 2-4 SCAP's DISA STIG profile selected

Architect automates most of the pre-installation fixes. In some cases, however, the user is asked to confirm a step to be taken and in other cases more manual intervention is required. An example of the latter is a rule requiring an additional file system be created. In that case, Architect will redirect the user to the Configure File System page so that the user may manually add the requested partition. Note that an (info) link is also available at the bottom of that page.

After SCAP has been configured and the rules applied or ignored, the target system image can be created.

Building the SCAP-configured Image

Refer to the general instructions in "Building an Image" on page 1-23 for building a target system image. The instructions that follow are specific to SCAP-configured target system images.

Note that before building the target system image, the target must be configured with SCAP and all the pre-installation rules applied or ignored. See the previous section "Configuring SCAP" on page 2-6.

To start the build, click on the **Build Image** button of the **Build** toolbox. Once the target system image is built, a remediation scan will automatically start in the chroot of the target system image on the host.

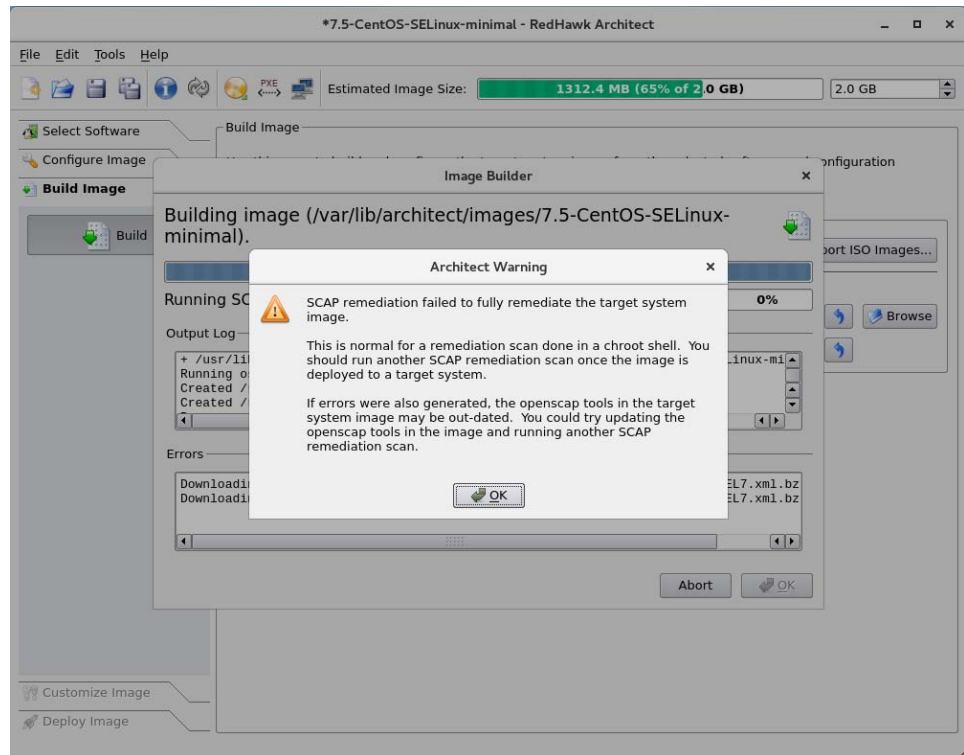


Figure 2-5 Architect's post-build remediation scan fails

The post-build auto-remediation scan in the chroot of the target system image on the host will most likely fail with a warning as in the figure above. This is expected as usually not all SCAP rules are fixed by auto-remediation in a chroot shell. The top window in the **Build Image** page can be scrolled to see the results of the scan or alternatively the output log can be saved to a file. Nonetheless, the target system image is now ready to be deployed and further scans can be run on the target system. More information on scans is

provided in the section "Running SCAP Scans on Deployed Target Systems" on page 2-10.

Customizing the SCAP-configured Image

After building an image, some customization may be necessary on the target system image before the target is deployed. The **Customize Image** toolbox provides several ways to customize the image and in general these are discussed in the section "Customizing an Image" on page 1-27.

Some customization can be done in the chroot of the target system image. To access the Chroot Shell, select the **Chroot Shell** button of the **Customize Image** toolbox and then click on **Run Chroot Shell**.

Below are some customizations, specific to SCAP-configured systems, that can be executed in the Chroot Shell:

- **Perform Manual Remediation of Failed Rules**

Most manual remediation involves editing system files. Information on steps to remediate a failed rule can be obtained from the reports generated by scans, mentioned in the "Run Additional Scans" item below.

- **Run Additional Scans**

Evaluation and remediation scans can be performed in the Chroot Shell on the host system. Scans can be performed in the Chroot Shell on the host system as follows.

To run an evaluation scan in the Chroot Shell, type these commands at the system prompt (specified here as '#'):

```
# cd /root/scap
# /run-eval-scan
```

To run a remediation scan in the Chroot Shell, type these commands:

```
# cd /root/scap
# /run-remediate-scan
```

The scan commands use the SCAP content file also placed in the **/root/scap** directory. Reports from system scans are generated as .html and .txt files and placed in that directory. Besides reporting on the pass/fail status, the reports provide information about the steps necessary to comply with each rule; information useful when manual steps are required. Initially, the report files from Architect's post-build remediation scan is found in the **/root/scap** directory. However, note that reports are overwritten each time a scan is initiated. If reports are to be saved, make sure to move them to another file before initiating a new scan.

- **Add Non-root Users**

Additional users can be added in the Chroot Shell using the **useradd (8)** command. This might be necessary since security-enhanced systems usually put restrictions on root logins.

- Re-enable root SSH login

Some SCAP remediation rules remove root user's privilege to remotely login to the system with `ssh(1)`. To circumvent this, edit the file `/etc/ssh/sshd_config` in the Chroot Shell, search for `PermitRootLogin` in that file and verify that it is set to a value of 'yes'. If not, change the 'no' to a 'yes'. Make sure that the entry does not have a '#' sign at the beginning. Those entries are ignored. Also note that this will change back to a 'no' each time a remediation scan is run.

- FIPS Support

SCAP rules for FIPS require special support in the kernel. Stock RedHawk kernels currently do not have FIPS support enabled. A custom kernel or RedHawk update is required to boot with FIPS enabled. Alternatively, the kernel fips boot options can be modified in order to boot. Click on the **Kernel Manager** button of the **Customize Image** toolbox to view the kernel boot options. If fips is set to one (`fips=1`), change its value to zero (`fips=0`). Note that you may not find this kernel option if, during the scan, the rule that enables FIPS fails.

Deploying the SCAP-configured Image

Refer to the general instructions in "Deploying an Image" on page 1-38 for deploying a target system image. The note that follows is specific to SCAP-configured images.

Some SCAP pre-installation rules require that additional file systems be configured on the **Configure File System** page. The Architect diskless deployment methods ignore all file system configuration settings made on the **Configure File System** page and a single mount point is used for the root '/' file system. For this reason diskless deployment methods may not be a good choice for some SCAP configurations.

Running SCAP Scans on Deployed Target Systems

It is possible and sometimes necessary to run SCAP evaluation and remediation scans on a target system once deployed.

SCAP Workbench is a graphical tool that can be used to perform SCAP scans on a target system. Scans can be done remotely over an SSH connection. SCAP Workbench can be invoked from the **Configure SCAP Security Policy** page of Architect once a target system image has been built.

To execute a remote scan of a target system:

1. Launch the SCAP Workbench GUI from Architect by clicking on the **SCAP Workbench** button. The SCAP content file (and optionally tailoring file) used by the session will be loaded into SCAP Workbench.
2. Select the **Checklist** and **Profile** of interest at the top of the page.
3. **IMPORTANT:** Click on **Remote Machine** (over SSH) and enter the SSH User and host on the target.

4. Check **Fetch remote resources** at the bottom of the page.
5. Check **Remediate** to perform a remediation scan; otherwise an evaluation scan will be done.
6. Click the **SCAN** button. Note that if you are using multiple profiles you will have to repeat steps 2 through 6 for each profile.

Warning

It is very important that the **Remote Machine (over SSH)** button be set on the **SCAP Workbench** page. If not, system changes will be applied to the host system.

Note

SCAP Workbench logs into the remote system via **SSH**. The remote system must be able to **ssh(1)** into the target system and the root user must have privileges to login. This privilege is sometimes removed by rules applied during the post-build remediation scan. If so, edit the file `/etc/ssh/sshd_config` in the chroot of the target system, search for `PermitRootLogin` in that file and verify that it is set to a value of 'yes'.

The following figure shows SCAP Workbench configured to run a remediation scan of the DISA STIG profile rules on a deployed target system specified by its IP address.

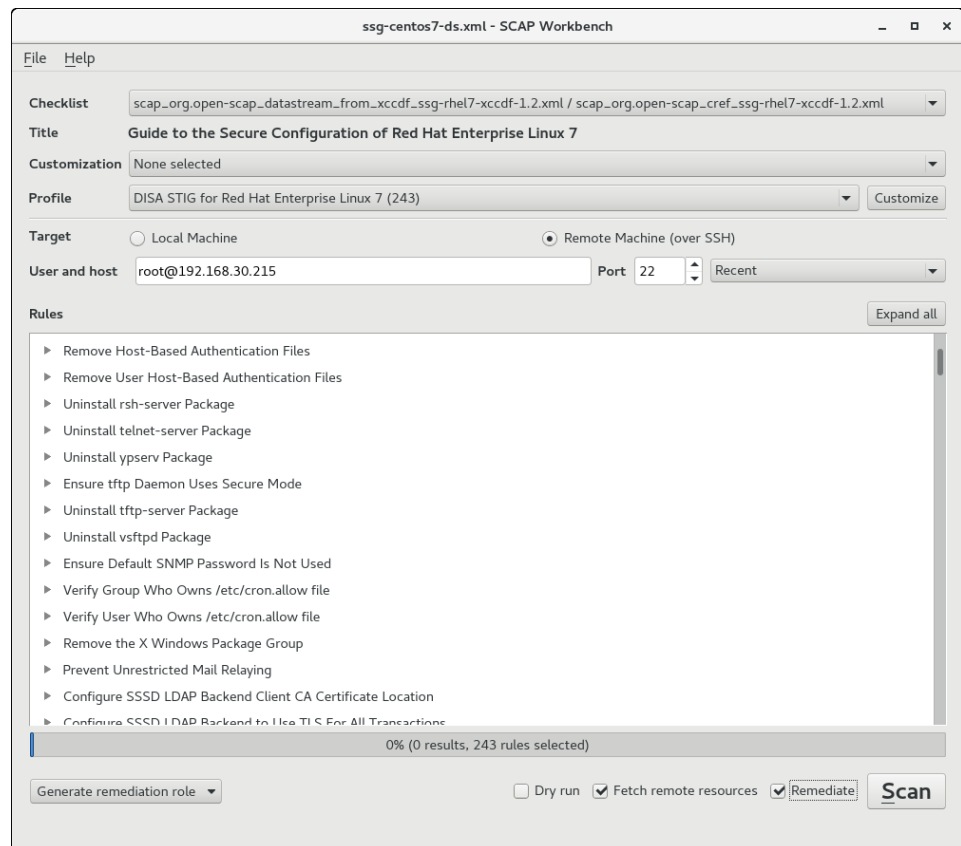


Figure 2-6 SCAP Workbench settings example

It is also possible to run SCAP scans from the command-line on a deployed target system.

To run an evaluation scan on a target system, type these commands at a shell prompt:

```
# cd /root/scap
# ./run-eval-scan
```

To run a remediation scan on a target system, type these commands at a shell prompt:

```
# cd /root/scap
# ./run-remediate-scan
```

More information about running these commands can be found in bullet item "Run Additional Scans" on page 2-9.

Customizing SCAP Content Using SCAP Workbench

The SCAP Workbench tool can be used to create SCAP tailoring files. To create a SCAP tailoring file, perform the following steps:

1. Run **scap-workbench**.

2. Load the content file of interest.
3. Select the Checklist and Profile of interest.
4. Click the **Customize** button (a new GUI window will be displayed where you can customize).
5. Save the customized profile to a SCAP tailoring file. Click on **File**, then **Save Customization Only**.

The following figure shows a customized version of the DISA STIG profile being created with the ntp rule “Configure Time Service Maxpoll Interval” deselected.

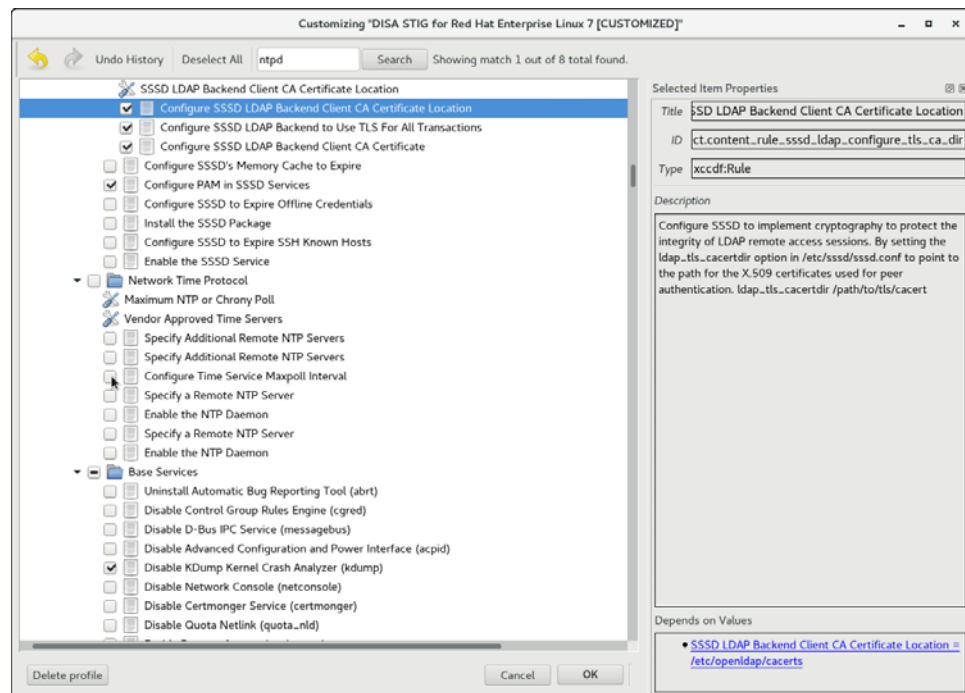


Figure 2-7 Using SCAP Workbench to create a customized profile

Once saved, the customized profile can be loaded via the **Customization** menu. The following figure shows the saved customized profile `/home/ssg-centos7-ds-tailoring.xml` loaded.

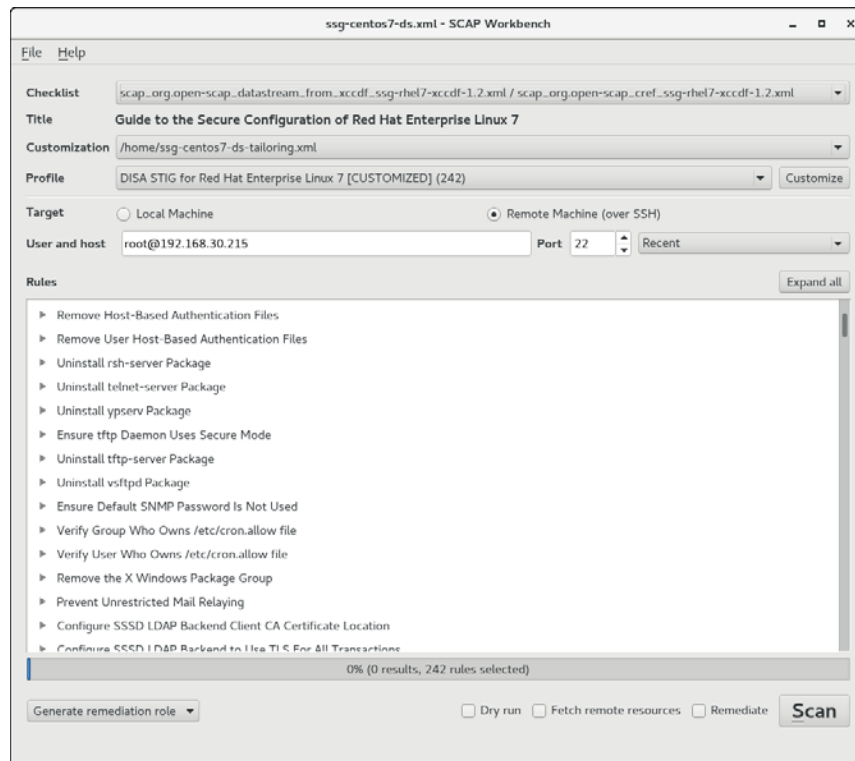


Figure 2-8 SCAP Workbench, customized profile loaded

More information about SCAP Workbench and customizing profiles can be found at <https://www.open-scap.org/resources/documentation/customizing-scap-security-guide-for-your-use-case> and other sites.

Known SCAP Issues

RHEL-based SCAP content files are required for pre-installation rules in Architect. If more recent SCAP content is required, RHEL-based SCAP files may be used in Architect for pre-installation, and the more recent SCAP files may be used to scan on the deployed targets.

If SCAP scanning produces errors, some SCAP rules may be broken. Verify that the latest SCAP software is installed on the system. Update SCAP content files by updating the **scap-security-guide** on the host: **"yum update scap-security-guide"** and on the target system update the **openscap-scanner** package **"yum update openscap-scanner"**.

If the root user is unable to login to the target system, verify that the root user has SSH permission to login remotely. Edit **/etc/ssh/sshd_config** in the chroot of the target system image on the host and make sure **PermitRootLogin** is set to 'yes' in that file. Non-root users may also be added in the Chroot Shell using the **useradd (8)** command.

If the target system gets FIPS errors and fails to boot, verify that the **fips** kernel option is set to zero. SCAP rules for FIPS require special kernel support not currently enabled in the

stock RedHawk kernel. A custom kernel or RedHawk update is required to boot with "*fips=1*" set.

The diskless deployment method should not be used when SCAP rules require multiple file systems on multiple partitions. Diskless images use a single root file system.

Red Hat also has problems posted to their issue-tracking system, Bugzilla. Below are a few.

SCAP Workbench problems:

https://bugzilla.redhat.com/show_bug.cgi?id=1464615

https://bugzilla.redhat.com/show_bug.cgi?id=1456429

openscap problems:

https://bugzilla.redhat.com/show_bug.cgi?id=1431186

Importing ISO Images

This chapter describes how to create or import on-disk ISO images to dramatically speed up and virtually automate target file-system image creation.

Importing ISO Images

Normally when building a target file-system image the user is prompted to insert various optical media discs containing the software that is required in order to create the initial target file-system image. If only one or two images are being produced, manually inserting optical media is generally acceptable.

However, if the user is generating and maintaining several different target file-system image configurations, it is often preferable to create on-disk ISO images of the various optical media discs. To accomplish this, select the **Media ISO Manager** item in the **Tools** menu, or click the **Import ISO Images** button on the **Build Image** page, and the following dialog will appear.

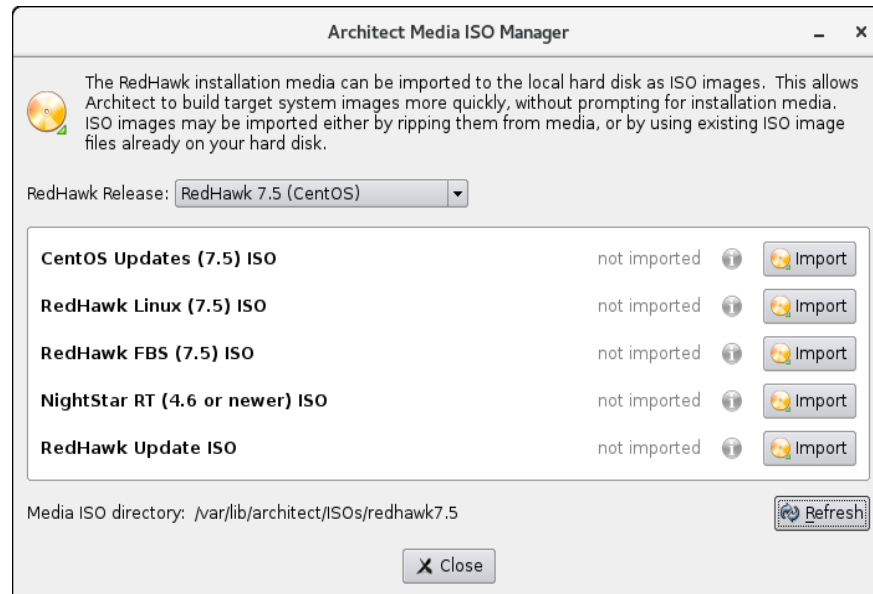


Figure 3-1 Import ISO Images Dialog

Pressing the **Import** button, will display a menu with three different options to import ISO images:

- Import ISO images directly from manually inserted optical media

- Copy ISO images from already existing ISO image files
- Link ISO images to already existing ISO images files

These various methods will be described in the following sections.

The user can import different sets of ISO images for different RedHawk release versions; use the **Select a RedHawk release** pull-down menu to select which version of RedHawk to import ISO images for.

In addition, different import methods can be used *within* a specific RedHawk release. For example, it is possible to use one import method to import the CentOS ISO image and a different import method to import the RedHawk and NightStar ISO images. All combinations are valid.

Importing ISO Images From Optical Media

To use this method select the **Rip ISO** from **media** import method and then press the **OK** button to begin the import process. A dialog similar to the following dialog will be displayed.

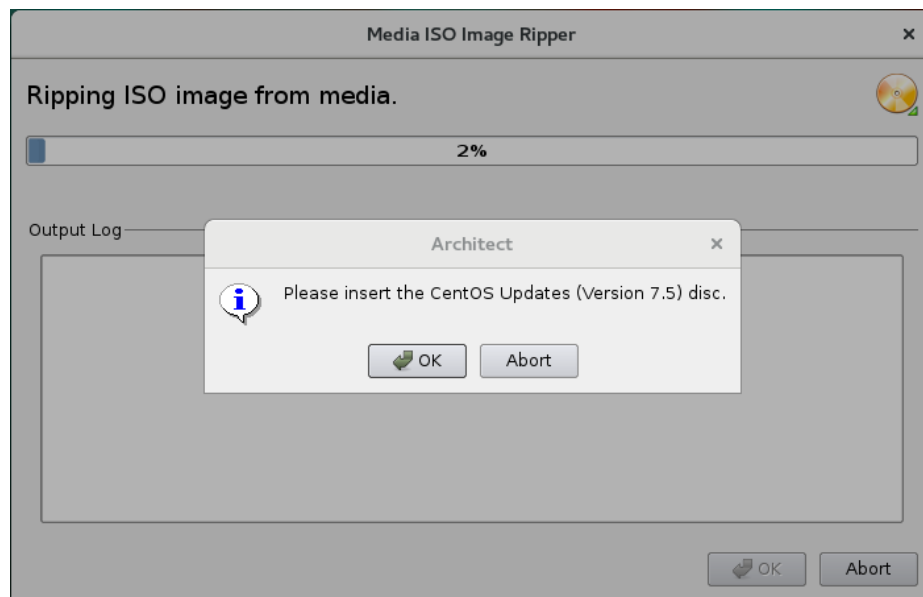


Figure 3-2 Rip ISO Images From Media

At this point the correct optical media disc for the requested item should be manually inserted into the host system's optical media tray. Once the optical media has been inserted, press **OK** to begin copying the ISO image from the optical media onto the host system's hard-drive. Various status messages will be displayed as the copy progresses.

Copying ISO Images From Existing ISO Images

If you already have the required media in ISO format on disk, Architect can import the ISO by creating Architect-specific copies of the ISO images; copying is useful when the original ISO images may be removed or unavailable at some point in the future.

To make ISO copies, select the **Copy existing ISO file on disk** import method and then press the **OK** button to begin the import process. A file selection dialog will be displayed. Navigate the file selection dialog to the appropriate directory to select the ISO image. An example ISO file selection is presented below. In this example, the ISO images are stored in the **/root/Downloads** directory and the CentOS Updates ISO image has been selected.




Figure 3-3 Copy ISO Image File Selector

Press the **Open** button to begin the process of copying the ISO image file into Architect's **/var/lib/architect/ISOs** directory. Once the copy is completed, the ISO image file that was copied is no longer needed and can be removed if necessary.

Linking To Existing ISO Images

If you already have the required media in ISO format on disk, Architect can import the ISO by creating symbolic links to the ISO images; linking is useful when you can be sure that the original ISO images will persist indefinitely.

To create ISO symbolic links, select the **Symbolically link to existing ISO file on disk** import method and then press the **OK** button. Navigate the file selection dialog presented to the appropriate directory and select the ISO image. An example ISO file selection is presented below. In the example, the ISO images are stored in the **/root/Downloads** directory and the RedHawk Linux ISO image has been selected.

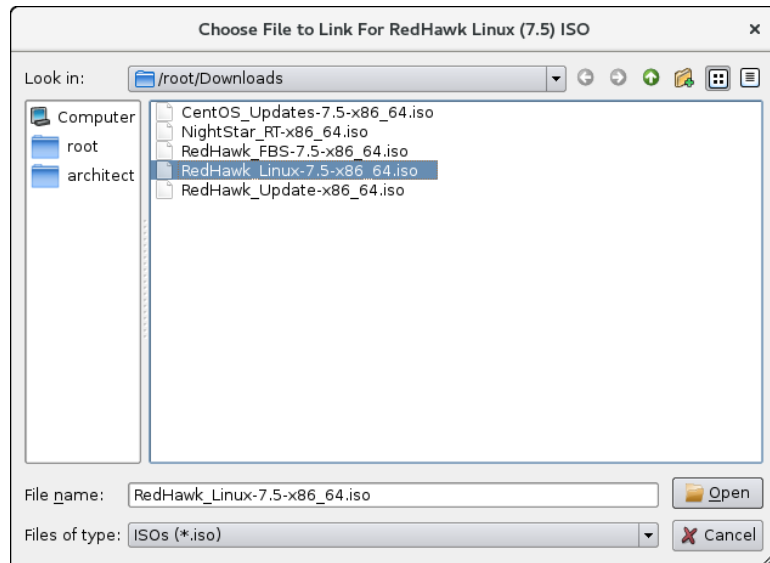


Figure 3-4 Symlink To ISO Image File Selector

Press the **Open** button to immediately create a symbolic link to the selected ISO image file. The symbolic link will be created and placed inside Architect's `/var/lib/architect/ISOs` directory. Once the copy is completed, the ISO image file that was linked to must be preserved and kept at the exact same file-system location in order for Architect's symbolic link to be valid.

NOTE

Architect will detect if it has symbolic links to ISO image files that have been erroneously removed and ISO image will no longer be shown as a valid ISO image in the list of imported ISO images.

If this happens, the ISO image must be imported again to be valid, otherwise Architect will prompt for the corresponding optical media disc during any subsequent target file-system image builds.

Deleting Imported ISO Images

Previously imported ISO images can be deleted at any time by pressing the **Delete** button of the corresponding ISO image. This is not generally necessary, but can be done in order to save disk space or to recover from the rare case of a file corruption.

This chapter describes how to manage PXE resources on the host and how targets in your network environment will use these PXE resources.

Enabling PXE on Targets

The Preboot eXecution Environment (PXE) provides a method for booting target systems using a network interface, without the requirement of having access to any local storage on the target system.

To use PXE, targets must first be configured to perform a PXE broadcast during boot. To enable the PXE broadcast perform the following steps:

1. Reboot the target and stop the system immediately after POST (Power-On Self-Test), normally by pressing Delete or F2, to get into the BIOS settings menu.
2. Each kind of computer has a slightly different BIOS settings menu, however the general rule is to navigate to the 'PCI Device' or the 'Integrated Devices' section of the BIOS menu and enable PXE boot on the first Ethernet interface that is present. Ensure that the chosen interface is connected to a switch that is present on the same network as the host system.
3. Record the MAC address of the target's Ethernet interface for later use with Architect's PXE Target Manager dialog. See "Managing PXE Targets" on page 4-7 for more information.

NOTE

Some older BIOSes do not provide an option to boot with PXE. The *Etherboot* utility can be used instead, however Concurrent Real-Time does not support this configuration. See <http://etherboot.org> for more information.

Initializing PXE Services

Various PXE-related services on the host system need to be properly initialized before any of the PXE-based image deployment methods can be used. To initialize these services, click on PXE Target Manager in the Tools menu and you will be presented with the following dialog.

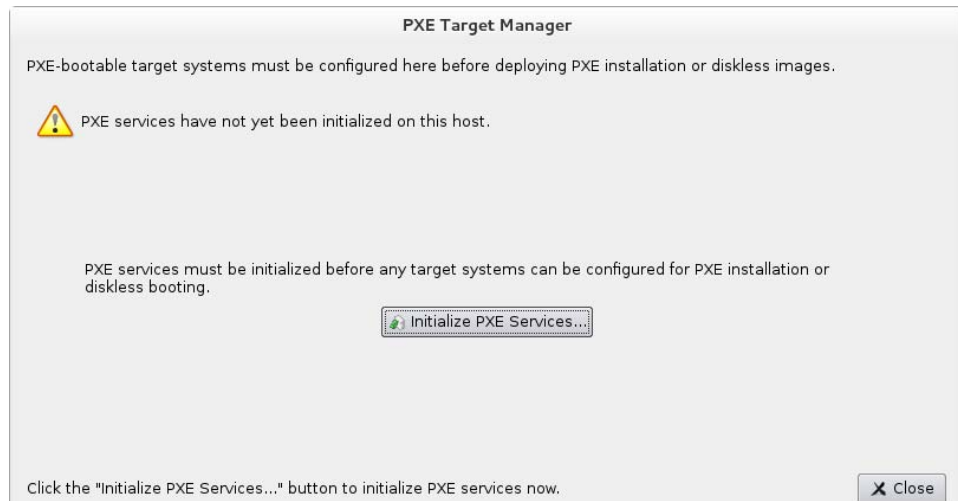


Figure 4-1 PXE Target Manager Uninitialized

Press **Initialize PXE Services...** to begin initializing the PXE services and you will then be presented with the following dialog.

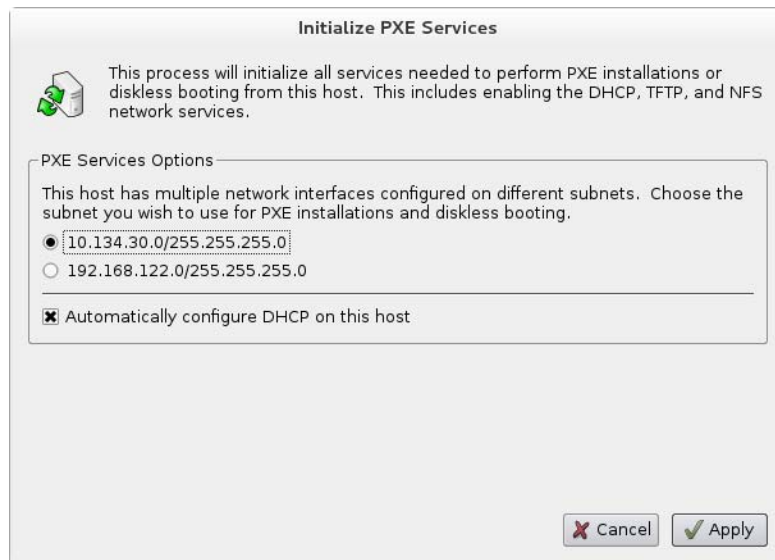


Figure 4-2 Initialize PXE Services Dialog

First, choose the network subnet that you wish to use for all PXE communications between the host and its targets. If only one subnet is available a choice cannot be made, however the information is still presented so that the user can verify that the subnet is the desired subnet.

By default DHCP services will be automatically configured and enabled on the host, and this is the recommended approach. However if another DHCP server already exists on the chosen subnet you will need to uncheck **Automatically configure DHCP on this host** or the two DHCP servers will conflict with each other. In this case, you will

need to manually merge the DHCP configuration files generated by Architect on the host with those of the actual DHCP server. See “Manual DHCP Configuration” on page A-1 for more information.

Once these settings are correct for your environment click **Apply** and the initialization will begin. Once the initialization has completed successfully you will see the following displayed in the dialog.



Figure 4-3 PXE Services Initializer Done

Press **OK** to return to the main PXE Target Manager window. At this point the host is now configured with the required networking services to enable PXE image deployments.

Managing PXE Images

PXE images that have been created with the PXE Installer and PXE Diskless tools in RedHawk Architect's Deploy Image toolbox are resources that can be inspected and managed with the PXE Image Manager.

Select **PXE Image Manager** from the **Tools** menu to access the PXE Image Manager. If no PXE images have yet been deployed you will be presented with the following empty dialog.



Figure 4-4 PXE Image Manager

The PXE Image Manager will remain empty until PXE images are created using the PXE deployment tools in the Deploy Image toolbox.

PXE Installer Images

The PXE Image Manager lists all of the installation images that have been deployed by the PXE Installer tool. The following dialog shows an example of the PXE Image Manager with installation images.

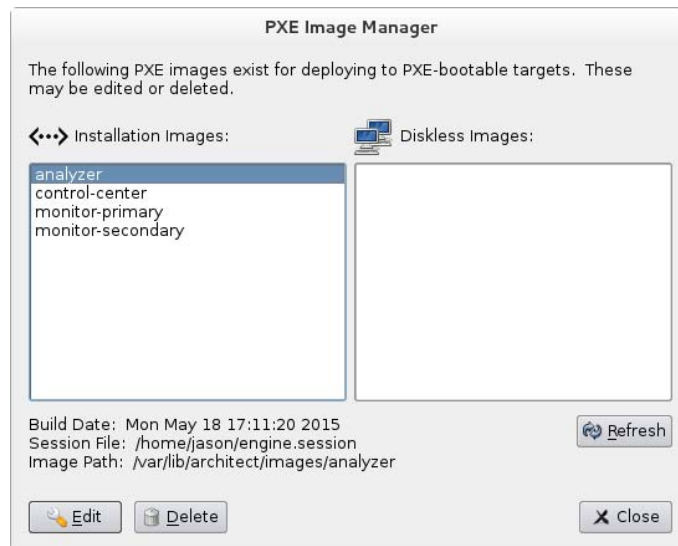


Figure 4-5 PXE Image Manager with Installation Images

Each PXE installation image that is created using the PXE Installer deployment method is effectively a snapshot in time of the root file system image being managed during an Architect session. These images can be inspected and/or individually removed.

Selecting an installation image from the list will display the following details about the image:

- The date the installation image was deployed
- The session file that was being used at the time of creation; if the session has not yet been saved the string `None` will be displayed instead
- The path of the root file system image that the installation image was created for

To delete an installation image, first select it in the list and then press the `Delete` button. You will be presented with a dialog asking for confirmation and simply press `Yes` to delete it.

To edit the attributes of an installation image, first select it in the list and then press the `Edit` button. You will be presented with a dialog allowing you to modify several attributes of the installation image including:

- The `Serial Console` for the install image.
- The `Automatically install image to disk when target is PXE-booted` checkbox.

Press `OK` to apply any changes made to the attributes of the installation image.

The `Refresh` button will refresh the list to match the resources currently on disk, however refresh is only useful if multiple copies of Architect are being used simultaneously to create and manage PXE installation images.

Press `Close` at any time to dismiss the dialog and return to the Architect main window.

PXE Diskless Images

The PXE Image Manager lists all of the diskless images that have been created by the PXE Diskless tool. The following dialog shows an example of the PXE Image Manger with diskless images.



Figure 4-6 PXE Image Manager with Diskless Images

Similar to PXE installation images, PXE diskless images created using the PXE Diskless deployment method create a snapshot in time of the root file system image being managed. These images can be inspected and/or individually removed.

Selecting a diskless image from the list will display the following details about the image:

- The date the diskless image was created
- The session file that was being used at the time of creation; if the session has not yet been saved the string `None` will be displayed instead
- The path of the root file system image that the diskless image was created for

To delete a diskless image, first select it in the list and then press the `Delete` button. You will be presented with a dialog asking for confirmation and simply press `Yes` to delete it.

To edit the attributes of a diskless image, first select it in the list and then press the `Edit` button. You will be presented with a dialog allowing you to modify several attributes of the diskless image including:

- The `PXE/DHCP Device` that the diskless image should use for all PXE and DHCP network traffic.
- The `Serial Console` for the diskless image.
- The `Kernel to Boot` for the diskless image.
- Any `Extra Kernel Options` for the diskless image's kernel to use.
- The `Boot Timeout` for the diskless image's boot menu to use.

Press `OK` to apply any changes made to the attributes of the diskless image.

The `Refresh` button will refresh the list to match the resources currently on disk, however refresh is only useful if multiple copies of Architect are being used simultaneously to create and manage PXE diskless images.

Press **Close** at any time to dismiss the dialog and return to the Architect main window.

Managing PXE Targets

PXE images that have been created with the PXE Installer and PXE Diskless tools in RedHawk Architect's **Deploy Image** toolbox are resources that can be assigned to targets using the PXE Target Manager.

Select **PXE Target Manager** from the **Tools** menu to access the PXE Target Manager. If no targets have yet been added you will be presented with the following dialog showing an empty list of targets.

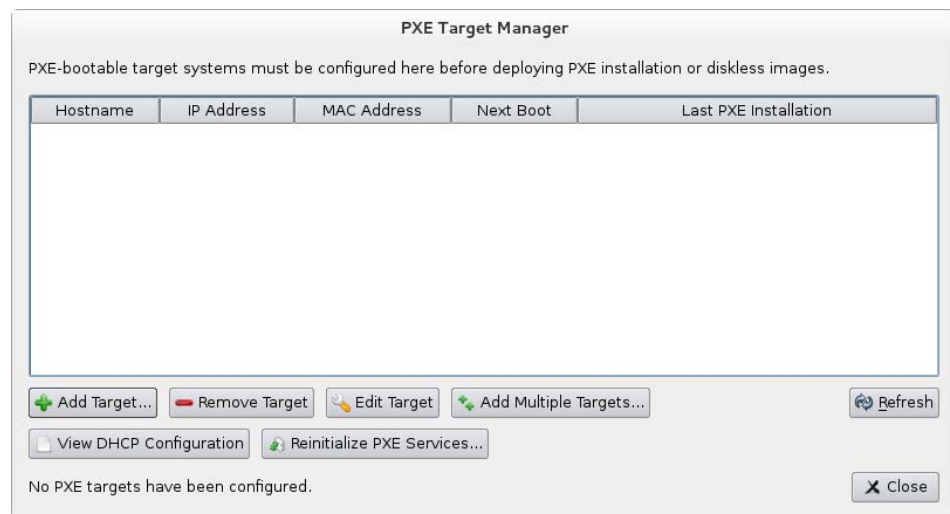


Figure 4-7 PXE Target Manager

The PXE Target Manager's target list will remain empty until targets are added using one of the **Add** buttons below the list.

Adding Targets

All targets that will be using PXE installation images and/or PXE diskless images must first be added to the PXE Target Manager. Targets can be added either individually or in groups, and these two methods are described in the following sections.

Adding Single Targets

A single target can be added to the PXE Target Manager by pressing the **Add Target...** button on the PXE Target Manager dialog. The following dialog will be shown.

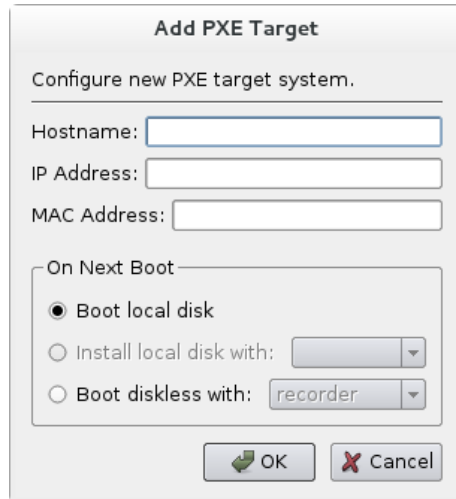


Figure 4-8 Add PXE Target Dialog

Enter the hostname, IP address and MAC address of the target in the corresponding fields.

In the **On Next Boot** area of the dialog, choose the desired target behavior that it will perform after its next reboot and subsequent PXE broadcast. The following behaviors are supported:

- Choose **Boot local disk** to have the target simply boot from its local disk upon next reboot.
- Choose **Install local disk with** and select a PXE installation image from the pulldown to have the target install the local disk with the selected PXE installation image upon next reboot. This option is only available if PXE installation images have previously been created; see “Installing via PXE over a Network” on page 1-45 for more information.
- Choose **Boot diskless with** and select a PXE diskless image from the pulldown to have the target boot disklessly with the selected PXE diskless image upon next reboot. This option is only available if PXE diskless images have previously been created; see “Booting Diskless via PXE over a Network” on page 1-47 for more information.

Press **OK** to add this target to the PXE Target Manager and dismiss the dialog.

After targets have been entered the PXE Target Manager will look similar to the following example:




Figure 4-9 PXE Target Manager with Targets

When you are finished adding targets press the **Close** button to return to the Architect main page.

Adding Multiple Targets

Multiple target can be added to the PXE Target Manager by pressing the **Add Multiple Targets...** button on the PXE Target Manager dialog. The following dialog will be shown.




Figure 4-10 Add Multiple PXE Targets Dialog

Enter the starting IP address in the corresponding field. This will be the address of the *first* target in the target group, and each additional target will simply increment this setting by one IP address.

Choose the number of targets to configure in the corresponding field. You can configure up to 256 targets simultaneously using this interface.

Enter the hostname prefix to be used for all targets in the Base Hostname field. This prefix will be used for the start of each hostname generated and a unique integer suffix will be appended for each target.

In the On Next Boot area of the dialog, choose the desired target behavior that it will perform after its next reboot and subsequent PXE broadcast. See the discussion of On Next Boot above on page 4-7 for more information.

Once the desired settings have been entered, click the Apply button. A dialog similar to the following will be shown:

Add Multiple PXE Targets

You may configure multiple new PXE target systems at one time.

These targets will be configured with IP addresses from a range of contiguous addresses on the PXE subnet (10.134.30.0/255.255.255.0). Hostnames for the targets will be generated automatically.

Starting IP Address:

How many targets would you like to configure?

Base Hostname:

On Next Boot

Boot local disk

Install local disk with:

Boot diskless with:

Hostname	IP Address	MAC Address
target-001	10.134.30.1	<input type="text"/>
target-002	10.134.30.2	<input type="text"/>
target-003	10.134.30.3	<input type="text"/>
target-004	10.134.30.4	<input type="text"/>
target-005	10.134.30.5	<input type="text"/>
target-006	10.134.30.6	<input type="text"/>
target-007	10.134.30.7	<input type="text"/>
target-008	10.134.30.8	<input type="text"/>
target-009	10.134.30.9	<input type="text"/>
target-010	10.134.30.10	<input type="text"/>

Figure 4-11 Add Multiple PXE Targets after Apply

Pressing Apply caused the dialog to generate hostname entries for all of the requested targets. Enter the MAC address of each target into its corresponding MAC Address field.

A MAC addresses is required for each target if Architect is directly managing DHCP services. However, MAC addresses are not necessary when you are *not* using Architect to directly manage DHCP services; in that case you can leave them blank. See “Manual DHCP Configuration” on page A-1 for more information.

Removing Targets

To remove a target that is currently being managed by the PXE Target Manager first select the target hostname in the list and then press the **Remove Target** button. You will be presented with a confirmation dialog. Press **Yes** to remove the target. Note that the target can be again added at any time if desired.

Editing Targets

To change the settings for a target currently being managed by the PXE Target Manager first select the target hostname in the list and then press the **Edit Target** button. You will be presented with a dialog similar to the following:

The screenshot shows a dialog box titled "Edit PXE Target". Inside, it says "Editing target 'analyzer'". There are three text input fields: "Hostname:" with "analyzer", "IP Address:" with "10.134.30.1", and "MAC Address:" with "00:1B:21:D8:51:0C". Below these is a section titled "On Next Boot" containing three radio button options: "Boot local disk", "Install local disk with:" (selected) with a dropdown menu showing "analyzer", and "Boot diskless with:" with a dropdown menu showing "recorder". At the bottom right are "OK" and "Cancel" buttons.

Figure 4-12 Edit PXE Target Dialog

With this dialog you can change the hostname, IP address and MAC address of the host. You can also change the **On Next Boot** setting to change the behavior of the target upon next reboot. Refer to the discussion of **On Next Boot** above on page 4-7 for more information.

Press **OK** to apply these settings and return to the PXE Target Manager.

Manual DHCP Configuration

This appendix describes how to add the required DHCP configuration for Architect PXE targets to an active existing DHCP server configuration. It is preferable to allow the Architect tool to administer DHCP by checking the box labeled **Automatically configure DHCP on this host**, however if another DHCP server already exists on the desired subnet you must follow the steps described in this section. Refer to “Initializing PXE Services” on page 4-1 and also the **dhcpcd.conf(5)** man page for more information.

Overview

The **View DHCP Configuration** button on the PXE Target Manager may be used to view the required DHCP configuration for Architect PXE targets. The information displayed may be cut and pasted into a text editor when editing the existing DHCP server configuration.

Alternatively, the DHCP configuration files maintained by Architect may be viewed or copied directly from the `/etc/dhcp/architect` directory on the host system where Architect is installed. This directory contains two files: `dhcpcd.conf` and `dhcpcd-targets.conf`. The `dhcpcd.conf` file contains a subnet stanza with all required DHCP parameters set for PXE targets and it will look similar to the following example:

```
option pxe-client-arch-type code 93 = unsigned integer
16;

subnet 10.134.30.0 netmask 255.255.255.0 {
    option subnet-mask 255.255.255.0;
    option broadcast-address 10.134.30.255;

    server-name cholula;
    next-server 10.134.30.166;
    if option pxe-client-arch-type = 00:09 {
        filename "architect/efi64/syslinux.efi";
    } elsif option pxe-client-arch-type = 00:07 {
        filename "architect/efi64/syslinux.efi";
    } else {
        filename "architect/bios/pxelinux.0";
    }

    use-host-decl-names on;
    include "/etc/dhcp/architect/dhcpcd-targets.conf";
}
```

In this example, the PXE subnet is the 10.134.30.0/24 subnet. The last line includes all PXE target host declarations from the `dhcpd-targets.conf` file, which will look similar to the following example:

```
host monitor2 {
    hardware ethernet 00:02:AC:55:88:A9;
    fixed-address 10.134.30.65;
}
host analyzer {
    hardware ethernet 00:1B:21:D8:51:0C;
    fixed-address 10.134.30.11;
}
host ccenter {
    hardware ethernet 84:2B:2B:9E:6E:1B;
    fixed-address 10.134.30.17;
}
host monitor1 {
    hardware ethernet 00:23:AE:D9:1C:AF;
    fixed-address 10.134.30.64;
}
host recorder {
    hardware ethernet 00:80:8E:02:9A:92;
    fixed-address 10.134.30.72;
}
```

This configuration data must be added to the active DHCP server configuration file(s). On most systems, the main DHCP configuration file is `/etc/dhcp/dhcpd.conf`.

Installing DHCP Configuration

The simplest way to add the Architect DHCP configuration to the DHCP server is to copy the files from `/etc/dhcp/architect` on the Architect host to the same location on the DHCP server host, and then add a single `include` line to the existing `/etc/dhcp/dhcpd.conf` file to include the Architect configuration. If creating the `/etc/dhcp/architect` directory on the DHCP server host is not possible, you may use any valid location on the file system; simply adjust the include lines accordingly.

To accomplish this perform the following steps:

1. Copy files from the Architect host to the DHCP server host. For example, run this command on the Architect host:

```
scp -r /etc/dhcp/architect dhcp_server:/etc/dhcp
```

where `dhcp_server` is the name or IP address of the DHCP server host.

2. Include this configuration in the main DHCP server configuration file. Edit `/etc/dhcp/dhcpd.conf` on the DHCP server host and add this line near the bottom of the file:

```
include "/etc/dhcp/architect/dhcpd.conf";
```


Note that most DHCP servers allow multiple subnet stanzas to be defined for the same subnet, each with different parameters defined within the scope of the stanza. Because of this, you are allowed to have the PXE target systems declared within one subnet stanza, and other DHCP clients or a dynamic IP address pool declared in another subnet stanza for the same subnet.

NOTE

You cannot have duplicate host declarations or reuse an IP address or MAC address in different host declarations anywhere in the entire DHCP configuration.

See the `dhcpcd.conf(5)` man page for more information.

Paths

/etc/dhcp/architect A-1, A-2
/etc/dhcp/dhcpd.conf A-2
/var/lib/architect 3-3, 3-4

A

Architect

Advanced Security Edition 2-1
introduction 1-1
main window 1-3
opening dialog 1-3
running 1-2

B

Base Distribution packages 1-5
baud rate 1-14
BIOS settings 4-1
Board Support Packages 1-35
building an image 1-23

C

ccur-config 1-32, 1-33
chroot shell 1-36
configuring an image 1-11
console 1-13
copy files into image 1-36
create new directories in image 1-36
creating a new session 1-3, 1-4
Custom Kernel
 build 1-33
 configuration 1-31
 remove 1-34
customizing an image 1-27

D

default gateway 1-16

delete files in image 1-36
DHCP
 automatic configuration 4-2
 manual configuration A-1
 view configuration A-1
dhcpd.conf A-1
dhcpd-targets.conf A-1
domains 1-16

E

editing an existing session 1-3, 1-4, 1-55
enabling PXE 4-1

F

File Manager 1-36
flash error 1-41
flashing an image 1-38

H

hostname 1-16

I

Image Cleanup 1-38
Initialize PXE Services 4-2
install additional RPMs 1-35
installing software 1-4, 1-23

K

Kernel
 configuring 1-31
 export configuration 1-33
 import configuration 1-32
Kernel Manager 1-30

M

main window 1-3

N

networking 1-15
new session 1-3, 1-4
New Session dialog 1-4
NightStar RT installation options 1-11
noatime file system option 1-18

O

On Next Boot behavior 4-8, 4-10, 4-11
Out-of-Sync Notice 1-13, 1-14, 1-16, 1-18, 1-21, 1-22

P

POST 4-1
Power On Self Test 4-1
primary DNS server 1-16
PXE
 broadcast 4-1, 4-8, 4-10
 diskless images 4-5
 enabling 4-1
 installer images 4-4
 manage images 4-3
 manage targets 4-1, 4-7

R

read-only root file system 1-18
RedHawk installation options 1-9
Remove RPM database 1-38
Remove selected from image 1-38
root password 1-12
running Architect 1-2

S

saving a session 1-3, 1-4, 1-55
SCAP (Security Content Automation Protocol)
 building target image 2-8
 configuring 2-6
 content file 2-4, 2-5, 2-6, 2-7, 2-9, 2-12, 2-14
 customizing target image 2-9
 deploying target image 2-10
 FIPS 2-10, 2-14

introduction 2-4
known issues 2-14
running scans 2-5, 2-9, 2-10
SCAP Workbench 2-5, 2-10, 2-12, 2-15
tailoring content 2-12
workflow overview 2-4
secondary DNS server 1-16
Security Extensions 2-1
SELinux 2-1
 configuring 2-1
serial port 1-14
software to install 1-4
system run level 1-12

T

time zone 1-12

V

Virtual Machine
 booting with QEMU 1-54
 deployment 1-54
 disk image 1-54
 image synchronization 1-54