

 REDHAWK *Architect™ User's Guide*

Copyright 2026 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, 800 NW 33rd Street, Pompano Beach, Florida 33064. 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.

<u>Revision History:</u>	<u>Level:</u>	<u>Effective With:</u>
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	92	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
June 2020	8.0	RedHawk Linux 8.0
June 2021	8.2	RedHawk Linux 8.2
March 2022	8.4	RedHawk Linux 8.4
June 2022	8.4-1	RedHawk Linux 8.4
June 2022	8.4-1a	RedHawk Linux 8.4
September 2022	8.4-2	RedHawk Linux 8.4
June 2023	8.4-3	RedHawk Linux 8.4
September 2023	9.2	RedHawk Linux 9.2
March 2024	9.2-1	RedHawk Linux 9.2
July 2024	9.2-2	RedHawk Linux 9.2
March 2025	9.2-3	RedHawk Linux 9.2
June 2026	9.6	RedHawk Linux 9.6
September 2026	9.6-1	RedHawk Linux 9.6

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 introduces you to RedHawk Architect and guides you through its use.
- Chapter 2 explains how to use the security extension of the Advanced Security Edition of Architect.
- Chapter 3 explains Importing ISO Images to avoid repetitive manual optical media insertion.
- Chapter 4 explains PXE Management.
- Appendix A explains Manual DHCP configuration for Architect PXE targets.
- Appendix B lists the python scripts that can be used to perform some of the functions provided by the GUI.
- Appendix C is an example of SCAP using the DISA STIG with GUI profile.

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

corresponding text. Clicking on Internet URLs provided in **blue** type will launch your web browser and display the web site. Clicking on publication names and numbers in **red** type will display the corresponding manual PDF, if 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 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	1-2
Running Architect	1-2
User Preferences	1-5
Creating a New Session	1-6
Editing and Saving an Existing Session	1-6
Selecting Software to Install in the Image	1-7
Selecting Base Distribution Linux Packages	1-7
Using the Base Environments View	1-8
Using the Categories and Groups View	1-9
Using the All Packages View	1-10
Using the Selected Packages View	1-10
Selecting RedHawk OS Options	1-11
Selecting NightStar Options	1-13
Configuring an Image	1-14
Configuring General Settings	1-14
Configuring a Console	1-16
Configuring Networking	1-18
Configuring File Systems	1-20
Simple Disk Partitioning	1-21
Advanced Disk Partitioning	1-22
Default File System Configuration	1-23
Optional File System Configurations	1-24
Example Using Optional Configurations	1-25
Steps in Configuring Additional Logical Volumes (LVM)	1-27
Example Configuring Logical Volumes (LVM)	1-28
Redundant Array of Independent Disks (RAID)	1-30
Steps in configuring a RAID	1-31
Example System Configuration Using RAID Partitions	1-32
Example Default Autogenerated RAID Configuration	1-40
Example Configuration Using multiple RAID levels	1-42
Read-Only Root Configuration	1-44
Configuring Special File Systems (tmpfs, bind)	1-45
Configuring SELinux	1-47
Building an Image	1-47
Customizing an Image	1-51
Software Updates	1-52
Additional RPMs	1-53
Installing Board Support Packages	1-54
System Services	1-54
Kernel	1-55
Configure Custom Kernel	1-57
Import Kernel Configuration	1-58
Export Kernel Configuration	1-59

Compile Custom Kernel	1-59
Remove Custom Kernel	1-61
File Manager	1-61
Chroot Shell	1-62
Image Cleanup	1-63
Deploying an Image	1-64
Deploying to USB Devices	1-66
Installing via USB drive	1-68
Installing via Disc media	1-72
Installing via PXE over a Network	1-74
Bootting Diskless via PXE over a Network	1-76
Remote Sync	1-81
Deploying to Virtual Machine	1-86
Extracting a session from Architect media	1-87

Chapter 2 Security Extensions

UEFI Secure Boot	2-1
Configuring SELinux	2-2
Customizing the kernel with FIPS	2-4
Security Content Automation Protocol (SCAP)	2-5
Introduction to SCAP	2-5
Understanding SCAP Evaluation and Remediation Scans	2-6
SCAP System Requirements	2-7
SCAP Workflow	2-7
Workflow Overview	2-7
1. Configure SCAP security policy	2-8
2. Build the target system image	2-11
3. Customize the target image	2-12
4. Run post-deploy scans	2-13
Directly from console	2-13
Remotely via ssh	2-14
Remotely from Architect	2-14
Customizing SCAP Content Using SCAP Workbench	2-15

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	3-4
Deleting Imported ISO Images	3-5

Chapter 4 PXE Management

Enabling PXE on Targets	4-1
Initializing PXE Services	4-1
Managing PXE Images	4-3
PXE Diskless Images	4-6
Managing PXE Targets	4-7
Adding Targets	4-8
Adding Single Targets	4-8
Adding Multiple Targets	4-9

Removing Targets	4-11
Editing Targets	4-11
Appendix A Manual DHCP Configuration	A-1
Overview	A-1
Installing DHCP Configuration	A-2
Appendix B Command Line Interface	B-1
Appendix C DISA STIG Example	C-1
DISA STIG System Requirements	C-1
DISA STIG Workflow Overview	C-2
Configure DISA STIG Security Policies:	C-2
Build the Image	C-3
Customize the Target Image	C-3
Run post-deployment scans	C-4
Directly from console	C-5
Remotely via ssh	C-5
Remotely from Architect	C-6
Manual Remediation	C-6
Sync from Target	C-7

Using RedHawk Architect

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

Introducing Architect

RedHawk Architect is a powerful tool with an easy-to-use GUI that lets a developer choose the Linux and application modules to be included in RedHawk target images.

Designed to scale from complete workstations to dedicated servers and even down to small embedded applications, users can select as few or as many packages as desired from many different package groups. Architect allows the file system to be customized using as much or as little space as desired.

Architect saves all the configuration choices made by the user in the session file. Architect processes the session file to create and configure the target system image by installing all necessary RPM packages and dependencies and making the desired configuration changes. Packages are installed from system installation media which may be saved on the host system as ISO images.

After building the target system image, Architect allows customization of the target system image and also the RedHawk kernel itself.

Architect provides several deployment methods for installing the target system image onto the target system hard drives and/or flash memories. Architect can boot diskless targets with a target system image. It can also build virtual machine images of target system images for use with QEMU/KVM, allowing complete target system images to be deployed without a physical target.

Architect's PXE target manager makes it simple for users to install and configure systems as highly-integrated, high-performance computing clusters. Architect uses PXE to remotely install any number of targets over a network and also for diskless-booting multiple nodes with the same version of RedHawk.

Architect greatly simplifies the following tasks to create and maintain a target system's runtime and development environment:

- installing custom configurations of the Rocky™, Oracle®, Fedora or Red Hat® Enterprise Linux distribution
- installing and configuring the RedHawk™ operating system
- installing target-specific board support packages (BSPs)
- installing NightStar™ RT application development tools
- installing RedHawk and NightStar software updates
- maintaining and reconfiguring a target root file system and kernel

- deploying target system images onto target systems or virtual machines

Creating a root File System for Target Systems

To create a target system image, 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 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. Note that the **sudo (8)** command can alternatively be used to run Architect.

You can start Architect from the command line or it can be launched from the GNOME or MATE desktop environments as follows:

1. as root user, from the command line:

architect

2. as non-root user, from the command line:

sudo -E architect

3. for GNOME and MATE desktops environments from the desktop's applications launcher search for Architect and click on it. If you are not logged in as the root user, you will be prompted for a password.

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

Scroll down and press the Accept button to continue.




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 1-6 for details.

A session can be saved at any time and loaded later to continue work on the target system image. To edit an existing session, click on the **Open** button. See “Editing and Saving an Existing Session” on page 1-6 for details.

A session and installer image can also be extracted from media created using Architect. See “Extracting a session from Architect media” on page 1-87 for more information. Click on the **Extract** button to use this tool.

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 1-6 for details.
- edit an existing session by clicking on the Open Session icon  or selecting **Open Session** in the **File** menu. See “Editing and Saving an Existing Session” on page 1-6 for details.

Place the cursor on an icon and a short description of its function shows up at the bottom of the page.

User Preferences

You may select your directory default preferences by selecting **Preferences** in the **Edit** menu. These defaults avoid having to enter a directory path each time.

The **Default session directory** is where you will find the session files created by Architect. The **Default image directory** is where the target system images are created. The **Default VM directory** and the **Default ISO directory** are the directories where, respectively, the virtual machine images and the disc installer ISO images are saved.

The following figure shows the default User Preferences settings:




Figure 1-4 User Preferences Page

If space is limited in your **/root** partition, you will want to change your defaults to a partition with more available disk space. In the following example, the default settings were changed to reside in the **/home/architect** directory.




Figure 1-5 User Preferences Modified

Creating a New Session

An Architect session file is used to store all the configuration choices for a target. It also contains data about recent deployments.

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 1-6 New Session Dialog

This dialog enables you to specify the version of RedHawk and the base distribution to be used for creating the target system image. Be sure that you have the correct version of the RedHawk and either the Rocky, Oracle or Red Hat Enterprise Linux media (or ISO files) necessary to create a target system image of the specified RedHawk release.

Editing and Saving an Existing Session

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

Selecting Software to Install in the Image

To select the software to install in the target 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 target 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.

Selecting Base Distribution Linux Packages

To select Rocky, Oracle or Red Hat packages for the target system image, click on the **Base Distro** selection from the **Select Software** toolbox.

Rocky, Oracle or Red Hat packages may be navigated by way of the “**Package View**” drop down menu on the top right hand of the page. Select the desired package view from the **Package View** 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.

Note that the following features are available in all the views except for the **Base Environments** view:

- 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-click 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 two packages selected as a group.

UEFI Support:				
Install	Package	Size (KB)	Description	
☐	efibootmgr	96	EFI Boot Manager (optional)	
☐	gdisk	713	An fdisk-like partitioning tool for	
☒	grub2-efi-x64	4957	GRUB for EFI systems.	
☒	shim-x64	4778	First-stage UEFI bootloader	
			☒ Select Package(s)	Return
			☐ Deselect Package(s)	Del
			Get Package Info	Alt+I
			Show Package Dependencies	Alt+D

Figure 1-7 Selecting Multiple Packages

Using the Base Environments View

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




Figure 1-8 Choosing Target Characterization, Base Environments View

For example, if the target will primarily be used to run a web server, choose the **Server** or **Server GUI** 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 Rocky, Oracle or Red Hat packages organized in a hierarchy of groups. The package group hierarchy is the standard Rocky, Oracle or Red Hat package group hierarchy. This view is shown in the following figure.




Figure 1-9 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 unchecking packages in that group in the lower pane. All the packages in a chosen group may be selected with the **Select Group** button.

All the packages that are in all the groups of the currently chosen package category may be selected with the **Select All Groups...** 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 provides a sorted linear list of all Rocky, Oracle or Red Hat packages.

The following figure shows the All Packages view.




Figure 1-10 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 provides a sorted linear list of all Rocky, Oracle or Red Hat packages that are currently selected for installation.

The figure below shows the Selected Packages view.




Figure 1-11 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 as in the above example. When this box is checked, only the optional packages are shown.

Selecting RedHawk OS Options

To select RedHawk Linux OS and kernels for the target system image, click on the RedHawk selection from the Select Software toolbox.

The RedHawk page is shown in the following figure.




Figure 1-12 Selecting RedHawk Options

Select which RedHawk kernel(s) to install by checking the appropriate check box(es): Standard, or Trace. 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 the debug kernel is no longer shipped but the user can build a custom kernel based on the debug kernel configuration. See “Configure Custom Kernel” on page 1-57. 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 (which includes the debug kernel).

Advanced users may wish to press the Select Individual 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 kernel debug information packages to install the kernel-matching packages with extra support for kernel debugging. These packages are required for RedHawk to be able to examine crash dumps.

Select Install development packages for building kernel modules to include headers needed to build 3rd party drivers.

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 1-13 Selecting NightStar Tools

By default no NightStar tools will be installed in the target image. Choose the **Install NightStar servers for remote analysis of target systems:** radio button if you wish to only install NightStar remote support for the various tools or 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 full NightStar tool suite (including GUI and servers) on target systems** 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 Tools** button located at the bottom right of page 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 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 target system image.

To configure the target 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
- SELinux

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 target system image, click on **General Settings** from the **Configure Image** toolbox.

The General Settings configuration page appears, as shown in the following figure.




Figure 1-14 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.

The Root Password is hashed and the hash value is saved in the session file. The actual password is not saved.

NOTE

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

In the System 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 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 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 target system image, click on **Console** from the **Configure Image** toolbox.

The Configure Console page appears, as shown in the following figure.




Figure 1-16 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 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 1-17 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 target system image, click on **Networking** from the **Configure Image** toolbox. The **Configure Networking** page appears, as shown in the following figure.




Figure 1-18 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 1-19 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 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 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 target system image, click on **File Systems** from the **Configure Image** toolbox. The **Configure File System** page appears, as shown in the following figure.




Figure 1-21 File System Configuration Page

There are two different partitioning modes supported: Simple configuration and Advanced configuration. This page defaults to the simple disk partitioning mode.

The following are the possible **Target Disk Device** names:

1. SATA disks: **sdX** where **X** can be from a-z. For example, **sda**.

2. NVMe disks: `nvmeXnY` where *X* can be from 0-9 and *Y* from 1-99. For example, **nvme0n1**.
3. eMMC disks: `mmcblkX` where *X* can be from 0-99. For example **mmcblk0**.
4. VirtIO disks: `vdX` where *X* can be from a-z. For example, **vda**.

NOTES

Architect attempts to match your disk device configuration. However, if it finds no matching disk device on the target system, Architect will map the configured disk device to the actual disk device found on the target system. When installing in interactive mode, Architect will show the mapping and prompt to confirm the chosen disk device.

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 1-76 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.

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 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 1-22 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 Partitioning** tab, Logical Volumes via the **LVM** tab, **RAIDs**, 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. You can also configure partitions and logical volumes with **LUKS2** encryption.

NOTE

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

NOTE

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

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




Figure 1-23 Advanced Disk Partitioning and Image Out-Of-Sync Notice

Default File System Configuration

Starting in the Configure File Systems page, click on the Advanced configuration button and the following menu appears the first time. Press OK for the default file system configuration.

Note that you can restart and reset the options anytime by pressing the Reset to Defaults... button at the bottom right of the Configure File System page.

All the partitions will be made with the same basic options. You can modify the settings later, on a per partition basis, using the Edit Partition button found in the Disk Partitioning tab.




Figure 1-24 File System Default Configuration

The Disk File Partitioning tab as shown below shows the disk partitioning for the basic default configuration.




Figure 1-25 Disk Partitioning with Basic Default Configuration

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 disk pull-down menu to choose the disk that you will be booting from if multiple disks have been defined.

Press Reset To Defaults to discard all the current settings and start from the beginning.

Optional File System Configurations

Optionally you can configure:

- the GPT partitioning table format

- logical volumes (LVM)
- the Linux Unified Key Setup (LUKS2) for encryption
- Redundant Array of Independent (or Inexpensive) Disks (RAID)

It is beyond the scope of this guide to fully document these features, however, documentation can be found on-line.

GPT stands for GUID Partition Table. It is a new standard that is gradually replacing MBR (Master Boot Record). It is required when using UEFI, which replaces the old BIOS.

LVM stands for Logical Volume Management. The main advantages of LVM are: increased abstraction, flexibility, and control in managing disk storage space. Volumes can be resized dynamically as space requirements change and volumes may be migrated between physical devices on a running system.

LUKS2 is the second version of the Linux Unified Key Setup for disk encryption management. LUKS enables you to encrypt block devices. It allows multiple user keys to decrypt a master key, which is used for the bulk encryption of the partition. Architect supports one passphrase per target for all LUKS2 configured partitions, however, multiple user keys may be configured by the user on the target system. Architect will prompt for the passphrase in the deployment stage.

A RAID may be constructed from any number of physical disk devices to form a single logical, large capacity storage device that offers a number of advantages over conventional hard disk storage device. It can improve performance, resiliency and costs, depending on the RAID level. I uses the techniques of striping, mirroring and parity.

A redundant bootloader is automatically created on all MSDOS disks and all GPT disks that are configured with a **/boot** partition and are used in a RAID. Optionally a BIOS_GRUB partition containing a bootloader may be added to any GPT disk.

Example Using Optional Configurations

For conciseness, the optional GPT table format, LVM and LUKS2 are all set in this example. Note that they can be chosen independent of each other.

When any of these options are set in the Default File System Configuration menu, all the system partitions will be made with these options. The settings can be modified at a later time, on a per partition or per logical volume basis, using the Edit Partition button in the Disk Partitioning tab or the Edit Logical Volume in the LVM tab.




Figure 1-26 File System Optional Configuration

When GPT, LVM, and LUKS2 encryption are all configured, the Disk Partitioning configuration looks as follows:




Figure 1-27 Disk Partitioning with GPT format, LVM and LUKS2 encryption

The LVM tab shows the logical volumes configured. The lock symbol on the left side of each logical volume shows that the LUKS2 encryption is enabled:




Figure 1-28 LVM tab showing LVM Logical Volumes and LUKS2 encryption

Press **Add Logical Volume** to add a new logical volume to the currently selected volume group.

Press **Remove Logical Volume** to remove the currently selected logical volume.

Press **Edit Logical Volume** to edit attributes of the currently selected logical volume.

Press **Edit Volume Group** to edit attributes of the currently selected volume group.

Steps in Configuring Additional Logical Volumes (LVM)

To setup additional logical volumes:

1. Configure a volume group. To add a new volume group, add a disk partition or a whole disk using the respective button in the **Disk Partitioning** tab. In the respective menus, set the **Use this partition for LVM** option and specify a **Volume Group** name.

If you are adding a logical volume to an existing volume group, then continue to step 2.

2. Configure the logical volumes. From the **LVM** tab, highlight the volume group by clicking on it and use the **Add Logical Volume** button to create a logical volume.

Example Configuring Logical Volumes (LVM)

In the next example, an additional disk partition will be configured for LVM. Click on the Use this Partition for LVM box and assign the Volume Group name 'DATA'. If instead of a disk partition, you want to use a whole disk, use the Add Disk, instead of the Add Disk Partition button.




Figure 1-29 Adding a new disk partition to be used for LVM

The following snippet shows a new partition, **sda2**, added in the Disk Partitioning tab.




Figure 1-30 New disk partition added to be used for LVM

The LVM tab will show a new volume group 'DATA' created:




Figure 1-31 New Volume Group 'DATA' created

From the LVM tab, click on the new volume group 'DATA' and press the Add Logical Volume to create a logical volume. In this example the mount point is set to '/private':




Figure 1-32 Adding logical volume 'private' to new volume group

The following snippet shows a new logical volume 'private' in the newly created volume group 'DATA':




Figure 1-33 Logical Volume 'private' created in new group 'DATA'

Redundant Array of Independent Disks (RAID)

A RAID device is a logical partition. It can either contain a file system or be used as an LVM physical volume.

RAID devices may be constructed from any number of physical disk devices, as long as the required minimum number of physical devices is met for a particular RAID level.

A RAID device can be made up of the following physical devices:

- a whole physical disk or
- a disk partition

All the physical devices used in a RAID must be of the same type; all whole disks or all disk partitions. If disk partitions are used, then all the partitions in a RAID must be on different disks.

If the **/boot** file system is on a RAID device then it must be configured for RAID level 1. This extends to the **root** file system if there is no **/boot** partition.

For Any GPT partition in a RAID that is used for the **/boot** file system, a BIOS_GRUB partition with a bootloader is automatically configured.

The only deployment methods that support RAIDs are the on-target installer methods: USB Installer, Disc Installer, and PXE installer. All other methods support a single disk only.

The following RAID levels are supported:

- RAID 0 : based on striping technique. This level does not provide fault tolerance but increases the system performance with higher read and write speeds. Minimum number of disk devices required is 2.
- RAID 1 : utilizes mirroring technique. This level provides fault tolerance in the loss of no more than one member disk and, in some cases, increases read speed. Minimum number of disk devices required is 2.
- RAID 10 (0+1) : based on the combination of striping and mirroring techniques. This level exhibits RAID 0 performance and RAID 1 fault tolerance. Minimum number of disk devices required is 4.

- RAID 5 : utilizes both striping and parity techniques. This level provides the read speed improvement as in RAID 0 and survives the loss of one RAID member disk. Minimum number of disk devices required is 3.
- RAID 6 : similar to RAID 5 but uses two different parity functions. This level can sustain two simultaneous RAID disk failures and still continue to function. Minimum number of disk devices required is 4.

Not supported at this time:

- The partitioning of a RAID
- Booting from a RAID
- Nesting of RAIDs

Steps in configuring a RAID

Steps to configure a RAID:

1. Configure the RAID by adding the appropriate number of physical devices. For example, for a level 1 RAID, this will be at the minimum two partitions or two whole disks. Use either the Add Partition... or Add Disk... buttons respectively from the Disk Partitioning tab.

Click respectively on Use this Partition for RAID or Use this disk for Raid button and give the RAID a name. Use the same RAID name when configuring all the physical devices in the RAID. If the RAID is using partitions, you must also specify the size of the partition.

NOTE

All the physical devices used by a RAID must be of the same type; either all whole disks or all partitions. If using partitions, all the partitions must be on different disks.

When adding a new disk for the purpose of creating a partition to be used for a RAID device, make sure to choose Create Partitions when adding the disk and Use this partition for RAID when adding the partition. The option Use this Disk for RAID should only be used when the entire disk is to be allocated to a RAID.


2. Configure the RAID device from the RAID tab. Double click on the specific RAID and, on the Edit RAID menu that pops up, specify the RAID level and click on your preference of: Use this RAID for a file system or Use this RAID for LVM.

If you choose to use a file system, you will be asked to enter file system related information. If you choose to use LVM, then you must enter the desired volume group name or use the default.

3. Configure the logical volumes from the LVM tab if you chose Use this RAID for LVM in the previous step. Click to highlight the volume group

name and then click on the Add Logical Volume button and enter the file system information requested.

When a RAID has the incorrect number of physical devices configured for the specific RAID level, the following error message will appear at the bottom of the Disk Partitioning screen. Switching over to the RAID tab, the error printed there will show the number of physical devices expected. When the RAID configuration is completed with the correct number of physical devices, the errors disappear.



✖ One or more RAID devices has an invalid configuration.
⚠ Some deployment methods cannot partition multiple disks.
⚠ If the size of any disk device is exceeded, you will get an error when the system is installed.

Figure 1-34 Invalid RAID configuration error message

Example System Configuration Using RAID Partitions

In this example, two RAIDs will be configured. For each RAID, two partitions will be allocated; one on the disk labeled **sda** and one on **sdb**.

The two RAIDs in this example are:

1. boot-raid: the RAID will be configured for RAID level 1 and a file system; both are requirements for the system's **/boot** partition.
2. lvm-raid: The RAID will be configured for RAID level 0. It will use logical volumes (LVM) for the system's **swap** and **/** directory.

To start, remove the current system configuration by clicking on the disk and using the Remove Disk button:




Figure 1-35 Remove system disk sda

Next, create two disks using the Add Disk... button. On the menu that pops up, click on the Create partitions on this disk and choose the partition table format of your choice. For this example, the MSDOS partition table format is used.

The image below shows the disk **sda** already created and now adding the disk **sdb**.




Figure 1-36 Adding disk sdb

Now add a partition to each disk using the **Add Partition...** button. Configure each partition for RAID and assign the partition size and a name for the RAID. In this example the size is 1 GiB and the name of the RAID is 'boot-raid'. Use the same size and RAID name for all the physical devices in the RAID.

In the following image, the boot-raid partition on disk **sda** has already been created with the same options as shown below for **sdb**.

The following snapshot shows the partition 'boot-raid' created on disk **sdb**:




Figure 1-37 Adding the boot-raid RAID partition on disk sdb

The same as with 'boot-raid', add another partition on each disk to be used for the RAID labeled 'lvm-raid'.

The following image shows the 'lvm-raid' partition has already been created on disk **sda** and now it is being created with the same options on disk **sdb**:




Figure 1-38 Adding the lvm-raid partition on disk sdb

The following snippet shows the Disk Partitioning tab after all the partitions have been created. Note that the RAID level always defaults to 0 when adding a disk or partition but it can be changed when configuring the RAID in the next step.

Disk Partitioning					
Configure disks and disk partitions here. A disk may be configured with a partition table and partitions, or be used entirely for LVM or RAID. Each partition may contain a file system or be used for LVM or RAID.					
Disk	Partition	Size	Mount Point	Type	Options
sda (MSDOS)					
The bootloader will be installed in the MBR of this disk.					
	sda1	1 GiB		RAID	"boot-raid" (RAID 0)
	sda2	maximum		RAID	"lvm-raid" (RAID 0)
All data on this disk will be overwritten.					
sdb (MSDOS)					
	sdb1	1 GiB		RAID	"boot-raid" (RAID 0)
	sdb2	maximum		RAID	"lvm-raid" (RAID 0)
All data on this disk will be overwritten.					

Figure 1-39 Disk Partitioning snippet after RAID partitions created

Next configure the RAIDs. Double click on the 'boot-raid', and on the **Edit RAID** menu that pops up, change the RAID level to 1 as it is the only run level that can be assigned to a

`/boot` system partition. It also requires a file system. Its mount point is `/boot` as shown below:




Figure 1-40 Configuring the RAID named boot-raid

Double click on the 'lvm-raid' to now configure the second RAID. Set the Use this RAID for LVM button. Use the RAID level 0 which is the default setting. The volume group name is 'redhawk' which is also the default setting.

The image below shows the Edit RAID menu settings for the 'lvm-raid' RAID:




Figure 1-41 Configuring the RAID 'lvm-raid' for LVM

The following snippet shows the RAID configuration. The RAID named 'boot-raid' is configured as a file system while 'lvm-raid' is configured for LVM.




Figure 1-42 Snippet of the RAID tab showing the RAIDs created

The final step is to configure logical volumes in the volume group 'redhawk' from the LVM tab. First add a logical volume for **swap**. The size is set to 8 GiBs and the file system type is **swap**.




Figure 1-43 Configuring the swap logical volume

Next, add a logical volume for the **root** file system and allocate to it the remaining space in the volume group.




Figure 1-44 Configure the root logical volume

The snippet that follows shows the contents of the LVM tab:

Volume Group	Logical Volume	Size	Mount Point	Type	Mount Options
▼ redhawk					
	swap	8 GiB		swap	
	root	maximum	/	xfs	defaults

Figure 1-45 Logical volumes swap and root configured

Example Default Autogenerated RAID Configuration

A level 1 RAID default configuration can be automatically created using the **Reset to Defaults...** button of the RAID tab. This 2-disk default configuration can be used as a starting point for larger RAIDs.

You are then prompted to choose from file system configuration options. Click the Use Software RAID option to create the RAID. You may also choose any or all of the other configuration options from the menu. In this example all the options were chosen, including the GPT partition table.





Figure 1-46 Default File System Configuration

The default RAID configuration that is created is shown in the snippet below.



The Disk Partitioning tab shows the default disks, **sda** and **sdb**, created in the RAID named ‘RAID 1’:




Figure 1-47 Disk devices for RAID ‘RAID 1’created

In the LVM tab, the default LVM volumes configured are shown with LUKS encryption.




Figure 1-48 LVM configuration with LUKS encryption

Example Configuration Using multiple RAID levels

In this example, four disks are added and partitioned. The partitions are shared by three different RAIDs, each configured at different RAID levels: one uses level 10, one level 5 and one level 6.

Each RAID is using the minimum physical devices required for the run level: four partitions for RAID levels 10 and 6, and three partitions for RAID level 5.

The four disks are created using the Add Disk... button and the Create partitions on this disk setting. The partitions are created using the Add Partition... button

and the Use this partition for RAID setting. The RAID level is configured from the RAID tab.

The image below shows the disk partitioning for this configuration. Note that this image is taken after the RAIDs are configured with their respective RAID levels, otherwise it would show 0 for all the RAID levels.




Figure 1-49 Disk partitioning for the three RAIDs configured

The image below shows the RAIDs configured.




Figure 1-50 RAID tab showing three RAIDs using different RAID levels

Read-Only Root Configuration

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 that are described below.

The first step is to configure the root file system as read-only via the **Edit Partition** button. Check the box that reads **Mount read-only**. Then, click again on the **Configure Read-only Root Settings** link for information on the next steps.

Temporary system storage is required when root (/) is configured as read-only. The files that require to be writable are specified in **/etc/rwtab** and **/etc/rwtab.d/***.

By default the system will create a RAM-based file system for storage that will be mounted on the target as **/var/lib/stateless/writable**. By default the maximum size this RAM-base file system can grow is to 50% of the size of RAM.

When root has been specified as read-only, the following dialog will be displayed when you click on **Read-Only Root Settings** button. Use the up and down arrows of the spin box to change the percentage **Maximum RAM Usage** setting.




Figure 1-51 Modifiable Maximum RAM Usage setting

Alternatively, to avoid using RAM, add a partition with the mount point **/var/lib/stateless/writable**, using the **Add Partition** button.

While these two temporary storage 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**. Files and directory paths listed in **/etc/statetab**

and `/etc/statetab.d/*` will automatically be bind-mounted from `/var/lib/stateless/state` to remain persistent over boots.

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 (**writable**) and the optional persistent disk partition (**state**), both under the `/var/lib/stateless/` directory.




Figure 1-52 Example root read-only partitioning scheme

Configuring Special File Systems (tmpfs, bind)

The Special File Systems tab is used to configure special (non-disk) file systems. See the mount (8) man page for more information on these special 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**.




Figure 1-53 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 1-54 Example All File Systems list

Configuring SELinux

The security-enhanced option is disabled in RedHawk but it can be enabled in the permissive or enforcing mode.

See the section “Configuring SELinux”, in Chapter 2: Security Extensions for more information.

Building an Image

To build the target 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 1-55 Build Image Page

Choose a directory in which to build the target 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 target 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 target 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 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 specific media, as shown in the following figure. Follow the directions to load the media, then click OK to begin.




Figure 1-56 Build Prompt to Insert Rocky Updates Media

NOTE

The 9.6 system release installation discs for Rocky, Oracle and Red Hat are created using Blu-ray technology and require a Blu-ray drive to read the data correctly.

When OK is selected, the Rocky 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 1-57 Status of Rocky 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 Rocky installation is complete, you are prompted to insert the RedHawk Linux disc:




Figure 1-58 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 below.




Figure 1-59 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 target 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
- Additional RPMs

- System Services
- Kernel
- 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 target system image, click on Software Updates in the Customize Image toolbox. The Software Updates page appears, as shown in the following figure.




Figure 1-60 Install Concurrent-RT Software Updates Page

Installing updates can be done from local media (Discs 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 1-61 Software Updater Dialog

Additional RPMs

To install additional RPMs into the target 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 1-62 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 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 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.

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 1-63 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 system image are available for customization on the **System Services** page.

You can bring up a search box by typing on the page CTRL-F or the '/' character, as noted at the bottom of the page.

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

Kernel

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** from the **Customize Image** toolbox.

The Customize Kernels page appears, as shown in the following figure.




Figure 1-64 Customize Kernels Page

The Customize Kernels page allows you to perform different functions with the kernel configuration in the target image.

The Default Kernel to Boot 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 system image immediately.

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

NOTE

When making changes to the boot options line, press the TAB key at the end of the line to have the boot options immediately evaluated and propagated to the PXE diskless pages.

At the bottom of this page there are functions to configure and build 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 system 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 1-65 Custom Kernel Dialog

This window shows the status of running the **ccur-config** command in the target system 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 1-66 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 1-67 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 1-68 Kernel Build Stages

Finally, once the entire build and install process is complete, you will be given the option to clean the source tree. At this point the build process is complete.




Figure 1-69 Kernel Build Completed

The custom kernel automatically becomes the default kernel to boot. If this choice is not desired, change the kernel to boot using the Default Kernel to Boot pull-down menu 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.

File Manager

To copy files to and from the host file system and the target system image, click on File Manager from the Customize Image toolbox. Manage Files In Image page appears, as shown in the following figure.




Figure 1-70 Manage Files In Image Page

The buttons just below the Host File System and Target System Image headings help you traverse directories in the host and target system, create new directories and do file name searches. To see a menu of the file system operations supported, right click while in either the target or host system boxes.

The default mode does not allow the host file system to be modified. To allow modifications, click on the bottom left of the page on the box labeled Allow host file system to be modified.

Chroot Shell

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




Figure 1-71 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.

The following options are enabled by default in the chroot but can be disabled by clicking on the check-marked box:

- Allow local X server to accept X connections launched from the shell. This option will allow any X application to run on the chroot.
- Mount DNF repositories from imported media ISOs. This option will automatically mount the ISOs you have imported in the build and configure the repositories so that you can run **dnf** operations inside the chroot, on any package in those ISOs.
- Mount /proc, /sys, and /dev from host (read-only) inside the chroot shell. This option will mount the mentioned file systems from the host in the chroot. The file systems are mounted read-only so that no changes can be made on the host system.





Figure 1-72 chroot Shell

This provides a shell with the root directory being the target system image directory. All changes made to system files (including software installed or removed) will be done in the target 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 target 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. **Clean up Image** page appears, as shown in the following figure.

**Figure 1-73 Clean Up Image Page**

Select the types of files to remove from the target 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 system images can be deployed onto target boards in several different ways with RedHawk Architect. Note that it is also possible to deploy target system images without the use of any target hardware with the **Virtual Machines** deployment method.




Figure 1-74 Deploy Image

- **USB Devices.** USB devices can be directly flashed with the target 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 “Deploying to USB Devices” on page 1-66 for more information.
- **USB Installer.** 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 “Installing via USB drive” on page 1-68 for more information.
- **Disc Installer.** Disc media installers can be created with the target system image on it. Architect creates a bootable installation disc that will boot on the target and install the root file system onto the target board’s local media. Once complete, the disc is removed and the board will boot the image upon a restart. See “Installing via Disc media” on page 1-72 for more information.
- **PXE Installer/PXE Diskless.** Architect can deploy the target system image over the network. It can deploy an installer that will install the target system image onto the target board’s local media, or it can deploy the target system image via NFS for fully diskless booting. See “Installing via PXE over a Network” on page 1-74 for more information on network installation, and see “Booting Diskless via PXE over a Network” on page 1-76 for more information on diskless booting.
- **Virtual Machines.** Architect can deploy the target system image directly to a virtual machine image that can be booted, via QEMU, directly

on the host. See “Deploying to Virtual Machine” on page 1-86 for more information.

The UEFI firmware target configuration is currently supported by all deployment methods. 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, USB Installer and Disc Installer), there is no UEFI configuration box as those methods work with either UEFI or BIOS systems.

Deploying to USB Devices

To copy a target 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 Devices** button.

This will display the **Deploy to USB Devices** page, as shown in the following figure.




Figure 1-75 Deploy to USB Devices Page

The **Flash USB Device...** button allows you to copy a target 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 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 1-76 Flash Device Confirmation

Press OK to confirm the operation and then the copy will begin, as shown in Figure 1-77.




Figure 1-77 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.

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 1-79.




Figure 1-78 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 1-79 Flash Copy Completed

Installing via USB drive

To create a bootable USB drive that will install the target 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 1-80 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.

Check the box labeled Automatically install image to disk when USB drive is booted to create a USB drive that will install the target system image onto a target board's local disks without any prompting or user interaction.

NOTE

This will destroy any data on the target system's local disks and therefore automatic installations should be used carefully; however, it is useful on systems without an attached console display or configured and connected serial console.

Check the box labeled Encrypt installation image (password required to install) if encryption of the target image is desired. You will be prompted for a password which will be required when installing the target. Note that the

Automatically install image to disk and the Encrypt install image options cannot both be set.

Also note that if there is more than one disk on the target system, Architect will take an educated guess at the most appropriate disk to use based on your configuration. For example, if you've configured an **sda** disk, it will choose a SATA disk over a NVMe disk. Nevertheless, it is suggested that you use the interactive installation method if you do not want to run the risk of over-writing a disk.

In an automatic installation you are still given 10 seconds to intervene. If you hit any key during the 10 seconds you will be presented with an interactive menu.



Figure 1-81 Automatic installation on target without user interaction

The default mode is an interactive installation. In this mode, the installation on the target will stop at the page shown below until a choice is entered. Press ENTER to continue with the interactive installation.




Figure 1-82 Interactive installation on Target

In an interactive installation you will be able to map disks configured in Architect with the physical disks on the target system.

The screen that follows shows two disks were configured in Architect but three disks were found on the target system. In this example, the two disks configured in Architect, **sda** and **sdb**, have been reassigned from the mapping suggested by Architect.

```

Would you like to change any disk assignments? (y/n) [n] y

Which disk assignment would you like to change? (sda/sdb) sda
Choose a local disk to use for "sda":
* 1) 465.8G ATA WDC WD5003ABYX-0 (/dev/sda)
  2) 1.8T ATA ST2000NM0008-2F3 (/dev/sdb)
  3) 29.3G ASolid USB (/dev/sdc)
: 2

The following disk assignments will be used for this installation:

sda => 1.8T ATA ST2000NM0008-2F3 (/dev/sdb)
sdb => 465.8G ATA WDC WD5003ABYX-0 (/dev/sda)

Would you like to change any disk assignments? (y/n) [n] y

Which disk assignment would you like to change? (sda/sdb) sdb
Choose a local disk to use for "sdb":
* 1) 465.8G ATA WDC WD5003ABYX-0 (/dev/sda)
  2) 1.8T ATA ST2000NM0008-2F3 (/dev/sdb)
  3) 29.3G ASolid USB (/dev/sdc)
: 3

The following disk assignments will be used for this installation:

sda => 1.8T ATA ST2000NM0008-2F3 (/dev/sdb)
sdb => 29.3G ASolid USB (/dev/sdc)

Would you like to change any disk assignments? (y/n) [n]

```

Figure 1-83 Remapping of configured disks to physical disks on target

When you have finished with the disk assignments, the installation will resume without further user interaction.

Installing via Disc media

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

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




Figure 1-84 Deploy via Installation Disc Page

Choose **Create Disc** tab to directly burn a disc that will install the target system image onto disc media. No ISO image will be saved on disk in this mode.

Choose **Create ISO Image** tab to create an ISO file that contains an installer image. This ISO image can be later burned to a disc, or it may be useful with other tools or for long term storage.

Choose **Burn ISO Image** tab to burn a previously created ISO image to disc media.

Check the box labeled **Automatically install image to disk when disc is booted** in the **Create Disc** tab or the **Automatically install image to disk when ISO is booted** in the **Create ISO Image** tab, to create a disc or ISO that will install the target system image onto a target system without any prompting or user interaction. In the “Installing via USB drive” section, starting on page 1-69, there is a full description of this option.

Check the box labeled **Encrypt installation image (password required to install)** if encryption of the target image is desired. You will be prompted for a password which will be required when installing the target. Note that the **Automatically install image to disk** and the **Encrypt install image** options cannot both be set.

NOTE

DVD and Blu-ray are the currently supported disc media. Writing a bootable installer image onto a CD disc is not supported at this time.

The size of the compressed image is printed during the process when you press **Make Installation Disc** from the **Create Disc** tab and also when you press **Make ISO Image** from the **Create ISO Image** tab.

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.

Installing via PXE over a Network

RedHawk Architect can deploy a target system image to a target system over an Ethernet network connecting the host machine to the target machine.

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.

Installation of the root file system is performed by first creating a PXE-bootable installation image. The PXE boot agent used for legacy systems is `pxelinux` while for modern UEFI systems `grub2` is used. A target machine can boot this installation image via PXE, which will then remotely copy the target system image into the target's local drive media.

To create a PXE-bootable installation image that will install the target 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 1-85 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.

Check the box labeled **Automatically Install Image to disk when target system is PXE-booted** to create a PXE image that will install the target system image onto a target system without any prompting or user interaction. In the “Installing via USB drive” section, starting on page 1-69 there is a full description of this option.

Check the box labeled **Encrypt installation image (password required to install)** if encryption of the target image is desired. You will be prompted for a password which will be required when installing the target. Note that the **Automatically install image to disk** and the **Encrypt install image** options cannot both be set.

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.

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 1-86 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

When you select **PXE Diskless** from the toolbox on the left side of the RedHawk Architect main window, the **Deploy to Diskless Systems** page will be displayed.

The **Deploy to Diskless Systems** page provides two different implementations for booting diskless. The first implementation **NFS** uses **NFS**, the second labeled **RAMDISK** uses a **Live RAMDISK**.

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.

With NFS, the target machine boots a diskless image via PXE, which will then mount the target system image via NFS. In a Live RAMDISK boot, the entire target system image 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 writable 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 writable 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.

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.

To create a PXE-bootable diskless image that will mount the target system image via NFS select the NFS tab and the page, shown in the figure below, will appear.




Figure 1-87 Initial PXE NFS Diskless Deployment Page

When the RAMDISK tab is selected, a different but similar page will appear as shown in the figure below.




Figure 1-88 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** and the **Target PXE/DHCP Interface** 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 4-3 for more information.

When choosing to create an NFS Diskless Image, Press **Make NFS Diskless Image** to begin building the named PXE diskless image.

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 1-89 PXE NFS Diskless Image Building Complete

Press **Okay** 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.

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 1-90 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 Section "Managing PXE Targets" on page 4-7 for more information.

Remote Sync

Synchronization between a target system and the target system image on the host is supported by the **Remote Sync** function found in the toolbox on the left side of the RedHawk Architect main window.

Changes made on the target can be applied to the target system image (sync from target) and changes made in the target system image can be applied to the target (sync to target).

Following is a summary of the steps:

1. Configure the target for syncing in the PXE Target Manager page.
2. Reboot the target.
3. Select the target in the Synchronize with Remote Targets page.
4. Click on the desired sync action button in the Synchronize with Remote Targets page.

To configure the target for synchronization, from the Architect PXE Target Manager page, click on the desired target and use the **Edit Target** button to select **Start remote sync process** as shown below:

Editing target 'speedy'.

Hostname: speedy

IP Address: 192.168.1.6

MAC Address: C8:F7:50:FC:47:BA

On Next Boot

- ☐ Boot local disk
- ☐ Install local disk with: RAM-analyzer.diskless
- ☐ Boot diskless with: RedHawk-8.2 (NFS)
- ☒ Start remote sync process

Cancel OK

Figure 1-91 Configure target 'speedy' for remote sync

Press OK and you will see the menu showing the target has been configured to sync on the next boot as shown below:

Architect PXE Target Manager

PXE-bootable target systems must be configured here before deploying PXE installation or diskless images.

Hostname	IP Address	MAC Address	Next Boot	Status
speedy	192.168.1.6	C8:F7:50:FC:47:BA	sync	
sphynx	192.168.1.11	C8:1F:66:06:84:A8	install (RAM-analyzer.diskless)	

1 target set to be installed on next PXE boot.

Close

Figure 1-92 Target 'speedy' configured for remote sync

Reboot the target system and you will see a splash screen with a count down of 10 seconds. If you hit any key during those 10 seconds, you will be presented with the menu shown below. Note that the default is a non-interactive sync as no interaction is really required but, for this example, an interactive sync was chosen.




Figure 1-93 Sync Menu with interactive selected

On the target system you will see some start-up messages and then the 'waiting for host' message. Note that you do not have to wait for this message as the host and target will wait for each other.



Figure 1-94 Waiting on host

Now on the host system, go to the Remote Sync Deployment Page and verify that the target's name appears in the Target to Sync with: box and then press on the desired action: Sync From Target or Sync to Target.

In Sync To Target, the target will be updated with changes made in the target system image. Files added or deleted in the target system image will be added or deleted on the target system.

In Sync From Target, the target system image will be updated with changes made in the target system. Files added or deleted in the target system will be added or deleted on the target system image.

For this example, Sync To Target, was selected:




Figure 1-95 Sync To Target ‘speedy’

Once you pressed one of the sync buttons on the host, you will get a prompt on the target asking if you wish to continue (if in interactive sync mode). Type ‘y’ to continue.

When the sync completes, the messages on the target system will look similar to the ones on the following page:

```
Start Remote Sync from Architect to continue.
(waiting for host...)

This tool will synchronize the file systems on this system
with the Architect image 'RedHawk-8.2'.

Do you wish to continue? (y/n) y

Mounting disk file systems.

Syncing target with Architect image 'RedHawk-8.2'.
(waiting for host to finish sync...)
Host finished syncing.
Unmounting disk file systems.

Syncing is complete!

Removing PXE config file (0A861EAC) from server.

Press ENTER to reboot. _
```

Figure 1-96 Synchronization complete; messages on the target

After the synchronization, the target system must be rebooted. This happens automatically in the non-interactive mode while in the interactive you have to press ENTER to reboot the system.

In this example, a new software package was installed in the target system image and it was applied to the target via the Sync to Target function. The output log below shows the **rpm** related files added to the target system:




Figure 1-97 Synchronization Complete Messages on Host

Deploying to Virtual Machine

To deploy a target 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 1-98 Deploy to Virtual Machines Page

The **Deploy to Virtual Machines** Page defaults to the **CREATE VM** tab. In this mode, pressing the **Make VM Image** button will simply create a virtual machine image file from the session's target 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 VM** tab 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 **Disk Image to Boot** text field or the file **Browse** button.

Selecting the **Sync VM** tab allows you to perform file synchronization between the target system image and the virtual machine disk image in both directions:

- Use the **Sync From VM** button to update the target system image with all file changes that have been made inside the booted virtual machine image.

- Use the Sync To VM button to update the virtual machine with all the changes that have been made in the target system 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 to VMM 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.

Selecting Export to KVM-RT allows you to utilize a previously created virtual machine image with KVM-RT. RedHawk KVM-RT is a Real-Time Hypervisor solution utilizing QEMU/KVM and RedHawk Linux real-time kernel features.

Extracting a session from Architect media

Session Extraction is a feature in Architect that allows users to extract a session and image from media that was previously created using Architect.

Session extraction is currently supported from an Architect generated installer disc (CD, DVD or Blu-ray) or an Architect generated ISO installer image. Extracting from a USB stick may be supported in the future.

NOTE

The Session Extraction tool is available in the 9.2.3 release and later releases. It can extract images generated in the 9.2 release and later.

To extract a session from a disc or ISO image, click on Session Extraction from the Tools Menu or click on the  icon.

Session extraction is also available from the **Welcome** page when Architect is first started:



Figure 1-99 Extract option available from the Welcome page

You are then presented with a page similar to the one below:




Figure 1-100 Session extraction from an Architect-generated Disc

Verify that the defaulted disc drive is correct and the disc is in the drive. Then choose a **Session Name** and **Image Name** and press **Extract**.

Note that all the directory paths specified come from User Preferences (See “User Preferences” on page 1-5). You may change the paths by using the **Browse** button or change the User Preferences defaults.

Press the **ISO** TAB to extract from an ISO file. A page similar to the one above appears except that instead of a **Device to Extract From:**, it will prompt you for an **ISO File to Extract From:**. Type the path or use the **Browse** button to find the ISO image. Choose a **Session Name** and **Image Name** and press **Extract**.

Once the extraction completes successfully, the user is given the option to open the session just extracted:




Figure 1-101 Prompt to open the session just extracted

Following are two example applications of the session extraction tool:

- The user deploys a target system image to disc to deploy to target(s) or for safe storage (disc or ISO). If the Architect host is damaged or the session data is unrecoverable, the user can use the extraction tool to restore the target system image from the media generated prior to the corruption.
- In some cases, Concurrent Real-Time generates the system using Architect and the customer only receives the installed system along with an Architect image installation disc. Architect can be used to extract the target installation image from the disc, make changes and redeploy the changes.

Security Extensions

Architect includes support for enhancing the security of deployed target systems. This includes configuring, creating, and deploying target system images with FIPS, SELinux and SCAP security policy enabled.

UEFI Secure Boot is also supported, however, Architect configuration is only required when deploying the target to a virtual machine.

Note that SCAP support is provided in the Advanced Security Edition of Architect by an optional package named **ccur-architect-security**. If this package is not installed on the system, the SCAP security extension will not be available. An example using the SCAP DISA STIG profile is included as an Appendix. See “DISA STIG Example” in Appendix C.

The following instructions are specific to enabling the various security extensions in a target image. These are in addition to the instructions in the previous chapters.

UEFI Secure Boot

Secure Boot is a UEFI firmware security feature that ensures only immutable and signed software are loaded during the boot time. To boot a target in UEFI Secure Boot mode, follow the instructions in the “UEFI Secure Boot” Appendix of the RedHawk Release Notes.

Architect only needs to know about UEFI Secure Boot when deploying a virtual machine image so that it can simulate firmware changes. In the case of booting a virtual machine target, check the **Secure Boot** box in the **Create VM** tab of the **Deploy to Virtual Machines** page as shown in the image below.




Figure 2-1 Secure Boot Configuration for virtual machines

Configuring SELinux

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. See figure below.




Figure 2-2 SELinux enabled in the enforcing mode

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 as shown below.




Figure 2-3 Target image out of sync with the session

SELinux is supported on both the NFS and RAM diskless implementations. For NFS, Architect uses NFSv4.2 which allows SELinux security contexts to pass between an NFS server and client.

Constraints of the SELinux NFS diskless implementation:

1. The host must have SELinux enabled even if just in permissive mode.
2. If `/var/lib/tftpboot` gets relabeled on the host then the PXE diskless image of the target must be recreated to reset its root context. This can happen if SELinux's autorelabel feature is enabled on the host.

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.

Customizing the kernel with FIPS

To enable support for the Federal Information Processing Standards (FIPS) in the kernel, the boot option `fips=1` must be set in the **Extra Boot Options** text field of the **Customize Kernels** page. This must be set for each kernel to be booted on the target. FIPS customization is supported by all the deployment methods.

Select the default kernel to boot and then add the `fips=1` kernel boot option as shown in the image that follows.

NOTE

When booting diskless, verify that `fips=1` is also listed in the **Extra Kernel Options:** field in the **PXE Diskless** deployment page.




Figure 2-4 Customizing kernel with support for FIPS

To verify that FIPS is enabled on the target, run:

fips-mode-setup --check

If fips mode is not enabled on the target as expected, verify the following:

- the kernel listed in the Default Kernel to Boot in the Kernel page of the Customize Image toolbar matches the kernel that is running on the target.
- fips=1 is set in the Extra Boot Options field for the kernel configured as the Default Kernel to Boot.
- if a diskless deployment, that fips=1 is also specified as extra kernel options in the PXE Diskless deployment page.

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.

Custom profiles can also be derived from existing profiles using the SCAP Workbench graphical tool. This is often referred to as SCAP content tailoring and is further explained in the "Customizing SCAP Content Using SCAP Workbench" on page 2-15.

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 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.

Red Hat also has problems posted to their issue-tracking system, Bugzilla (<https://bugzilla.redhat.com>). You can search on that site for existing bugs related to any of the SCAP packages.

SCAP System Requirements

SCAP support is provided by an optional package (**ccur-architect-security**) which is included in the Advanced Security Edition of Architect. If this package is not installed, SCAP functionality will be missing and the SCAP menu options will not be visible in the Architect GUI.

Optionally, depending on your site's requirements, you may choose to perform a system update on the host and target systems. We do recommend to get the latest SCAP updates on the host as follows:

dnf update scap-security-guide

Instructions to update the target's SCAP packages are given later in this document after the target's image is built.

SCAP Workflow

Workflow Overview

Note that step by step instructions in performing a SCAP evaluation are given in the example provided in Appendix C, "DISA STIG Example".

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

1. Configure 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's pre-installation rules, before the target system image is built.
2. Build the target image. A post-build SCAP remediation scan is done automatically in the target system image when the image is built.
3. Customize the target image. Some customization may be needed and, optionally, you can run additional post-build SCAP scans 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 post-deployment scans. These scans may be evaluation scans or remediation scans. 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.

NOTE

Some SCAP pre-installation rules require that additional file systems be configured. Therefore, the diskless deployment methods may not be a good choice for some SCAP policies since those methods ignore all file system configuration settings and instead use a single mount point for the root '/' file system.

Each step is explained further in the sections that follow.

1. Configure SCAP security policy

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.

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.

In the figure below the following default SCAP settings are shown for a session building a RedHawk 9.6 (RHEL) target system image.




Figure 2-5 Configure SCAP Security Policy page

The default SCAP content file in the above figure provides a choice of benchmarks and profiles. The SCAP content and tailoring files are listed at the bottom of the page. Use the **Change** button on the right hand corner of the page to load a different content file or a tailoring file. For information on tailoring files, see "Customizing SCAP Content Using SCAP Workbench" on page 2-15.

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. Note that each

one has an (info) link also. For each pre-installation 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.

The Benchmark and SCAP content file are initially set by Architect. The profile must be selected. In this example, we have selected the “DISA STIG with GUI for Red Hat Enterprise Linux 9” profile as shown in the figure below. The DISA STIG profile is designed for safeguarding the Department of Defense (DoD) IT network and systems. Note that the instructions that follow are the same for any profile.

NOTE

A complete example using the “DISA STIG with GUI “ profile is included as an appendix. See “DISA STIG Example” in Appendix C.

Following are the corresponding SCAP settings when building a RedHawk 9.6 (RHEL) target system image and selecting the "DISA STIG with GUI for RHEL 9" profile.

content file: `/usr/share/xml/scap/ssg/content/ssg-rhel9-ds.xml`

benchmark: Guide to the Secure Configuration of Red Hat Enterprise Linux 9

profile: DISA STIG with GUI for Red Hat Enterprise Linux 9




Figure 2-6 DISA STIG with GUI profile selected

If you have a tailoring file, use the **Change...** button to include it before starting the pre-installation fixes. The tailoring file is explained in the section "Customizing SCAP Content Using SCAP Workbench" on page 2-15.

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 that an additional file system be created.

In the case where additional file systems are required, Architect will redirect the user to the **Configure File Systems** page so that the user may manually add the requested partition. Once finished adding the partition, the user is returned back to the **Configure SCAP Security Policy** page.

Note that (info) links are also available for the rules.

Use the **Fix All...** button on the bottom right of the screen to automate the fixing of all the rules shown. If something needs manual intervention Architect will redirect the user as stated above and afterwards continue to the next rule.

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

2. Build the target system image

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 2-7 Architect's post-build remediation completes

The post-build auto-remediation scan in the chroot of the target system image on the host will most likely complete with failures as in the figure above. This is expected as usually not all SCAP rules are fixed by auto-remediation in a chroot shell.

The results of the scan can be viewed and further scans can be run from the SCAP Scanner tab of the Customize Image toolbox. See the section “3. Customize the target image” that follows.

3. Customize the target image

Optionally, you may update the target system. We do recommend that the SCAP packages be updated. In the Chroot Shell from the Customize Image toolbox update the the packages as follows:

```
subscription-manager register  
dnf update openscap openscap-scanner
```

The results from previous remediation and evaluation scans can be viewed from the SCAP Scanner page of the Customize Image toolbox. Both evaluation and remediation scans on the target's chroot can be run from this page. Also a scan can be started remotely from the host on a deployed target using the Launch SCAP Workbench GUI... button. The SCAP Scanner page is shown below.




Figure 2-8 SCAP Scanner page

After customizing the build image, the user may proceed to deploy it.

4. Run post-deploy scans

Evaluation and remediation scans can be performed by executing commands on the target system. Post-deploy scans can be done in various ways which are listed below and explained in the sections that follows.

1. Directly, logging in to the console of the target system, and running local commands.
2. Remotely, logging in from another system via **ssh**, and running local commands.
3. Remotely, from Architect, using the **scap-workbench** tool which uses an **ssh** connection.

The required files are stored under the **/root/scap** directory of the target. The scan commands use the SCAP content file which can also be found in the **/root/scap** directory.

Post-deployment reports from system scans, are generated as .html and .txt files and placed in a subdirectory under the directory **/root/scap/results** on the target system.

NOTE

scap-workbench is an optional package and has its own way of generating reports, hence, you will not find those reports in the **/root/scap/results** 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 the last scans done in Architect are found under the **/root/scap/results** directory.

NOTE

Reports generated are overwritten each time a scan is initiated. If reports need to be saved, make sure to move the **/root/scap/results** directory to a new name before initiating a new scan.

Directly from console

Login to the console directly and perform any system configuration that may be required on the target. Then execute the scans as follows:

```
# login as root. You may be asked to change your passwd
cd /root/scap
./run-remediate-scan
# View the reports. Manual remediation may be necessary
./run-eval-scan
```

Remotely via ssh

To run an evaluation scan and/or a remediation scan, **ssh** into the target system as a non-root user and perform any system configuration that may be required on the target follows:

```
ssh -X scapuser@target
sudo -s
cd /root/scap
./run-remediate-scan
# View the reports. Manual remediation may be necessary
./run-eval-scan
```

Remotely from Architect

SCAP Workbench is a graphical tool that can be used to perform SCAP scans on a target system. Scans can be done remotely from the Architect host over an **ssh** connection.

SCAP Workbench requires a non-root user with NOPASSWD **sudo** privileges.

To execute a remote scan of a target system using SCAP Workbench, follow the steps below.

1. Launch SCAP Workbench from the button labeled Launch SCAP Workbench GUI... found in the SCAP Scanner page of the Customize Image toolbox of Architect. The SCAP content file, and the optional tailoring file, used in the session will be loaded into SCAP Workbench.
2. Select the Profile of interest at the top of the page.
3. Click on Remote Machine (over SSH) and enter the User and host of the target. Click on the user is sudoer box.

IMPORTANT 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.

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.

Click the Scan button. Note that if you are using multiple profiles you will have to repeat steps 2 through 6 for each profile.

NOTE

scap-workbench may fail with “cannot open display :0”. This is an issue where DISPLAY is hard-coded to :0 in **ssh-askpass**. See "Scap-workbench: remote scans fail with cannot open display :0" in the Release Notes for help.

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




Figure 2-9 SCAP Workbench with settings for Remote Machine

Customizing SCAP Content Using SCAP Workbench

Customization, also called tailoring, allows existing SCAP policies to be customized without having to rewrite the policy. Tailoring is useful to change or exclude certain security rules in a profile. You can deselect rules that do not apply to your company, select optional rules not selected by default, or change modifiable settings in the policy. All these policy changes are then saved in the tailoring file. Architect can then load the tailoring file and apply the changes in policy.

NOTE

Tailoring files must be created outside of Architect and be available to be used by Architect during the SCAP configuration step, before the build image is created.

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

1. Run the **scap-workbench** command-line tool on the host system.
2. Select the content file of interest and press the **Load Content** button.
3. Select the **Profile** of interest.
4. Click on the **Customize** button to the right of the **Profile** box. A new window will be displayed.
5. Enter the new profile ID. Choose the name provided or follow the instructions specified for naming the profile.
6. Make the desired changes by making the desired modifications. Then press the **OK** button to close the window.
7. 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 profile being created with the FIPS rule “Ensure ‘/etc/system-fips’ exists” selected.




Figure 2-10 Using SCAP Workbench to create a customized profile

Once saved, the tailoring file can be used in the SCAP configuration. In the Configure SCAP Security Policy page, click on the Change... button at the right bottom of the page and enter the tailoring file which, in this example, was saved as `/root/Documents/ssg-rhel9-ds-tailoring.xml` file. See the following image.




Figure 2-11 SCAP Workbench, customized (tailoring) 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.

Importing ISO Images

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

Importing ISO Images

Normally when building a target 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 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 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 Rocky 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.

NOTE

The 9.6 system release installation discs for Rocky and Oracle are created using Blu-Ray technology and require a Blu-ray drive to read the data correctly.

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 Rocky 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 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 as in the figure above, 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 Appendix A, 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. The PXE Image Manager will remain empty until PXE images are created using the PXE deployment tools in the **Deploy Image** toolbox.

The PXE Image Manager lists all of the installation images that have been deployed by the PXE Installer tool. The Remote Sync Preferences allows you to change global parameters affecting remote syncs.




Figure 4-4 PXE Image Manager.

Clicking on the Edit button of the Architect PXE image Manager page will bring up the Edit Remote Sync Preferences page shown below. Note that changes are not applied to a particular target but globally.




Figure 4-5 Global parameters for Sync Operations

The following dialog shows an example of the PXE Image Manger with installation images.




Figure 4-6 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 target 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 target 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 Manager with diskless images.




Figure 4-7 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 target 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 target 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 4-8 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.


The image shows a dialog box titled "Add PXE Target" with a close button (X) in the top right corner. The dialog contains the text "Configure new PXE target system." followed by three input fields: "Hostname:", "IP Address:", and "MAC Address:". Below these fields is a section titled "On Next Boot" which contains four radio button options: "Boot local disk" (which is selected), "Install local disk with:" (followed by a dropdown menu), "Boot diskless with:" (followed by a dropdown menu showing "RedHawk-8-RAM (RAMDISK)"), and "Start remote sync process". At the bottom right of the dialog are two buttons: "Cancel" and "OK".

Figure 4-9 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 1-74 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 1-76 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 4-10 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.

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 (192.168.1.0/255.255.255.0). Hostnames for the targets will be generated automatically.

Starting IP Address: 192.168.1.20

How many targets would you like to configure? 10

Base Hostname: target-

Apply

On Next Boot

☒ Boot local disk

☐ Install local disk with: RAM-analyzer.diskless

☐ Boot diskless with: RedHawk-8.2 (NFS)

☐ Start remote sync process

Hostname

IP Address

MAC Address

Cancel

OK

Figure 4-11 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 4-8 for more information.

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

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 (192.168.1.0/255.255.255.0). Hostnames for the targets will be generated automatically.

Starting IP Address: 192.168.1.20

How many targets would you like to configure? 10

Base Hostname: target-

Apply

On Next Boot

- ☒ Boot local disk
- ☐ Install local disk with: RAM-analyzer.diskless
- ☐ Boot diskless with: RedHawk-8.2 (NFS)
- ☐ Start remote sync process

Hostname	IP Address	MAC Address
target-020	192.168.1.20	
target-021	192.168.1.21	
target-022	192.168.1.22	
target-023	192.168.1.23	
target-024	192.168.1.24	
target-025	192.168.1.25	
target-026	192.168.1.26	
target-027	192.168.1.27	
target-028	192.168.1.28	
target-029	192.168.1.29	

Cancel OK

Figure 4-12 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 Appendix A, 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 dialog box is titled "Edit PXE Target" with a close button (X) in the top right corner. Below the title bar, it says "Editing target 'sphynx'". There are three text input fields: "Hostname:" with the value "sphynx", "IP Address:" with the value "192.168.1.11", and "MAC Address:" with the value "C8:1F:66:06:84:A8". Below these fields is a section titled "On Next Boot" containing four radio button options: "Boot local disk", "Install local disk with:" (selected), "Boot diskless with:", and "Start remote sync process". The "Install local disk with:" option has a dropdown menu showing "RAM-analyzer.diskless". The "Boot diskless with:" option has a dropdown menu showing "RedHawk-8.2 (NFS)". At the bottom right of the dialog are two buttons: "Cancel" with a red circle icon and "OK" with a green checkmark icon.

Figure 4-13 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 4-8 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.

Command Line Interface

This appendix describes the python scripts that can be run independent of the Architect GUI tool. Note that the scripts coordinate with the Architect tool by reading and updating the session file. These scripts are found under the directory: `/usr/lib/architect/cli-tools`.

build-image

Builds an architect target image based on settings specified in the session file. It will replace the root path in the session file if a different path was previously used.

Usage: `build-image [-n][-l log] root session`

Arguments:

root Specifies the root directory of the build image to be created.
session Specifies the architect session file to use.

Options:

-n Run non-interactively.
-l log Log output to the file log.
--help Display this help information.

Example:

```
cd /usr/lib/architect/cli-tools
sudo ./build-image -n -l /tmp/log.out /var/lib/architect/images/rocky-9.2 \
/root/Documents/sessions/rocky-9.2.session
```

chroot

Executes a command in a chroot shell for the *root* system image directory specified.

Usage: `chroot [-m][-k kernel-release] root [cmd]`

Arguments:

root Specifies the *root* directory of the image to chroot to.
cmd An optional command to execute. If not specified, *cmd* defaults to `/bin/bash`.

Options:

-m Bind-mount read-only the host system's `/proc`, `/sys`, and `/dev` directories.
-k [kernel-release] Fake the kernel release name in `/bin/uname`, necessary when building kernels in the chroot. The **kernel-release** format should resemble the output from the command `uname -r`; for example `6.1.19-rt8-RedHawk-9.2-trace`.
--help Display this help information.
--version Display software version.

Examples:

```
cd /usr/lib/architect/cli-tools
sudo ./chroot /var/lib/architect/images/rocky9.2 blscfg
sudo ./chroot -m -k 6.1.19-rt8-RedHawk-9.2-trace \
/var/lib/architect/images/rocky-9.2
```

make-pxe-diskless-image

Builds a PXE-bootable diskless image for an existing architect target system image based on settings specified in the *session* file. If it is the first time a PXE diskless image is built for this session, the script will use the corresponding GUI page defaults.

The settings can be modified using the GUI or editing the session file directly. Note that if you choose to modify the settings using the GUI, you must also build the PXE diskless image and then save the session.

NFS is the default diskless type. In order to switch to RAMDISK, you must build a RAMDISK diskless image using the GUI and save the session or edit the session file directly.

Usage: `make-pxe-diskless-image [-n][-l log] session pxe-name`

Arguments:

session Specifies the architect session file to use.
pxe-name Specifies the name of the PXE diskless image to be created.

Options:

-n Run non-interactively.
-l log Log output to the file **log**.
--help Display this help information.

Example:

```
cd /usr/lib/architect/cli-tools
sudo ./make-pxe-diskless-image -n -l /tmp/log.out \
/root/Documents/rocky9.2.session rocky9.2-NFS
```

make-pxe-install-image

Builds a PXE-bootable installation image for an existing architect target system image based on settings specified in the *session* file. If it is the first time a PXE installation image is built for this session, the script will use the corresponding GUI page defaults.

The settings can be modified using the GUI or editing the session file directly. Note that if you choose to modify the settings using the GUI, you must also build the PXE install image and then save the session.

Usage: `make-pxe-install-image [-n][-l log][-L LUKS-passphrase] [-E encrypt-password] session pxe-name`

Arguments:

session Specifies the architect session file to use.
pxe-name Specifies the name of the PXE installation image to be created.

Options:

-n Run non-interactively.
-l log Log output to the file **log**.

-L passphrase

Use this passphrase for any LUKS volumes that are defined in the session file.

-E password

Use this password to encrypt the image.

--help

Display this help information.

Example:

```
cd /usr/lib/architect/cli-tools
```

```
sudo ./make-pxe-install-image -n -l /tmp/log.out -L londonbridgeisfallingup\  
-E WxYZ1976 /root/Documents/rocky9.2.session rocky9.2-pxe
```

setup-pxe

Configures PXE services on the current host. When executed with no arguments it will print the available subnets. The *dhcp_subnet* is specified in the form *ip_address/mask*.

Usage: **setup-pxe [-d] *dhcp_subnet***

Arguments:

dhcp_subnet

Specifies which subnet on the host should be used when PXE booting from architect.

Options:**-d**

Do not configure DHCP server on this host.

--help

Display this help information.

Examples:

```
cd /usr/lib/architect/cli-tools
```

```
sudo ./setup-pxe
```

```
sudo ./setup-pxe -d 192.0.2.0/255.255.255.0
```


DISA STIG Example

This appendix covers using RedHawk Architect's Advanced Security Edition to defense-harden a target that has been created and deployed by Architect. This appendix also provides guidance to running post-deployment scans and some suggestions for manual remediations.

It is assumed the user is familiar with the Architect tool and has read the SCAP chapter. See “Security Content Automation Protocol (SCAP)” on page 2-5.

DISA STIG System Requirements

In this example, the following are assumed and recommended but not required:

- the target image uses the RHEL distribution
- the *host* system is installed with a RHEL distribution
- the latest RHEL updates are installed on both the host and the target image
- the latest RedHawk updates are installed

Based on the above recommendations, note the following:

- a Red Hat Enterprise Linux 9.6 installation disc is needed to import the software into Architect.
- a Red Hat software subscription is needed to update the host and target to the latest Red Hat releases.

Additionally, below are some target configuration settings for DISA STIG. While you may choose other configuration settings, these have been tested and proven to work.

BIOS:

If the Intel SW Guard Extension (SGX) is enabled in your BIOS firmware, we recommend you disable it or the scans will fail due to the `/dev/sgx*` device node found unlabeled with a proper SELinux type.

New Session Prompt:

Select a RedHawk Linux Release: This choice is made at the first prompt when you start a new session. Choose any RedHawk (RHEL) release.

Base Distro/Base Environments:

Choose Workstation with all the add-ons packages listed on the right selected.

Base Distro/Select Base Distribution Packages:

Select the group for UEFI Support. Optionally, you may select other packages.

Select Software/RedHawk:

Select the **ccur-selinux** package. Optionally you can select all the RedHawk packages.

Configure Image/General Settings:

Change System Run Level: 5 Graphical

Configure Image/File Systems:

Change Partition Table Format to GPT.

DISA STIG Workflow Overview

Note that before starting Architect, it is recommended that you update the SCAP packages on the host as follows:

dnf update scap-security-guide

The following steps are typically performed to run DISA STIG scans on a target system:

1. Configure DISA STIG security policy
2. Build the target image
3. Customize the target image
4. Run post-deployment scans
5. Manual remediation
6. Sync from Target (Optional)

Configure DISA STIG Security Policies:

Select the “DISA STIG with GUI for Red Hat Enterprise Linux 9” profile, add a SCAP tailoring file if you previously created one (see “Customizing SCAP Content Using SCAP Workbench” on page 2-15) and then press the **Fix All** button at the bottom right of the page.

Some manual intervention is required to add partitions. Instead of using the **Add Partition** button, an easier way is to select a partition in the **Disk partitioning** box before or after where you want to add the new partition. Then click the right mouse button and you are prompted to choose the placement of the new partition; before or after the partition highlighted. Once the placement is selected you are presented with the **Add Disk Partition** menu. Note that Architect prints a SCAP line indicating the file system

that needs to be configured at the bottom of the Configure File Systems page. Add all the partitions as requested.

Once DISA STIG has been configured, the Configure File Systems page will look similar to the one below. The user-specified partition sizes and type may differ from this example. Unless an option is specifically requested, the rules will automatically add the required options for the partitions.




Figure C-1 Example File System Page after configuring SCAP

Build the Image

Click on the Build Image button from the Build Image toolbox to start the build. The target image will be built and a remediation scan will be automatically run at the end of the build.

Customize the Target Image

It is recommended that you update the following SCAP packages in the Chroot Shell of the Customize Image toolbox.

```
subscription-manager register
dnf update openscap openscap-scanner
```

To install the RedHawk Software Updates bring up the **Software Updates** page of the **Customize Image** toolbox, make a selection from the **Update Options** and press the **Install Updates** button.

Some additional customization may be required depending on how you plan to log in to the target system after it is deployed. You may choose one way or both:

1. login as root on the target system console
2. login remotely as a non-root user via **ssh**

Logging in as root on the target system console is preferable as the remediation scan disables the entry for the non-root user in the **/etc/sudoers** file and you must replace it after each remediation scan.

If you are planning to login as root using a console connected with USB devices (keyboard, mouse), disable the **usbguard** service now. Bring up the **System Services** page of the **Customize Image** toolbox, find the **usbguard** service and disable it (clear the check mark). Once the target is deployed, you will be instructed to re-enable the **usbguard** service back after some configuration.

If you plan to login remotely via **ssh**, use the **Chroot Shell** option of the **Customize Image** toolbox to add a non-root user and after add the user to the sudoers file. Note that the user should be added after all remediation scans are completed on the chroot of the target system as the remediation scan removes the user from the **/etc/sudoers** file. In this example we use *stiguser* to represent the non-root user's login name:

```
useradd stiguser -d /home/stiguser
passwd stiguser
echo "stiguser    ALL=(ALL)    NOPASSWD:    ALL" >> /etc/sudoers
```

NOTE

STIG-compliant password rules are complex and the rules can be found online. The STIG-compliant password may be configured in the chroot or you may be prompted to change it when you log in to the target. Note also that STIG enforces strict login lockouts.

You may also run additional scans or view results from previous scans from the **SCAP Scanner** option of the **Customize Image** toolset. However, note that the most profitable scanning and remediation will happen on the actual target system.

The target may now be deployed. If you configured a non-root user, verify that the entry for the user in the **/etc/sudoers** has not been removed by a remediation scan. If it has, make sure you add it before deploying the target image.

Run post-deployment scans

Post-deploy scans can be done in various ways which are listed below and explained in the sections that follow. Note that there is an example for **scap-workbench** in the SCAP

chapter, so there is only a reference provided here. Note also that some system configuration may have to be done prior to running the scans.

1. Directly, logging in to the console of the target system, and running local commands.
2. Remotely, logging in from another system via **ssh**, and running local commands.
3. Remotely, from Architect, using the **scap-workbench** tool which uses an **ssh** connection.

Directly from console

The instructions below will first update the RHEL system. Second, it will generate rules to exempt USB devices on the target, re-enable the **usbguard** service and verify that it is enabled and active. After a reboot, you may proceed to run scans on the system.

```
# login as root
# Run these only on the first boot
usbguard generate-policy > /tmp/rules.conf
install -m 0600 -o root -g root /tmp/rules.conf /etc/usbguard/rules.conf
# Verify the usbguard rules file is not empty and was installed correctly
# before starting the usbguard service
ls -l /tmp/rules.conf /etc/usbguard/rules.conf
systemctl enable --now usbguard.service
systemctl status usbguard.service
reboot
# login back in as root
# Ready to run scans
cd /root/scap
./run-remediate-scan
# Perform manual remediation steps. See "Manual Remediation" on page C-6
./run-eval-scan
```

NOTE

The above changes to the system, while persistent over boots, will be gone when you re-make your build image or installation image.

Remotely via ssh

The non-root user, noted as *stiguser* for this document, was appended to the **/etc/sudoers** file when the user was added. However, remediation scans will disable that entry, so it is crucial to add that entry back after each remediation scan.

The evaluation scan, in turn, will fail one of the sudo rules if it finds a user with the NOPASSWD set in that file. To avoid that failure the user must be removed from the sudoers file before running an evaluation and add it back after the evaluation scan. Alternatively, you can simply ignore that failure and just run the evaluation scan.

NOTE

If the user *stiguser* gets logged out of the system without root privileges specified in */etc/sudoers* (for example in an inactivity timeout), it will not be able to assume those privileges again unless you have root access to the console.

The commands that follow update the RHEL system. After a reboot then you can run scans. The code for *adduser-sudo.sh* and *rmuser-sudo.sh* scripts are included for you below the commands.

```
# Login via ssh
ssh -X stiguser@target
sudo -s
# Ready to run scans
cd /root/scap
# Run remediation scan and immediately after execute the adduser-sudo.sh script
./run-remediate-scan; ./adduser-sudo.sh
# Perform manual remediation steps here. See "Manual Remediation" on page C-6
# If you can ignore the NOPASSWD test failure, run the run-eval-scan script by itself
./rmuser-sudo.sh; ./run-eval-scan; ./adduser-sudo.sh
```

adduser-sudo.sh script:

```
#
# If there isn't one already, append an entry for stiguser in /etc/sudoers
#
grep ^stiguser /etc/sudoers > /devnull || \
    echo "stiguser    ALL=(ALL)        NOPASSWD:    ALL" >> /etc/sudoers
tail -2 /etc/sudoers
```

remove-sudo.sh script:

```
#
# remove lines starting with stiguser from the /etc/sudoers file
#
sed -i '/^stiguser/d' /etc/sudoers
tail -2 /etc/sudoers
```

Remotely from Architect

An example is provided in the SCAP section. Refer to "Remotely from Architect" on page 2-14.

Manual Remediation

Some manual remediation is required in order to get closer to the 100% passing mark. You can find the failures in the **fix-report.html** file that is generated when you run the remediation script. This file is placed under a subdirectory in the */root/scap/results* directory.

When you view the **fix-report.html** file and click on the link to a specific rule that failed, it will open a help page. In the Description section of the help page, you will find the description of the rule and sometimes instructions how to resolve the issue.

Following are some of the more obscure failures and suggestions on how to correct. For other failures, follow the directions given in the **fix-report.html**.

1. “Enable FIPS mode” failure. To fix run the following command:

blscfg --kopt-expand

2. “Ensure No Device Files are Unlabeled by SELinux”. If the device unlabeled is **/dev/sgx***, disable the Intel SW Guard Extension (SGX) in the BIOS.
3. “Ensure Chrony is only configured with the server directive” failure. To fix add a server entry with a time source in **/etc/chrony.conf** as follows:
server <ipaddress> maxpoll 16

NOTE

After manual remediations you can run more evaluation scans but note that a remediation scan after manual remediation may undo some of the fixes applied.

Sync from Target

This is an optional step to update the Architect's target system image with the changes that have been made on the target.

The user may want to preserve the target's updated image with the remediation fixes. A sync would be appropriate in this case. Conversely, the user may want to preserve the pure Architect image to install other systems. It is up to the user to decide. Refer to “Remote Sync” on page 1-81 for help in sync'ing from the target.

