

Red Hat Network Satellite 5.3.0 Reference Guide

Red Hat Network Satellite

Red Hat Network Satellite 5.3.0 Reference Guide

Red Hat Network Satellite

Edition 2

Copyright © 2010 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.









All other trademarks are the property of their respective owners.










1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588 Research Triangle Park, NC 27709 USA

Introduction to the Guide	xi
1. More to Come	xi
1.1. Send in Your Feedback	xii
1. Red Hat Network Overview	1
1.1. Update	2
1.2. Management	2
1.3. Provisioning	3
1.4. Monitoring	4
1.5. Errata Notifications and Scheduled Package Installations	4
1.6. Security, Quality Assurance, and Red Hat Network	5
1.7. Before You Begin	5
2. The <code>rhn_register</code> Client	7
2.1. Using <code>rhn_register</code>	7
2.1.1. Command-line version of <code>rhn_register</code>	13
3. Package Updater	15
3.1. Using the Package Updater	15
3.2. The Package Updater Applet	17
3.3. Updating Packages from the Command Line with <code>yum</code>	18
3.3.1. <code>yum</code> Commands	18
4. Red Hat Update Agent	21
4.1. Starting the Red Hat Update Agent	21
4.2. Registration	24
4.2.1. Registering a User Account	25
4.2.2. Activate	27
4.2.3. Channels	29
4.2.4. Packages Flagged to be Skipped	31
4.2.5. Available Package Updates	31
4.2.6. Retrieving Packages	33
4.2.7. Installing Packages	34
4.3. Command Line Version	36
4.3.1. Installing the Red Hat GPG key	39
4.3.2. Manual Package Installation	40
4.3.3. Synchronizing Your System Profile	40
4.3.4. Log File	41
4.4. Configuration	41
4.4.1. Using the Red Hat Update Agent Configuration Tool	41
4.4.2. Command Line Version	45
4.5. Registering with Activation Keys	46
4.6. Registering a System to an Organization	47
5. Red Hat Network Daemon	49
5.1. Configuring	49
5.2. Viewing Status	49
5.3. Disabling	49
5.4. Troubleshooting	49
6. Red Hat Network Alert Notification Tool	51
6.1. Configuring the Applet	51
6.2. Notification Icons	53
6.3. Viewing Updates	53

6.4. Applying Updates 54
 6.5. Launching the RHN Website 54

7. Red Hat Network Website 55

7.1. Navigation 55
 7.1.1. Entitlement Views 55
 7.1.2. Categories and Pages 56
 7.1.3. Errata Alert Icons 59
 7.1.4. Quick Search 59
 7.1.5. Systems Selected 59
 7.1.6. Lists 60
 7.2. Logging into the RHN Website 60
 7.3. Overview 62
 7.3.1. Your Account 63
 7.3.2. Your Preferences 64
 7.3.3. Locale Preferences 65
 7.3.4. Subscription Management 65
 7.3.5. Organization Trusts 67
 7.4. Systems 67
 7.4.1. Overview —  67
 7.4.2. Systems 67
 7.4.3. System Groups —  84
 7.4.4. System Set Manager —  87
 7.4.5. Advanced Search —  93
 7.4.6. Activation Keys —  95
 7.4.7. Stored Profiles —  98
 7.4.8. Custom System Info —  98
 7.4.9. Kickstart —  99
 7.5. Errata 112
 7.5.1. Relevant Errata 113
 7.5.2. All Errata 113
 7.5.3. Advanced Search 115
 7.6. Channels 117
 7.6.1. Software Channels 117
 7.6.2. Package Search 122
 7.6.3. Manage Software Channels 123
 7.7. Configuration 124
 7.7.1. Preparing Systems for Config Management 125
 7.7.2. Overview 125
 7.7.3. Configuration Channels 126
 7.7.4. Configuration Files 127
 7.7.5. Locally-Managed Files 128
 7.7.6. Systems 130
 7.8. Schedule 131

7.8.1. Pending Actions	131
7.8.2. Failed Actions	132
7.8.3. Completed Actions	132
7.8.4. Archived Actions	132
7.8.5. Actions List	132
7.9. Users — 	133
7.9.1. User List <input type="checkbox"/> Active — 	134
7.9.2. User List <input type="checkbox"/> Deactivated — 	139
7.9.3. User List <input type="checkbox"/> All — 	139
7.10. Monitoring — 	139
7.10.1. Probe Status — 	139
7.10.2. Notification — 	141
7.10.3. Probe Suites	143
7.10.4. Scout Config Push — 	144
7.10.5. General Config — 	145
7.11. Admin	145
7.11.1. Admin <input type="checkbox"/> Organizations	145
7.11.2. Admin <input type="checkbox"/> RHN Satellite Configuration	145
7.12. Help	146
7.12.1. Reference Guide	146
7.12.2. Satellite Installation Guide	146
7.12.3. Proxy Guide	147
7.12.4. Client Configuration Guide	147
7.12.5. Channel Management Guide	147
7.12.6. Release Notes	147
7.12.7. API	147
7.12.8. Search	148
8. Monitoring	151
8.1. Prerequisites	151
8.2. Red Hat Network Monitoring Daemon (rhnm d)	151
8.2.1. Probes requiring the daemon	152
8.2.2. Installing the Red Hat Network Monitoring Daemon	152
8.2.3. Configuring SSH	153
8.2.4. Installing the SSH key	154
8.3. mysql package	154
8.4. Notifications	155
8.4.1. Creating Notification Methods	155
8.4.2. Receiving Notifications	155
8.4.3. Redirecting Notifications	156
8.4.4. Filtering Notifications	157
8.4.5. Deleting Notification Methods	157
8.5. Probes	157

8.5.1. Managing Probes	158
8.5.2. Establishing Thresholds	158
8.5.3. Monitoring the RHN Server	159
8.6. Troubleshooting	159
8.6.1. Examining Probes with rhn-catalog	159
8.6.2. Viewing the output of rhn-runprobe	160
9. Multiple Organizations	163
9.1. Recommended Models for Using Multiple Organizations	163
9.1.1. Centrally-Managed Satellite for A Multi-Department Organization	163
9.1.2. Decentralized Management of Multiple Third Party Organizations	164
9.1.3. General Tips for Multi-Org Usage	165
9.2. Admin \square Organizations	166
9.2.1. Admin \square Organizations \square Details	167
9.3. Creating an Organization	168
9.4. Managing Entitlements	169
9.4.1. Admin \square Subscriptions \square Software Channel Entitlements	169
9.4.2. Admin \square Subscriptions \square System Entitlements	170
9.5. Configuring Systems in an Organization	170
9.6. Organizational Trusts	171
9.6.1. Establishing an Organizational Trust	171
9.6.2. Sharing Content Channels between Organizations in a Trust	172
9.6.3. Migrating Systems from One Trusted Organization to Another	173
9.7. Admin \square Users	176
9.7.1. Admin \square Organizations \square Details \square Users	176
10. Virtualization	177
10.1. Setting Up the Host System for Your Virtual Systems	177
10.1.1. Create a Kickstart Profile for the Guest Systems	177
10.1.2. Kickstart Your Host System	178
10.2. Setting Up Your Virtual Systems	182
10.2.1. Create a Kickstart Profile for the Guest Systems	183
10.2.2. Provision Your Guest Systems	184
10.3. Working With Your Virtual Systems	185
10.3.1. Logging into Virtual Systems Directly via SSH	185
10.3.2. Gaining Console Access Via the Host	185
10.3.3. Installing Software Via the Satellite Web Interface	186
10.3.4. Installing Software Via Yum From the Virtual System	186
10.3.5. Restarting Guests when Host Reboots	186
10.3.6. Deleting Virtual Systems	187
11. Cobbler	189
11.1. Cobbler Requirements	189
11.1.1. Using cobbler check	189
11.1.2. Configuring Cobbler with /etc/cobbler/settings	190
11.1.3. Cobbler and DHCP	190
11.1.4. Xinetd and TFTP	191
11.1.5. Configuring SELinux and IPTables for Cobbler Support	191
11.1.6. Syncing and Starting the Cobbler Service	192
11.2. Adding a Distribution to Cobbler	193
11.3. Adding a Profile to Cobbler	193
11.4. Adding a System to Cobbler	194
11.5. Cobbler Templates	194

11.5.1. Using Templates	195
11.5.2. Kickstart Snippets	195
11.6. Using Koan	196
11.6.1. Using Koan to Provision Virtual Systems	196
11.6.2. Using Koan to Re-install Running Systems	196
12. UNIX Support Guide	199
12.1. Introduction	199
12.1.1. Supported UNIX Variants	199
12.1.2. Prerequisites	199
12.1.3. Included Features	199
12.1.4. Differences in Functionality	200
12.1.5. Excluded Features	200
12.2. Satellite Server Preparation/Configuration	200
12.3. Client System Preparation	202
12.3.1. Download and Install Additional Packages	203
12.3.2. Deploying Client SSL Certificates	206
12.3.3. Configuring the clients	206
12.4. Registration and Updates	207
12.4.1. Registering Systems	207
12.4.2. Obtaining Updates	208
12.5. Remote Commands	211
12.5.1. Enabling Commands	211
12.5.2. Issuing Commands	212
A. Red Hat Network Registration Client	213
A.1. Configuring the Red Hat Network Registration Client	213
A.2. Starting the Red Hat Network Registration Client	215
A.3. Registering a User Account	218
A.4. Registering a System Profile	220
A.4.1. Hardware System Profile	220
A.4.2. Software System Profile	222
A.5. Finishing Registration	224
A.6. Entitling Your System	226
A.7. Text Mode RHN Registration Client	227
B. Command Line Config Management Tools	229
B.1. Red Hat Network Actions Control	229
B.1.1. General command line options	229
B.2. Red Hat Network Configuration Client	230
B.2.1. Listing Config Files	230
B.2.2. Getting a Config File	231
B.2.3. Viewing Config Channels	231
B.2.4. Differentiating between Config Files	231
B.2.5. Verifying Config Files	232
B.3. Red Hat Network Configuration Manager	232
B.3.1. Creating a Config Channel	233
B.3.2. Adding Files to a Config Channel	233
B.3.3. Differentiating between Latest Config Files	234
B.3.4. Differentiating between Various Versions	235
B.3.5. Downloading All Files in a Channel	235
B.3.6. Getting the Contents of a File	235
B.3.7. Listing All Files in a Channel	236

B.3.8. Listing All Config Channels	236
B.3.9. Removing a File from a Channel	236
B.3.10. Deleting a Config Channel	237
B.3.11. Determining the Number of File Revisions	237
B.3.12. Updating a File in a Channel	237
B.3.13. Uploading Multiple Files at Once	238
C. RHN API Access	239
C.1. Using the auth Class and Getting the Session	239
C.2. Obtaining the system_id	239
C.3. Determining the sid	239
C.4. Viewing the cid	239
C.5. Getting the sgid	239
C.6. Channel Labels	240
C.7. Sample API Script	240
D. Probes	243
D.1. Probe Guidelines	243
D.2. Apache 1.3.x and 2.0.x	244
D.2.1. Apache::Processes	244
D.2.2. Apache::Traffic	245
D.2.3. Apache::Uptime	245
D.3. BEA WebLogic 6.x and higher	246
D.3.1. BEA WebLogic::Execute Queue	246
D.3.2. BEA WebLogic::Heap Free	247
D.3.3. BEA WebLogic::JDBC Connection Pool	247
D.3.4. BEA WebLogic::Server State	248
D.3.5. BEA WebLogic::Servlet	248
D.4. General	249
D.4.1. General::Remote Program	249
D.4.2. General::Remote Program with Data	249
D.4.3. General::SNMP Check	250
D.4.4. General::TCP Check	251
D.4.5. General::UDP Check	251
D.4.6. General::Uptime (SNMP)	251
D.5. Linux	252
D.5.1. Linux::CPU Usage	252
D.5.2. Linux::Disk IO Throughput	252
D.5.3. Linux::Disk Usage	253
D.5.4. Linux::Inodes	253
D.5.5. Linux::Interface Traffic	254
D.5.6. Linux::Load	254
D.5.7. Linux::Memory Usage	255
D.5.8. Linux::Process Counts by State	255
D.5.9. Linux::Process Count Total	256
D.5.10. Linux::Process Health	256
D.5.11. Linux::Process Running	257
D.5.12. Linux::Swap Usage	258
D.5.13. Linux::TCP Connections by State	258
D.5.14. Linux::Users	259
D.5.15. Linux::Virtual Memory	259
D.6. LogAgent	260

D.6.1. LogAgent::Log Pattern Match	260
D.6.2. LogAgent::Log Size	261
D.7. MySQL 3.23 - 3.33	262
D.7.1. MySQL::Database Accessibility	262
D.7.2. MySQL::Opened Tables	262
D.7.3. MySQL::Open Tables	263
D.7.4. MySQL::Query Rate	263
D.7.5. MySQL::Threads Running	263
D.8. Network Services	264
D.8.1. Network Services::DNS Lookup	264
D.8.2. Network Services::FTP	264
D.8.3. Network Services::IMAP Mail	265
D.8.4. Network Services::Mail Transfer (SMTP)	265
D.8.5. Network Services::Ping	265
D.8.6. Network Services::POP Mail	266
D.8.7. Network Services::Remote Ping	267
D.8.8. Network Services::RPCService	267
D.8.9. Network Services::Secure Web Server (HTTPS)	268
D.8.10. Network Services::SSH	268
D.8.11. Network Services::Web Server (HTTP)	269
D.9. Oracle 8i, 9i, and 10g	270
D.9.1. Oracle::Active Sessions	270
D.9.2. Oracle::Availability	271
D.9.3. Oracle::Blocking Sessions	271
D.9.4. Oracle::Buffer Cache	271
D.9.5. Oracle::Client Connectivity	272
D.9.6. Oracle::Data Dictionary Cache	272
D.9.7. Oracle::Disk Sort Ratio	273
D.9.8. Oracle::Idle Sessions	273
D.9.9. Oracle::Index Extents	274
D.9.10. Oracle::Library Cache	274
D.9.11. Oracle::Locks	275
D.9.12. Oracle::Redo Log	275
D.9.13. Oracle::Table Extents	276
D.9.14. Oracle::Tablespace Usage	276
D.9.15. Oracle::TNS Ping	277
D.10. RHN Satellite	277
D.10.1. RHN Satellite::Disk Space	277
D.10.2. RHN Satellite::Execution Time	278
D.10.3. RHN Satellite::Interface Traffic	278
D.10.4. RHN Satellite::Latency	278
D.10.5. RHN Satellite::Load	279
D.10.6. RHN Satellite::Probe Count	279
D.10.7. RHN Satellite::Process Counts	279
D.10.8. RHN Satellite::Processes	280
D.10.9. RHN Satellite::Process Health	280
D.10.10. RHN Satellite::Process Running	281
D.10.11. RHN Satellite::Swap	281
D.10.12. RHN Satellite::Users	282

Reference Guide

E. Revision History	287
Index	289

Introduction to the Guide

Welcome to the *Red Hat Network Satellite 5.3.0 Reference Guide*. The *RHN Reference Guide* guides you through registering systems with Red Hat Network and using its many features.

Since Red Hat Network offers a variety of service levels, from the most basic Update module to the most advanced Monitoring package, some content of this guide may be inapplicable to you. This is particularly true of the RHN website, which displays selected categories, pages, and tabs depending on the entitlement level of the account used to log in. Refer to [Chapter 7, Red Hat Network Website](#) to determine what is available to you.

Depending on the version of Red Hat Enterprise Linux installed and the addition of new features, the **Red Hat Network Registration Client** and the **Red Hat Update Agent** may differ from the descriptions in this manual. Use Red Hat Network to update these applications before referring to the latest version of this manual.

All versions of this manual are available in HTML and PDF formats at <http://www.redhat.com/docs/manuals/satellite/>.



Warning

Systems running Red Hat Enterprise Linux 2.1 must use the **Red Hat Network Registration Client** before starting the **Red Hat Update Agent**. Refer to [Appendix A, Red Hat Network Registration Client](#) for instructions. Systems running Red Hat Enterprise Linux 3, 4, and 5 or later register with the **Red Hat Update Agent**. Refer to [Chapter 4, Red Hat Update Agent](#) for instructions.

For an overview of RHN Satellite offerings, please review the descriptions available at https://www.redhat.com/systems_management/ and <http://www.redhat.com/rhn/>.

1. More to Come

The *Red Hat Network Reference Guide* is constantly expanding as new Red Hat Network features and service plans are launched. HTML and PDF versions of this and other manuals are available within the **Help** section of the RHN Satellite website and at <http://www.redhat.com/docs/manuals/satellite>.



Note

Although this manual reflects the most current information possible, read the *RHN Release Notes* for information that may not have been available prior to the finalization of the documentation. The notes can be found on the RHN website and at <http://www.redhat.com/docs/manuals/satellite/>.

The following RHN documentation has been translated for this RHN Satellite release: RHN Satellite Reference Guide, RHN Satellite Installation Guide, RHN Client Configuration Guide, RHN Channel Management Guide, and RHN Satellite Release Notes. Translated documentation is available at <http://www.redhat.com/docs/> under **Red Hat Network Satellite**.

1.1. Send in Your Feedback

If you would like to make suggestions about the *Red Hat Network Satellite Reference Guide*, please submit a report in Bugzilla (<http://bugzilla.redhat.com/bugzilla>) against the component `Documentation_Reference_Guide` (Product: Red Hat Network Satellite, Version: 520).

Red Hat Network Overview

Have you ever read about a new version of a software package and wanted to install it but could not find it?

Have you ever tried to find an RPM through an Internet search engine or an RPM repository and been linked to an unknown site?

Have you ever tried to find an RPM but instead found only source files that you had to compile yourself?

Have you ever spent hours or even days visiting different websites to see if you have the latest packages installed on your system, only to have to do it again in a few months?

Those days are over with Red Hat Network (RHN). RHN provides the solution to all your system software management needs.

Red Hat Network is an Internet solution for managing a single Red Hat Enterprise Linux system or a network of Red Hat Enterprise Linux systems. All Security Alerts, Bug Fix Alerts, and Enhancement Alerts (collectively known as Errata Alerts) can be downloaded directly from Red Hat or your own custom collection. You can even schedule updates for delivery to your system immediately after release.

The main components of Red Hat Network are as follows:

- the **Red Hat Update Agent**
- the Red Hat Network website, whether this is hosted by the central RHN Servers, an RHN Satellite, or fed through an RHN Proxy Server
- Red Hat Network Daemon
- the **Red Hat Network Registration Client** - for systems running Red Hat Enterprise Linux 2.1 only.

The **Red Hat Update Agent (up2date)** provides your initial connection to Red Hat Network. Red Hat Enterprise Linux 3 and newer systems use the Red Hat Update Agent to register with RHN. Registration involves creating a unique RHN username and password, probing the hardware on your system to create a Hardware Profile, and probing the software packages installed on your system to create a Package Profile. This information is sent to RHN and RHN returns a unique System ID to your system. Once registered, the Red Hat Update Agent enables channel subscription, package installs, and management of System Profiles. See [Chapter 4, Red Hat Update Agent](#) for further information.

The Red Hat Update Agent, as the base component of RHN, is designed to manage a single system. It allows the system's superuser to view and apply Errata to the system. The RHN web interface facilitates the management, monitoring, and provisioning of a large deployment of systems, including the configuration of the Red Hat Update Agent for each system.

The **Red Hat Network Daemon (rhnstd)** runs in the background as a service and probes the Red Hat Network for notifications and updates at set time intervals (see [Chapter 5, Red Hat Network Daemon](#) for further information). This daemon is necessary in order to schedule updates or other actions through the website.

Red Hat Enterprise Linux 5 uses the `rhn_register` application documented in [Chapter 2, The `rhn_register` Client](#), while Red Hat Enterprise Linux 3 and 4 have registration functionality built into the **Red Hat Update Agent**.

Many Red Hat Network terms are used throughout this manual. As you read the *Red Hat Network Reference Guide*, refer to the [Glossary](#) as necessary for an explanation of common terms.



Tip

For a comparison chart of RHN service levels, refer to <http://www.redhat.com/rhn/compare/>.

1.1. Update

The RHN Update service is ideal for a user with one Red Hat Enterprise Linux system or a small number of Red Hat Enterprise Linux systems. Updated Subscription to Update can be purchased at <https://www.redhat.com/apps/commerce/rhn/>.

With each Update subscription, you receive the following services:

- **Download Software** — For customers who have purchased subscriptions to Red Hat Network, ISO images are available for immediate download.
- **Priority Access** during periods of high load — When Red Hat releases a large erratum, users with Priority Access can be guaranteed that they will be able to access the updated packages immediately.
- **RHN Support Access** — All paying customers of Red Hat Network receive web based support for their RHN questions.
- **Errata Notification, Multiple Systems** — Subscriptions for multiple systems means Errata notification for Errata to all of those systems. Note that only one email is distributed per each Erratum, regardless of the number of systems affected.
- **Errata Updates, Multiple Systems** — Get quick updates for multiple systems with an easy button click for each system.

1.2. Management

In addition to the features offered in the RHN Update subscription level, the RHN Management subscription service allows you to manage your network of Red Hat Enterprise Linux systems, users, and system groups through its **System Set Manager** interface.

RHN Management is based upon the concept of an organization. Each Management-level Red Hat customer has the ability to establish users who have administration privileges to system groups. An Organization Administrator has overall control over each Red Hat Network organization with the ability to add and remove systems and users. When users other than the Satellite Administrator log into the Red Hat Network website, they see only the systems they have permission to administer.

To create an account that can be used to entitle systems to RHN Management, go to <https://rhn.redhat.com/>¹ and click on the **Create Login** link under the **Sign In** fields. On the *Create a Red*

¹ <https://rhn.redhat.com/>

Hat Login page, click **Create a new Business Login**. After creating a business account, you may add users within your organization to it.

The Red Hat Network features available to you depend on the subscription level for each Red Hat Enterprise Linux system. With each Management subscription, you receive the functionality provided to Update users, plus:

- **Package Profile Comparison** — Compare the package set on a system with the package sets of similar systems with one click.
- **Search Systems** — Search through systems based on a number of criteria: packages, networking information, even hardware asset tags.
- **System Grouping** — Web servers, database servers, workstations and other workload-focused systems may be grouped so that each set can be administered in common ways.
- **Multiple Administrators** — Administrators may be given rights to particular system groups, easing the burden of system management over very large organizations.
- **System Set Manager** — You may now apply actions to sets of systems instead of single systems, work with members of a predefined system group, or work with an ad-hoc collection of systems. Install a single software package to each, subscribe the systems to a new channel, or apply all Errata to them with a single action.
- **Batch Processing** — Compiling a list of outdated packages for a thousand systems would take days for a dedicated sysadmin. Red Hat Network Management service can do it for you in seconds.

1.3. Provisioning

As the highest management service level, RHN Provisioning encompasses all of the features offered in the RHN Update and Management subscription levels. It is designed to allow you to deploy and manage your network of Red Hat Enterprise Linux systems, users, and system groups.

Like Management, Provisioning is based upon an organization. It takes this concept a step further by enabling customers with Provisioning entitlements to kickstart, reconfigure, track, and revert systems on the fly.

In addition to all of the features mentioned in lower service levels, Provisioning provides:

- **Kickstarting** — Systems with Provisioning entitlements may be re-installed through RHN with a whole host of options established in kickstart profiles. Options include everything from the type of bootloader and time zone to packages included/excluded and IP address ranges allowed. Even GPG and SSL keys can be pre-configured.
- **Client Configuration** — RHN Satellite Customers may use RHN to manage the configuration files on Provisioning-entitled systems. Users can upload files to custom configurations channels on the Satellite, verify local configuration files against those stored on the Satellite, and deploy files from the Satellite.
- **Snapshot Rollbacks** — Provisioning-level users have the ability to revert the package profile and RHN settings of systems. RHN Satellite customers can also roll back local configurations files. This is possible because snapshots are captured whenever an action takes place on a system. These snapshots identify groups, channels, packages, and configuration files.

- Custom System Information — Provisioning customers may identify any type of information they choose about their registered systems. This differs from System Profile information, which is generated automatically, and the Notes, which are unrestricted, in that the Custom System Information allows you to develop specific keys of your choosing and assign searchable values for that key to each Provisioning-entitled system. For instance, this feature allows you to identify the cubicle in which each system is located and search through all registered systems according to their cubicle.

1.4. Monitoring

Monitoring entitlements are available to RHN Satellite customers with Red Hat Enterprise Linux systems.

Monitoring allows an organization to install probes that can immediately detect failures and identify performance degradation before it becomes critical. Used properly, the Monitoring entitlement can provide insight into the applications, services, and devices on each system.

Specifically, Monitoring provides:

- Probes — Dozens of probes can be run against each system. These range from simple **ping** checks to custom remote programs designed to return valuable data.
- Notification — Alerts can be sent to email and pager addresses with contact methods identified by you when a probe changes state. Each probe notification can be sent to a different method, or address.
- Central Status — The results of all probes are summarized in a single **Probe Status** page, with the systems affected broken down by state.
- Reporting — By selecting a probe and identifying the particular metric and a range of time, you can generate graphs and event logs depicting precisely how the probe has performed. This can be instrumental in predicting and preventing costly system failures.
- Probe Suites — Groups of probes may be assigned to a system or set of systems at once rather than individually. This allows Administrators to be certain that similar systems are monitored in the same way and saves time configuring individual probes.
- Notification Filters — Probe notifications may be redirected to another recipient, halted, or sent to an additional recipient for a specified time based on probe criteria, notification method, scout or organization.

1.5. Errata Notifications and Scheduled Package Installations

You can configure Red Hat Network to send you email notifications of new and updated software packages as soon as the packages are available through RHN. You receive one email per Erratum, regardless of the number of affected systems. You can also schedule package installs or package updates. The benefits include:

- Reduced time and effort required by system administrators to stay on top of the Red Hat Errata list
- Minimized security vulnerabilities in your network through the application of updates as soon as Red Hat releases them

- Filtered list of package updates (packages not relevant to your network are not included)
- Reliable method of managing multiple systems with similar configurations

1.6. Security, Quality Assurance, and Red Hat Network

Red Hat Network provides significant benefits to your network, including security and quality assurance. All transactions made between your systems and Red Hat Network are encrypted and all RPM packages are signed with Red Hat's GNU Privacy Guard (GPG) signature to ensure authenticity.

Red Hat Network incorporates the following security measures:

1. Your System Profile, available at <http://rhn.redhat.com>, is accessible only with an RHN-verified username and password.
2. A Digital Certificate is written to the client system after registration and is used to authenticate the system during each transaction between the client and Red Hat Network. The file is only readable by the root user on the client system.
3. Red Hat signs all communications with an electronic signature using GPG. RPM can be used to verify the authenticity of the package before it is installed.
4. Red Hat encrypts all transactions using a Secure Sockets Layer (SSL) connection.
5. The Red Hat Quality Assurance Team tests and verifies all packages before they are added to the Red Hat Errata list and Red Hat Network.

1.7. Before You Begin

By default, all software packages necessary to access Red Hat Network are installed with Red Hat Enterprise Linux distributions. However, if you chose not to install them during the installation process, you must obtain the **Red Hat Update Agent (up2date)** and possibly the **Red Hat Network Registration Client (rhn_register)**. In Red Hat Enterprise Linux 3 and later, registration functionality is built into the **Red Hat Update Agent**, while Red Hat Enterprise Linux 2.1 users will need the **Red Hat Network Registration Client**.



Warning

The SSL certificate packaged with older versions of the **Red Hat Update Agent** and the **Red Hat Network Registration Client** reached its end of life August 28, 2003. Users attempting to connect using this certificate will receive SSL connection or certificate verification errors. You may view and obtain the versions of these applications containing new certificates at the *RHN Client Software*² page. In the RHN website, click **Help** at the top-right corner, **Get RHN Software** in the left navigation bar, and scroll down to examine the packages and versions.

To determine the versions of the client applications installed, run the `rpm -q` command followed by the package name. For instance, for the **Red Hat Network Registration Client**, type the following command:

```
rpm -q rhn_register
```

If the **Red Hat Network Registration Client** is installed, it will return something similar to:

```
rhn_register-2.9.3-1
```

The version number might differ slightly.

If you do not have the **Red Hat Network Registration Client** installed, the command will return:

```
package rhn_register is not installed
```

Perform this check for every package in [Table 1.1, “Red Hat Network Packages”](#) that is relevant to your system. Remember, only Red Hat Enterprise Linux 2.1 users need **Red Hat Network Registration Client**. If you prefer to use the command line versions, the two packages ending in **gnome** are not required..

Package Name	Description
rhn_register	Provides the Red Hat Network Registration Client program and the text mode interface
rhn_register-gnome	Provides the GNOME interface (graphical version) for the Red Hat Network Registration Client ; runs if the X Window System is available
up2date	Provides the Red Hat Update Agent command line version and the Red Hat Network Daemon
up2date-gnome	Provides the GNOME interface (graphical version) for the Red Hat Update Agent ; runs if the X Window System is available

Table 1.1. Red Hat Network Packages

The `rhnc_register` Client

Red Hat Enterprise Linux 5 features an application called `rhnc_register`. This application works with the `yum`-based RHN Hosted and RHN Satellite client called **Package Updater** (or `pup`) that replaces `up2date`. For more information about `pup`, refer to [Chapter 3, Package Updater](#).

The `rhnc_register` application normally runs as part of the `firstboot` configuration process just after installation. The first time a newly-installed Red Hat Enterprise Linux 5 system is booted, `firstboot` uses `rhnc_register` to register your system with RHN.

2.1. Using `rhnc_register`

If you should ever need to re-register your system at a later time (or you chose not to register during `firstboot`), you can use `rhnc_register` to do so. You can execute the command `rhnc_register` from the command line as root. If you have never registered, you can start `rhnc_register` by selecting **Applications** (the main menu on the panel) ▢ System Tools ▢ Package Updater. (You will be asked to enter the root password.) The Package Updater, when run on a system that has not yet been registered, triggers `rhnc_register` if there is no `/etc/sysconfig/rhn/systemid` file on the system.

If you have already registered before and `/etc/sysconfig/rhn/systemid` exists on the system, `rhnc_register` first asks if you are sure that you would like to register again. Doing so may create a duplicate system profile in RHN Satellite. Consider using `rhncreg_ks` and activation keys to re-register a system without creating a duplicate entry. Refer to [Section 7.4.2.9.1.4, "System Details ▢](#)

[Details ▢ Reactivation](#) —  " for more information.




Figure 2.1. Verifying Registration

If you are certain you would like to re-register this way, select the **Yes, Continue** button.

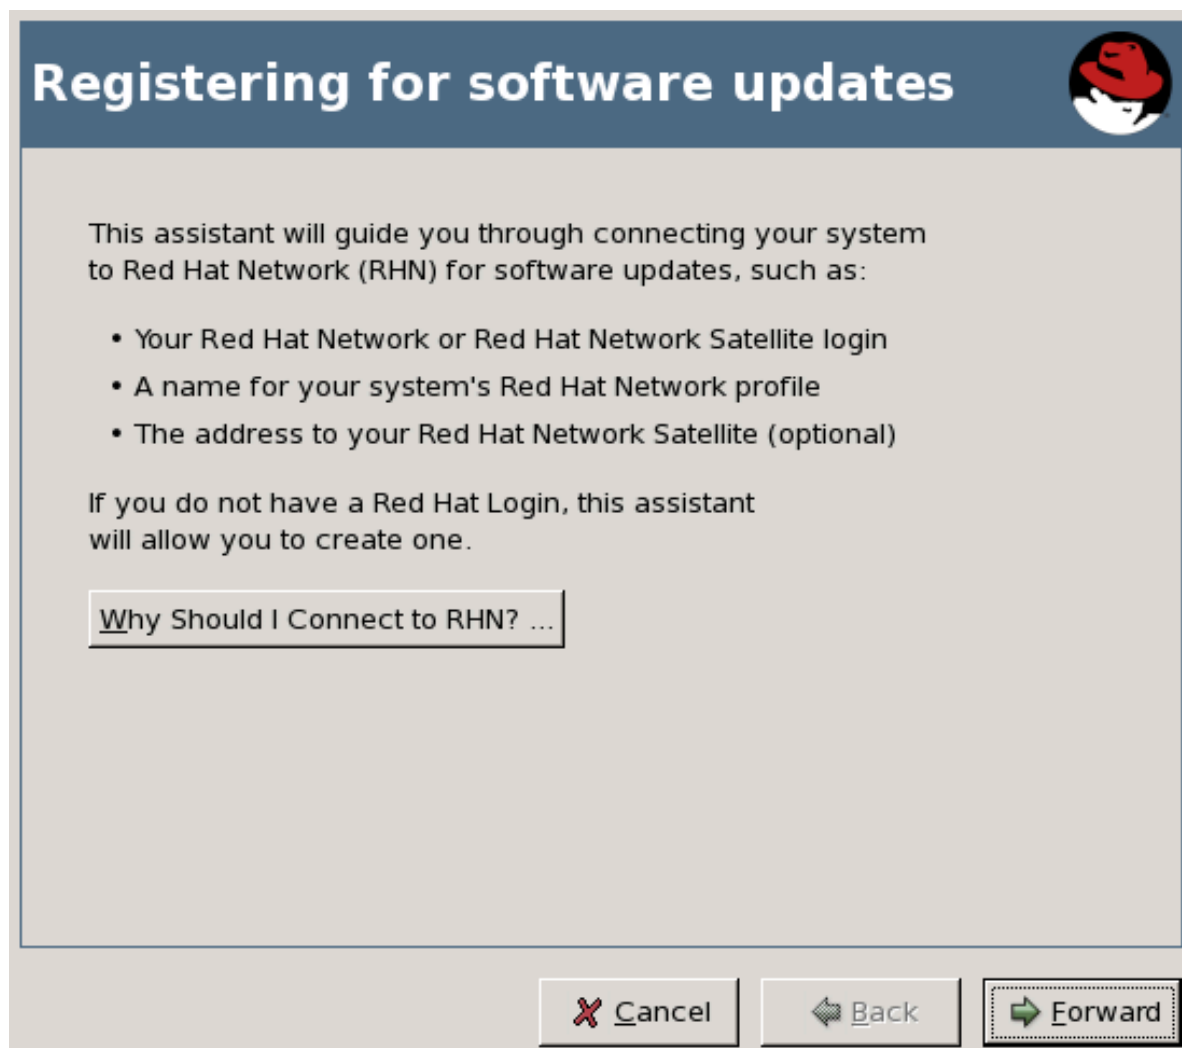


Figure 2.2. Registering for Software Updates

The **Registering for Software Updates** page summarizes the steps involved in the registration process. To learn more about the benefits of Hosted and Satellite, press the **Why Should I Connect to RHN** button. Otherwise, press the Forward button to continue.

Choose an update location

You may connect your system to **Red Hat Network** (<https://rhn.redhat.com/>) or to a **Red Hat Network Satellite** or **Red Hat Network Proxy** in order to receive software updates.

I'd like to receive updates from **Red Hat Network**. (I don't have access to a Red Hat Network Satellite or Proxy.)

I have access to a **Red Hat Network Satellite** or **Red Hat Network Proxy**. I'd like to receive software updates from the Satellite or Proxy below:

Red Hat Network Location:

Example: <https://satellite.example.com>

Figure 2.3. Choose an Update Location

The **Choose an Update Location** page allows you to select the source of your software updates - either from RHN Hosted or from RHN Satellite Server or Proxy Server. For Satellite or Proxy, the associated radio button and enter the URL of your Satellite or Proxy into the **Red Hat Network Location** field.

If you connect to the internet through an HTTP Proxy, press the **Advanced Network Configuration** button. In the subsequent pop-up window, use the appropriate fields for your HTTP proxy; if your proxy requires authentication, enter the username and password here. When finished, press the **Close** button to continue. You are returned to the **Choose an Update Location** page. Press **Forward** to continue.



Enter your account information

Please enter your account information for **Red Hat Network** (<http://rhn.redhat.com/>)

Login:

Password:

i Tip: Forgot your login or password? Look it up at <https://www.redhat.com/wapps/sso/rhn/lostPassword.html>

Create a New Login

Figure 2.4. Enter Your Account Information

The **Enter Your Account Information** page requires you to enter your RHN login information (if you already have one), or to create a new account if you do not. To create a new RHN account, press the Create a New Account button. Fill in the fields that are marked with an asterisk and any other information you wish to enter. Press the Create New Login button to create your new login.



Note

If you are registering to RHN Hosted as part of an organization, please do not create a new account through this screen. Contact your Organization Administrator and request that they create an account for you, then provide that information in the Enter Your Account Information page. Otherwise, you may not be correctly associated with your organization or its resources.



The screenshot shows a web form titled "Create your system profile" with the Red Hat logo in the top right corner. The form is divided into two main sections: "System Name" and "Profile Data".

System Name

You'll want to choose a name for this system so you'll be able to identify it in the Red Hat Network interface.

System Name:

Profile Data

You'll need to send us a profile of what packages and hardware are installed on your system so we can determine what updates are available.

Send hardware profile

Send package profile

At the bottom of the form, there are three buttons: "Cancel" (with a red X icon), "Back" (with a left arrow icon), and "Forward" (with a right arrow icon).

Figure 2.5. Create Your System Profile

The **Create Your System Profile** page allows you to select a profile name for the system you are registering. The default name for any system is that system's hostname, although you may change it as you like. You can also select whether to report hardware and package information to RHN. It is recommended that you choose to report this information; doing so allows RHN to automatically subscribe your system to the base and child channels most appropriate to your system. If you wish, you may press either the **View Hardware Profile** or **View Package Profile** button to view the information that `rhncat` uploads to RHN or Satellite in this step.



Note

This automatic registration does not automatically subscribe your system to optional child channels, such as the RHN Tools channel. If you wish to register a system and have them automatically subscribed to a set of channels of your choice, consider using a kickstart profile or `rhncat_ks` and activation keys to do so.



Figure 2.6. Review System Subscription Details

The **Review System Subscription Details** page displays the base and child channel information to which your system has been subscribed. Take a moment to review the channels, and then press **Forward** to continue.



Figure 2.7. Finish Setting Up Software Updates

The **Finish Setting Up Software Updates** page indicates that you have successfully registered a Red Hat Enterprise Linux 5 system with RHN. From this point, you do not have to do anything to receive software updates. A package icon will appear in the upper right corner of your desktop when updates are available. Click on the icon to apply available updates. Click **Finish** to exit the wizard.



Note

If you do not have any entitlements available for this system, this final page indicates that the registration has failed. This does not mean that the system profile has not been stored with RHN, only that you will not receive automatic updates without manual intervention. You can always login to the RHN or Satellite Web interface and either purchase additional entitlements or get an entitlement from your Satellite administrator. Click the **Exit software update setup** button to exit the wizard.

2.1.1. Command-line version of `rhncregister`

There is also a command-line version of `rhncregister` that allows you to register your system for access to RHN or Satellite without a graphical desktop environment.

Type **rhn_register** at a shell prompt. If you are on shell terminal window and want to run the non-graphical version, you must type **rhn_register --nox** to prevent opening the graphical client.

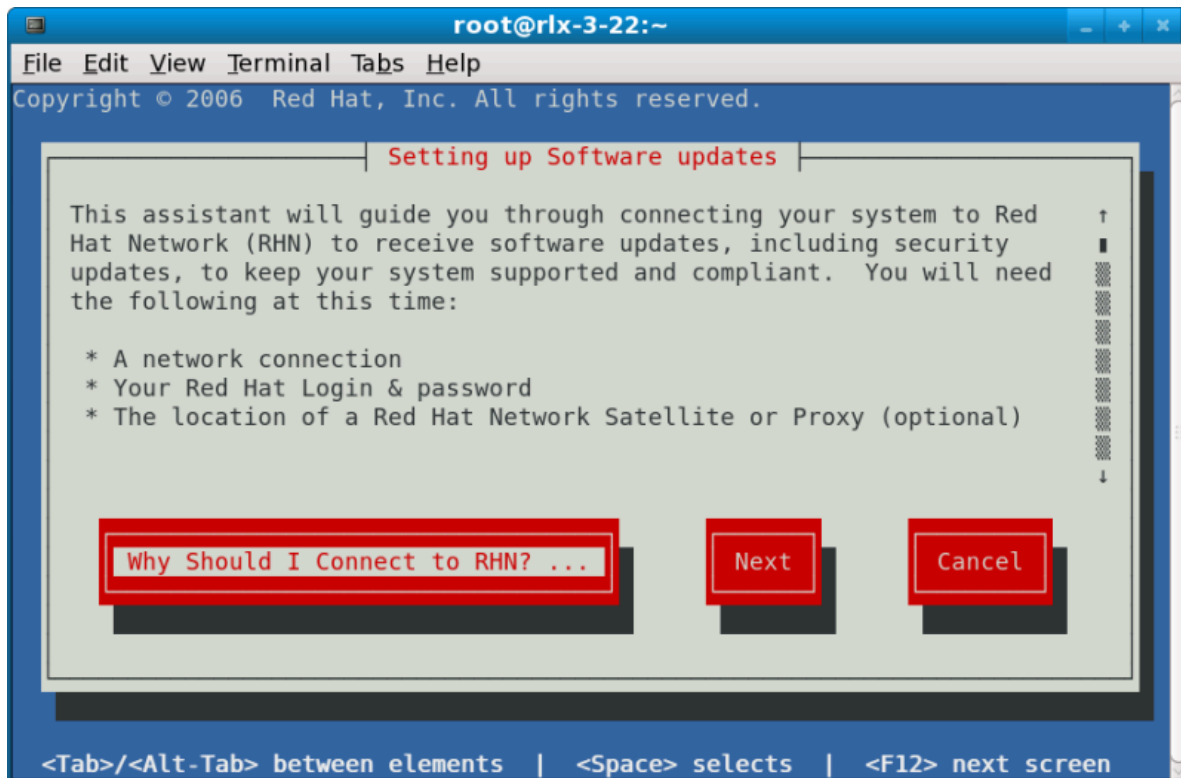


Figure 2.8. **rhn_register** Command-line version

The interface of the command-line version of **rhn_register** has the same configuration screens as the graphical desktop version. However, to navigate the screen, use the directional keys on the keyboard to move left or right and highlight the selections. Press the Space Bar to select an action. Press **Tab** to move through different navigational elements such as text boxes, checkboxes (which are marked with an **x** when selected), and radio buttons (which when selected will be marked with an asterisk).

Package Updater

Depending on your version of Red Hat Enterprise Linux, systems registered to a Satellite can update client systems directly using various tools and applications installed on the system. For Red Hat Enterprise Linux 5, you can use the **Package Updater** (or **pup**) to keep systems updated.

The **Package Updater** (**pup**) is the desktop update application for Red Hat Enterprise Linux 5. Using this tool, you can update packages and read details on the updated packages, such as bug fix information, security alerts, enhancements, and more.

3.1. Using the Package Updater

To start the **Package Updater** from the desktop, open **Applications** (the main menu on the panel) ▢ **System Tools** ▢ **Package Updater**.

If you are at a shell prompt window, type **pup** to open the **Package Updater**.

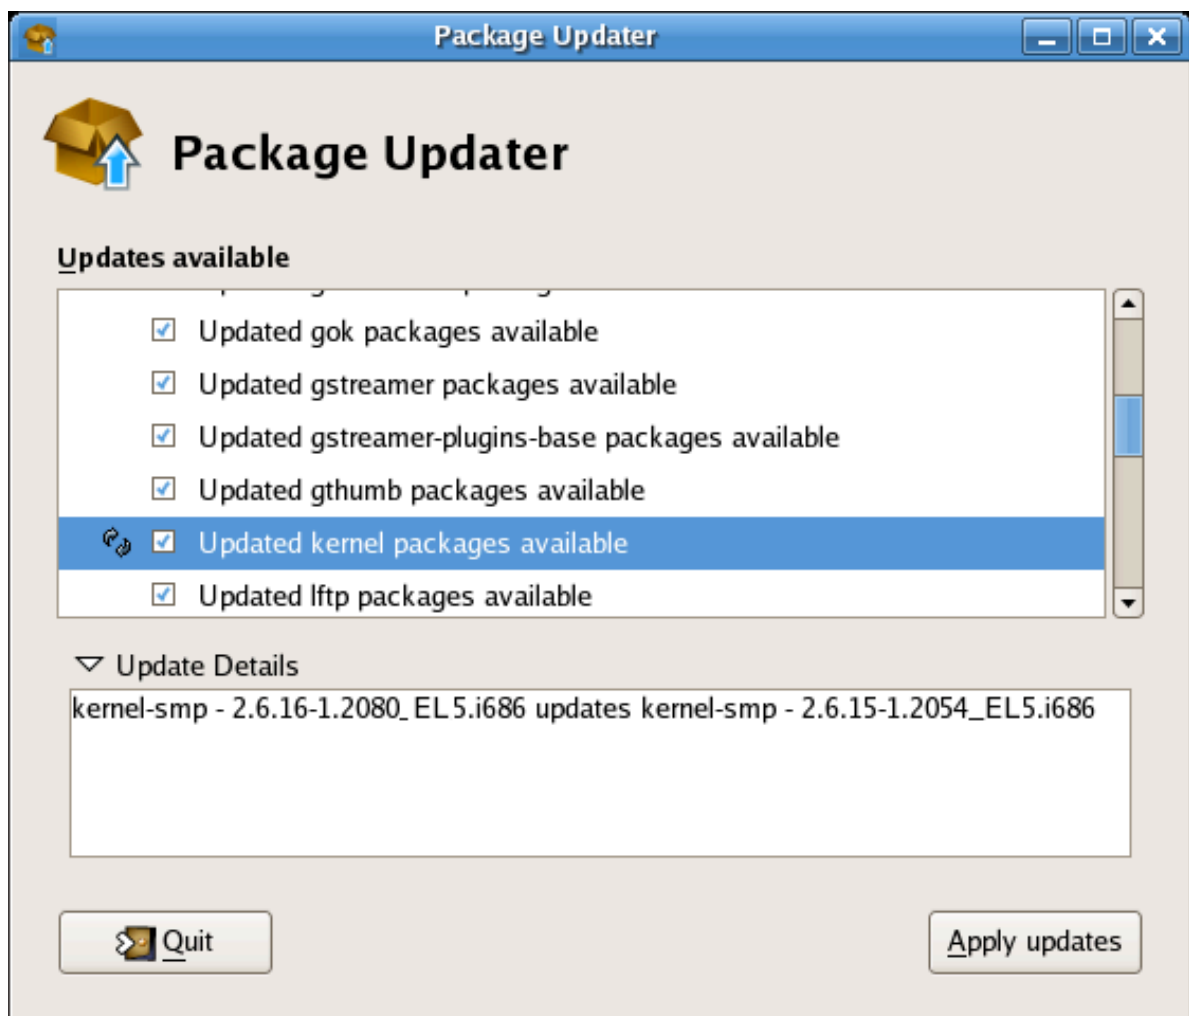


Figure 3.1. Package Updater Interface

If there are multiple package updates, they will be listed with checkmarks next to them so that you can choose which files to update. Some packages (for example, kernel packages) may have a circular arrow icon next to them, indicating that you are required to reboot your system after updating the package.

To view the update details of any package, highlight the package and click the arrow next to **Update Details**.

When you are ready to update the packages, click **Apply updates**. The Updater will resolve any dependencies, and notify you when a package must be installed to meet a dependency for an updated package.

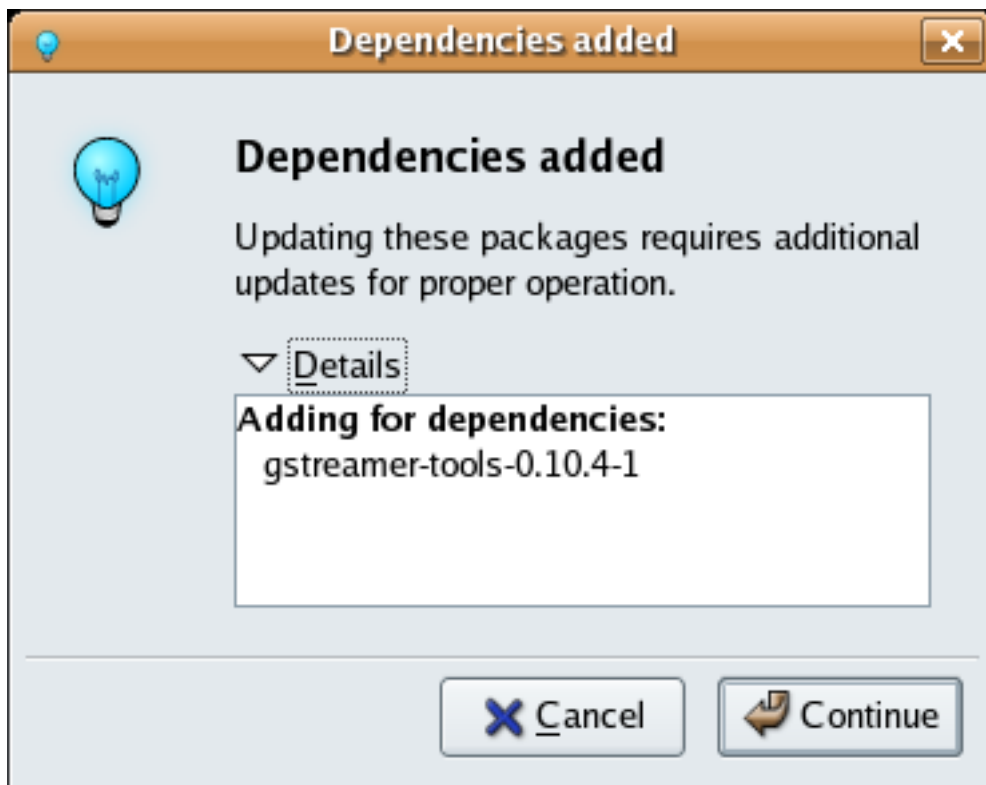


Figure 3.2. Package Dependency

Click **Continue** to accept the dependency and resume the update.

If this is the first time using the **Package Updater**, the program will prompt you to import the Red Hat GPG security key that verifies that a package has been signed and is certified for Red Hat Enterprise Linux.



Figure 3.3. Import the GPG Key

Click **Import Key** to accept the Key and continue with the update.

When the update completes, you may be prompted to reboot your system for the changes to take effect.



Figure 3.4. Reboot Prompt

You can choose to reboot now or later, but it is recommended to click **Reboot Now** to start using the updated packages.

3.2. The **Package Updater** Applet

Red Hat Enterprise Linux 5 also features a a running program on the graphical desktop panel that periodically checks for updates from the RHN or Satellite server and will alert users when a new update is available.




Figure 3.5. **Package Updater** Applet

The **Package Updater** Applet stays in the notification tray of the desktop panel and checks for new updates periodically. The applet also allows you to perform a few package maintenance tasks from the applet by clicking the notification icon and choosing from the following actions:

- **Refresh** — Check RHN or the Satellite for new updates
- **View Updates** — launches the **Package Updater** application so that you can see any available updates in more detail and configure the updates to your specifications

- **Apply Updates** — Download and Install all updated packages.
- **Quit** — close the applet

3.3. Updating Packages from the Command Line with yum

The foundation of the **Package Updater** is the Yum package manager, developed by Duke University to improve the installation of RPMs. **yum** searches supported repositories for packages and their dependencies so they may be installed together in an effort to alleviate dependency issues. Red Hat Enterprise Linux 5 uses **yum** to fetch packages and install packages.

up2date is not available on Red Hat Enterprise Linux 5, which uses Yum (Yellowdog Updater Modified). The entire stack of tools that installs and updates software in Red Hat Enterprise Linux 5 is now based on Yum. This includes everything from the initial installation via **Anaconda** installation program to host software management tools like **pirut**.

3.3.1. yum Commands

yum commands are typically typed as the following:

```
yum command [package_name]
```

By default, Yum will automatically attempt to check all configured repositories to resolve all package dependencies during an installation or upgrade. The following is a list of the most commonly-used **yum** commands. For a complete list of available yum commands, refer to **man yum**.

yum install package_name

Used to install the latest version of a package or group of packages. If no package matches the specified package name(s), they are assumed to be a shell wildcard, and any matches are then installed.

yum update package_name

Used to update the specified packages to the latest available version. If no packages are specified, then **yum** will attempt to update all installed packages.

If the **--obsoletes** option is used (i.e. **yum --obsoletes package_name**), yum will process obsolete packages. As such, packages that are obsoleted across updates will be removed and replaced accordingly.

yum check-update

This command allows you to determine whether any updates are available for your installed packages. **yum** returns a list of all package updates from all repositories if any are available.

yum remove package_name

Used to remove specified packages, along with any other packages dependent on the packages being removed.

yum provides package_name

Used to determine which packages provide a specific file or feature.

yum search keyword

This command is used to find any packages containing the specified keyword in the description, summary, packager and package name fields of RPMs in all supported repositories.

yum localinstall *absolute path to filename*

Used when using yum to install a package located locally in the machine.

Red Hat Update Agent

The **Red Hat Update Agent** is your connection to Red Hat Network on Red Hat Enterprise Linux 4. It enables you to register your systems, create System Profiles, and alter the settings by which your organization and RHN interact. Once registered, your systems can use the **Red Hat Update Agent** to retrieve the latest software packages from Red Hat. This tool allows you to always have the most up-to-date Red Hat Enterprise Linux systems with all security updates, bug fixes, and software package enhancements.

Remember, this tool must be run on the system you wish to update. You cannot use the **Red Hat Update Agent** on the system if it is not entitled to an RHN service offering.



Warning

Only systems running Red Hat Enterprise Linux 3 and later can use the **Red Hat Update Agent** to register with RHN. Systems running Red Hat Enterprise Linux 2.1 must use **Red Hat Network Registration Client** before starting the **Red Hat Update Agent**. Refer to *Chapter 2, The `rhn_register` Client* for instructions, then return to this chapter for **Red Hat Update Agent** instructions.



Important

You must use **Red Hat Update Agent** Version 2.5.4 or higher to upgrade your kernel automatically. It installs the updated kernel and configures LILO or GRUB to boot the new kernel the next time the system is rebooted. To ensure that you are running the latest version, execute the command `up2date up2date`. If you do not have the latest version installed, this command updates it.

4.1. Starting the Red Hat Update Agent

If you are not running the X Window System or prefer the command line version of the **Red Hat Update Agent**, skip to [Section 4.3, “Command Line Version”](#).

You must be root to run the **Red Hat Update Agent**. If started as a standard user, Red Hat Update Agent prompts you to enter the root password before proceeding. The **Red Hat Update Agent** can be started using one of the following methods:

For Red Hat Enterprise Linux 5:

- On the GNOME and KDE desktops, **Applications** (the main menu on the panel) => **Add/Remove Software**.
- At a shell prompt (for example, an `xterm` or `gnome-terminal`), type the command `system-config-packages`.

For Red Hat Enterprise Linux 4:

- On the GNOME and KDE desktops, go to Applications (the main menu on the panel) => **System Tools** => **Red Hat Network**.

- At a shell prompt (for example, an **xterm** or **gnome-terminal**), type the command **up2date**.

If you choose the last option and start the application from a shell prompt, you can specify the options in [Table 4.1, “Graphical Update Agent Options”](#). To view these options, type the command **up2date --help**.

For example, use the following command to specify the directory in which to download the updated packages (temporarily overriding your saved configuration):

```
up2date --tmpdir=/tmp/up2date/
```

Option	Description
--configure	Configure Red Hat Update Agent options. Refer to Section 4.4, “Configuration” for detailed instructions.
-d, --download	Download packages only; do not install them. This argument temporarily overrides the configuration option Do not install packages after retrieval . Use this option if you prefer to install the packages manually.
-f, --force	Force package installation. This option temporarily overrides the file, package, and configuration skip lists.
-i, --install	Install packages after they are downloaded. This argument temporarily overrides the configuration option Do not install packages after retrieval .
-k, --packagedir	Specify a colon separated path of directories in which to look for packages before trying to download them.
--nosig	Do not use GPG to check package signatures. This option temporarily overrides the saved configuration option.
--tmpdir=directory	Temporarily override the configured package directory. The default location is /var/spool/up2date . This option is useful if you do not have enough space in the configured location.
--dbpath=dir	Specify an alternate RPM database to use temporarily.

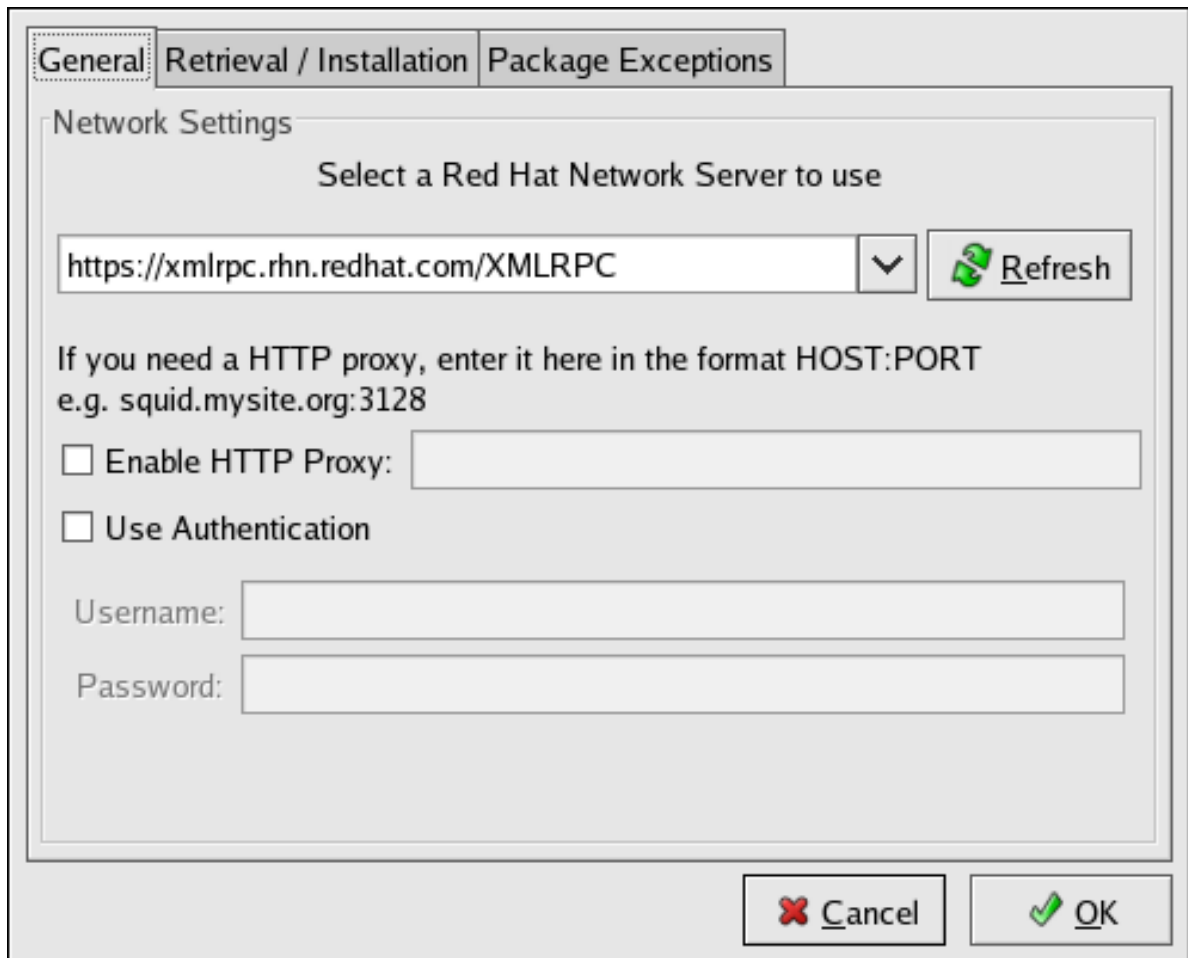
Table 4.1. Graphical Update Agent Options

The first time you run the **Red Hat Update Agent**, two dialog boxes appear that you will not see in subsequent startups: **Configure Proxy Server** and **Install GPG Key**.

As shown in [Figure 4.1, “Configure Proxy Server”](#), the first dialog box to appear prompts you for HTTP Proxy Server information. This is useful if your network connection requires you to use a proxy server to make HTTP connections. To use this feature, select the **Enable HTTP Proxy** checkbox and type your proxy server in the text field with the format **HOST:PORT**, such as **squid.mysite.org:3128**. Additionally, if your proxy server requires a username and password, select the **Use Authentication** checkbox and enter your username and password in the respective text fields.

An HTTP Proxy Server is not required by Red Hat Network. If you do not want to use this feature, click the **OK** button without making any selections. Note that the Red Hat Network Server dropdown menu at the top of the dialog box is only useful to RHN Proxy and Satellite customers. These customers

should refer to the *RHN Client Configuration Guide* for registration steps. Also note that this dialog box is actually the **General** tab of the **Red Hat Update Agent Configuration Tool**. Refer to [Section 4.4, “Configuration”](#) for detailed instructions.



The screenshot shows a dialog box with three tabs: "General", "Retrieval / Installation", and "Package Exceptions". The "General" tab is selected. Inside the dialog, there is a section titled "Network Settings" with the instruction "Select a Red Hat Network Server to use". Below this is a text box containing the URL "https://xmlrpc.rhn.redhat.com/XMLRPC" and a dropdown arrow. To the right of the text box is a "Refresh" button with a green circular arrow icon. Below the text box, there is a note: "If you need a HTTP proxy, enter it here in the format HOST:PORT e.g. squid.mysite.org:3128". There are two checkboxes: "Enable HTTP Proxy:" followed by an empty text box, and "Use Authentication". Below these are two more text boxes labeled "Username:" and "Password:". At the bottom right of the dialog are two buttons: "Cancel" with a red 'X' icon and "OK" with a green checkmark icon.

Figure 4.1. Configure Proxy Server

The second dialog box to appear prompts you to install the Red Hat GPG key, as shown in [Figure 4.2, “Install GPG Key”](#). This key is used to verify the packages you download for security purposes. Click **Yes** to install the key, and you will not see this message again.



The screenshot shows a dialog box with a question mark icon in a blue speech bubble. The text inside reads: "Your GPG keyring does not contain the Red Hat, Inc. public key. Without it, you will be unable to verify that packages Update Agent downloads are securely signed by Red Hat." Below this, it says: "Your Update Agent options specify that you want to use GPG." and "Install key?". At the bottom right are two buttons: "No" with a red 'X' icon and "Yes" with a green checkmark icon.

Figure 4.2. Install GPG Key

4.2. Registration

Before you begin using Red Hat Network, you must create a username, password, and System Profile. Upon launch, the Red Hat Update Agent senses whether these tasks have been accomplished. If not, it guides you through the registration process.

If you ever need to force the Red Hat Update Agent into registration mode, such as to re-register an existing system, you may do so by issuing the following command at a shell prompt:

```
up2date --register
```



Important

If your username is part of a larger organizational account, you should take caution when registering systems. By default, all systems registered with the **Red Hat Update Agent** end up in the Ungrouped section of systems visible only to Satellite Administrators. To ensure you retain management of these systems, Red Hat recommends that your organization create an activation key associated with a specific system group and grant you permissions to that group. You may then register your systems using that activation key and find those System Profiles within RHN immediately. Refer to [Section 4.5, “Registering with Activation Keys”](#) for instructions.


After installing the Red Hat GPG Key, the screen shown in [Figure 4.3, “Welcome Screen”](#) appears. It appears each time you start the Red Hat Update Agent. Click **Forward** to continue.



Figure 4.3. Welcome Screen

4.2.1. Registering a User Account

Before you create a System Profile, you must create a user account. Red Hat recommends that you do so through the website at <https://rhn.redhat.com/newlogin/>, but you may also do so via Red Hat Update Agent (**up2date**).



Important

Users may access and read Red Hat's privacy statement from this screen. Click the **Read our Privacy Statement** button to do so. Red Hat is committed to protecting your privacy. The information gathered during the registration process is used to create a System Profile, which is essential to receiving update notifications about your system. When finished, click **OK**

Those users that have created a Red Hat login previously may enter their username and password and click the **Forward** button to continue.

Users that have registered at least one system with Red Hat Network can add new machines to the same account. To do so, run the Red Hat Update Agent on the new machine and enter the existing Red Hat username and password at this screen.

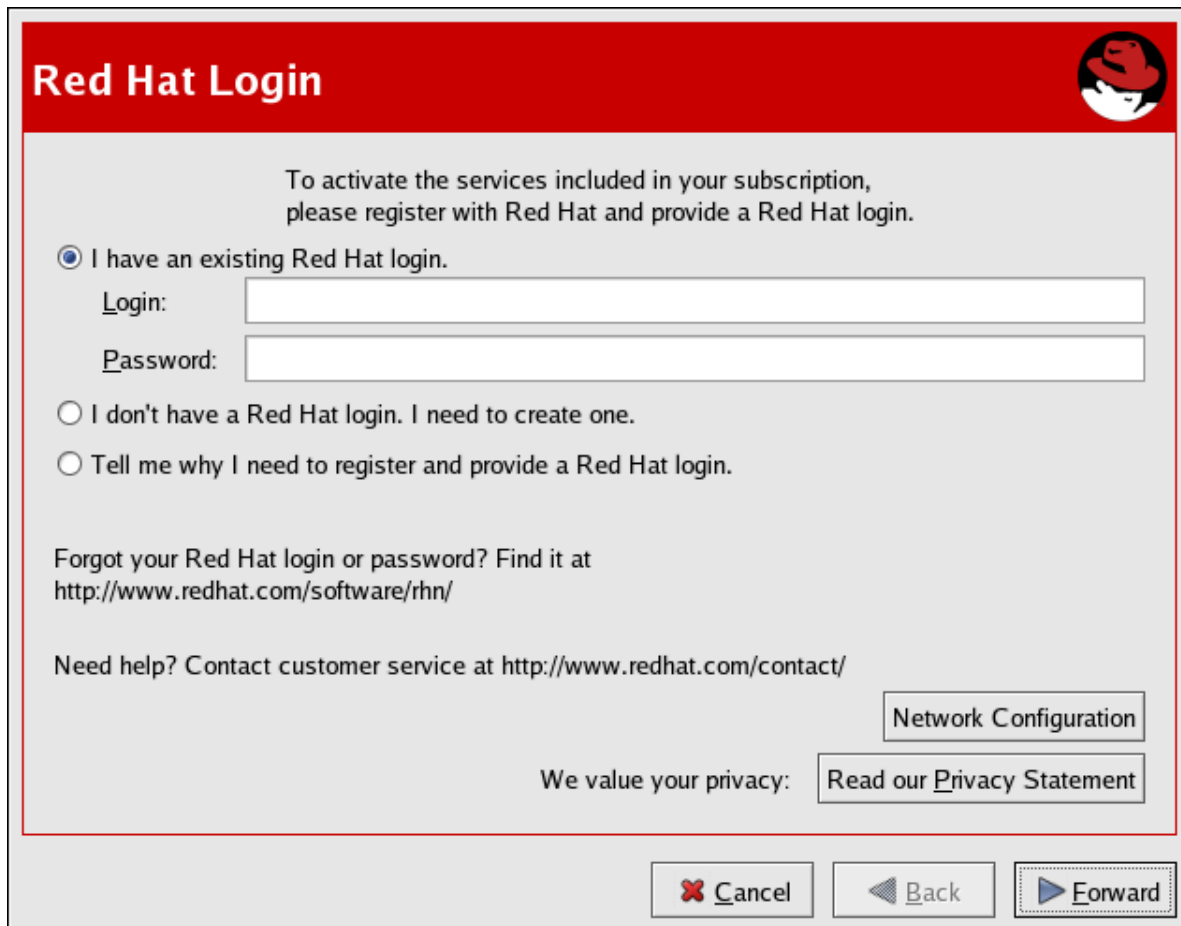


Figure 4.4. Red Hat Login Screen

New users must select the **I don't have a Red Hat login. I need to create one.** radio button and click the **Forward** button. Add details about yourself and your business to the screen shown in [Figure 4.5, "Create a User Account"](#), and identify the methods by which you may be reached.

Your username has the following restrictions:

- Cannot contain any spaces
- Cannot contain the characters & +, %, or '
- Is not case-sensitive, thereby eliminating the possibility of duplicate usernames differing only by capitalization

In addition, the following restrictions apply to both your username and password:

- Must be at least four characters long
- Cannot contain any tabs
- Cannot contain any line feeds

Passwords are case-sensitive for obvious reasons.



Note

You must choose a unique username. If you enter one already in use, you will see an error message. Try different usernames until you find one that has not been used.

Complete all fields marked by an asterisk (*). The address and email addresses are required so that Red Hat may communicate with you regarding your account. You may select to receive monthly copies of Red Hat Magazine, a valuable source of tips, insights, and Red Hat news.

When finished, click **Forward**.

Figure 4.5. Create a User Account

4.2.2. Activate

The Activation screen allows you to select various details of your registration. If you have a subscription number, enter it in the appropriate field. If not, select the **Use one of my existing, active subscriptions** radio button.

In the **Connect Your System** option group, select whether to send a hardware or software profile.

After creating a username and password for your Red Hat Network account, the **Red Hat Update Agent** probes your system for the following information:

- Red Hat Enterprise Linux version
- Hostname
- IP address

- CPU model
- CPU speed
- Amount of RAM
- PCI devices
- Disk sizes
- Mount points

The software System Profile consists of a list of RPM packages for which you wish to receive notifications. The **Red Hat Update Agent** displays a list of all RPM packages listed in the RPM database on your system and then allows you to customize the list by deselecting packages.

To see the details of the information gathered from your system, click the **Details** button next to the profile. When finished, click **OK**. If you uncheck the box to the left of the profile, that information is not sent to RHN.



Note

If you do not send a Software Profile, this system will receive no Errata Updates.

Click **Forward** to send the information to RHN.

Figure 4.6. Activate

Figure 4.7, “Sending System Profile to Red Hat Network” shows the progress bar displayed as the System Profile is sent.

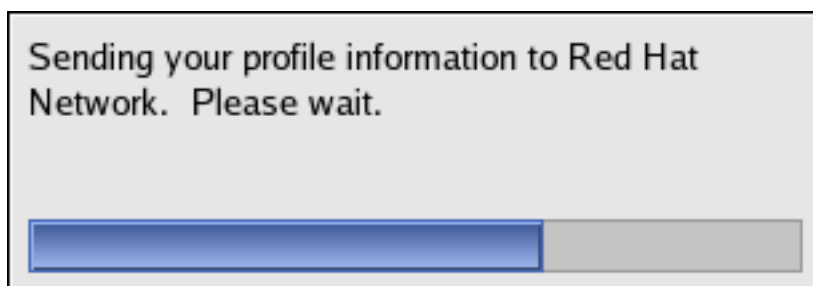


Figure 4.7. Sending System Profile to Red Hat Network

4.2.3. Channels

Red Hat Update Agent next displays all package channels to which you have access. The channels you select from this screen must match the base operating system of the system you are registering. If any child channels are available, such as the **RHEL AS (v.4 for x86) Extras** channel in the figure, you may select them as well. Additional information regarding the selected channel is displayed in the **Channel Information** pane. When finished, click **Forward** to continue.

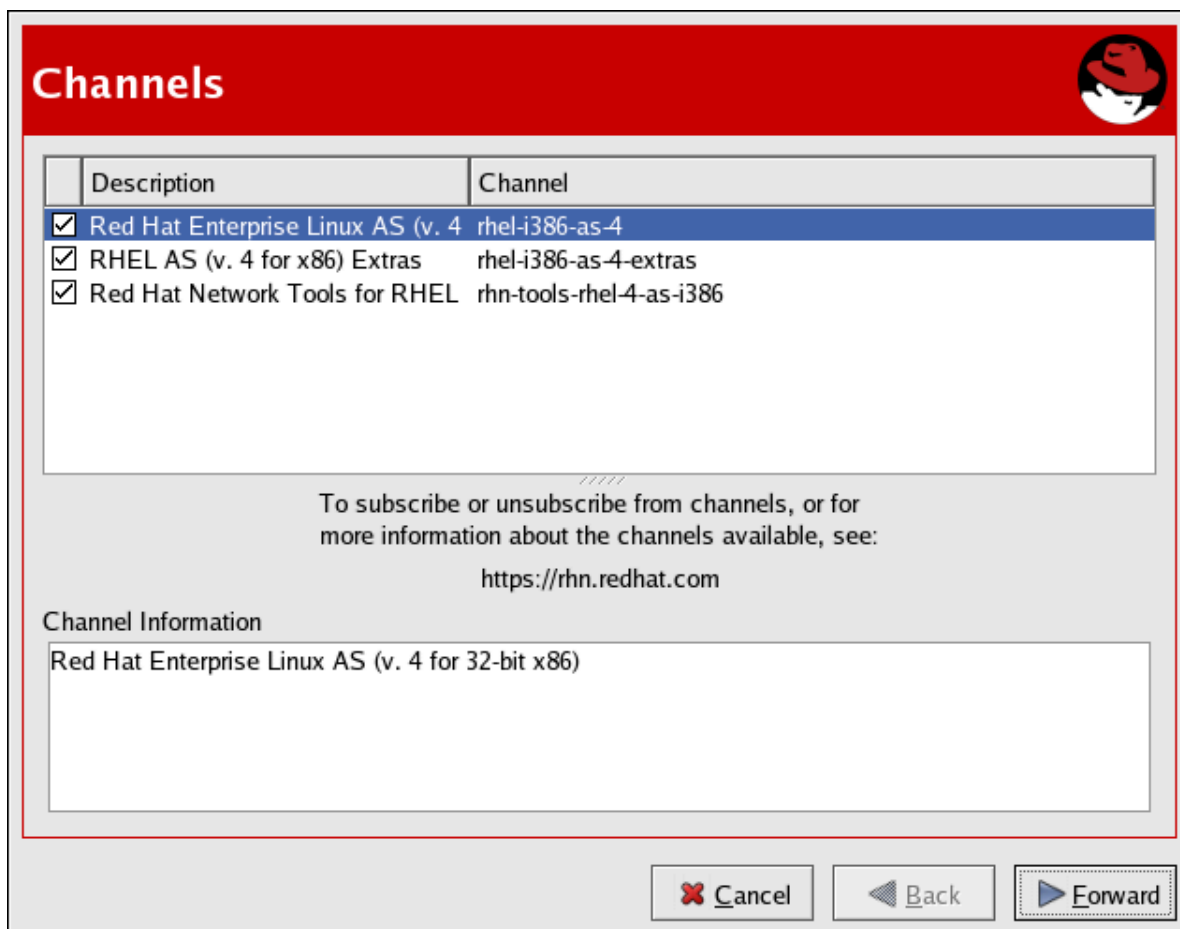


Figure 4.8. Channels

Red Hat Update Agent now compares the packages in your RPM database with those available from the Channel you selected. The progress bar shown in *Figure 4.9, "Fetching package list"* is displayed during this process.

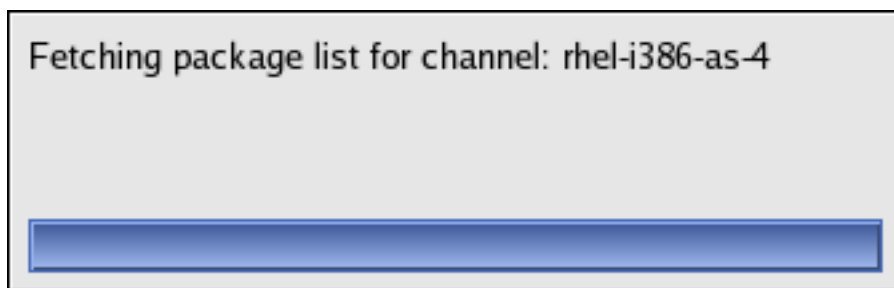


Figure 4.9. Fetching package list

Note

If the version of **up2date** on your system is older than the one in your selected channel, the Red Hat Update Agent asks whether you would like to update it. If you agree, the only package that will be updated is the **up2date** package. This is equivalent to executing the **up2date up2date** command from a shell prompt. Once the updated process has completed, the Red Hat Update Agent restarts and completes the initial update of the system.

4.2.4. Packages Flagged to be Skipped

The next step in the initial update is the selection of files to be skipped. Any packages checked here will not be downloaded and updated by the Red Hat Update Agent. This screen is displayed whenever packages are available that are currently selected to be ignored. You may change these settings at any time from the Red Hat Network Alert Notification Tool. Refer to [Chapter 6, Red Hat Network Alert Notification Tool](#) for additional information.

Make your selections and click **Forward** to continue.

Packages Flagged to be Skipped

Select all packages

	Package Name	Version	Old Version	Arch	Size	Reason Skipped
<input checked="" type="checkbox"/>	kernel	2.6.9-11.EL	2.6.9-5.EL	i686	9854 kB	Pkg name/pattern
<input checked="" type="checkbox"/>	kernel-utils	2.4-13.1.66	2.4-13.1.48	i386	543 kB	Pkg name/pattern

Package Information View Advisory

According to your preferences you have chosen not to automatically update the above packages. If you would like to override your settings and include one of the above packages in the list of packages to retrieve, select its checkbox.

Figure 4.10. Packages Flagged to be Skipped

4.2.5. Available Package Updates

The Red Hat Update Agent next displays all available updates except those you chose to skip in the previous screen. Select those you wish to download and click **Forward** to continue. To view the complete Errata Advisory text for an update, highlight the relevant package and click the **View Advisory** button. When finished, click **OK**.

Select those you wish to download and click **Forward** to continue.

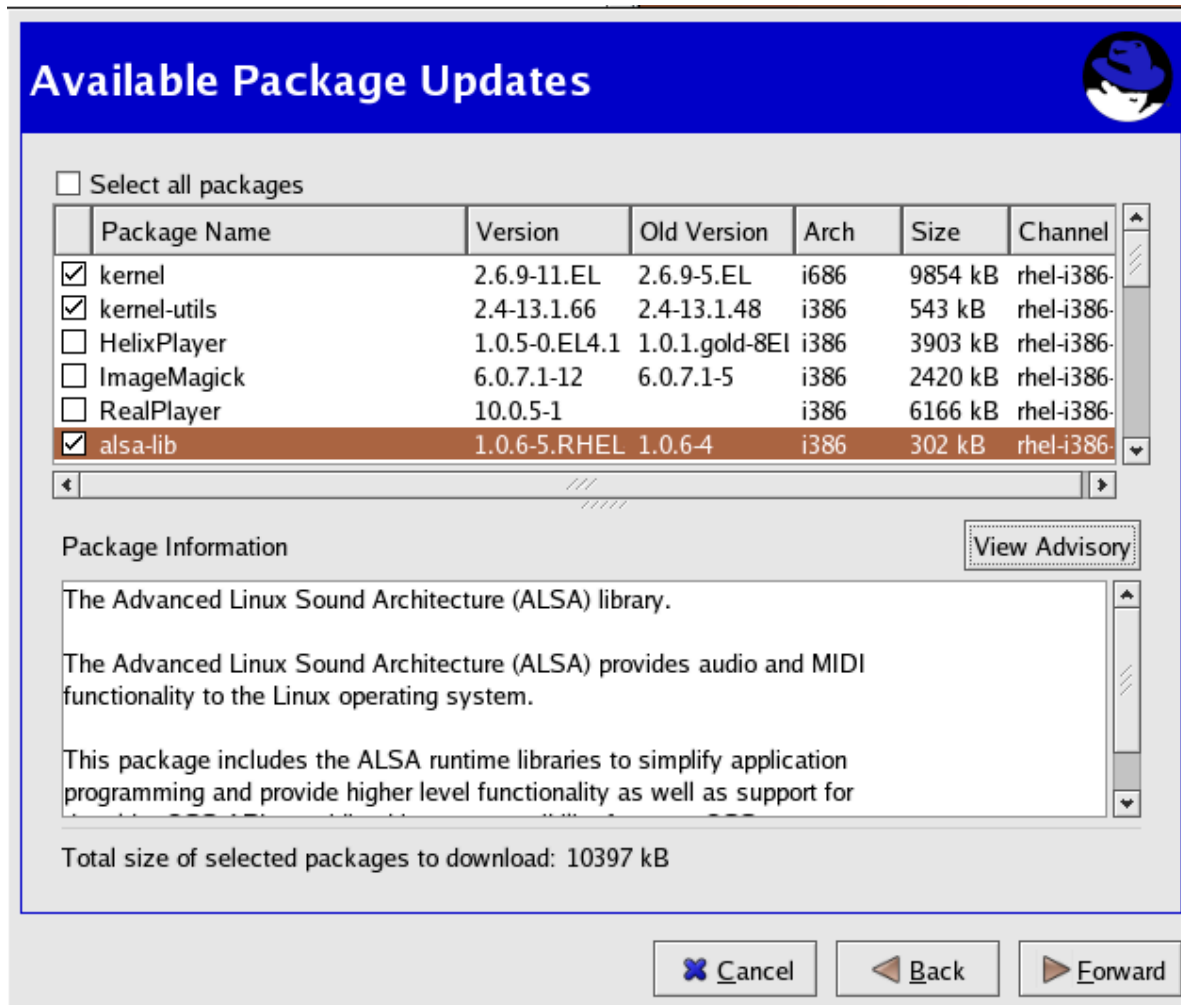


Figure 4.11. Available Package Updates

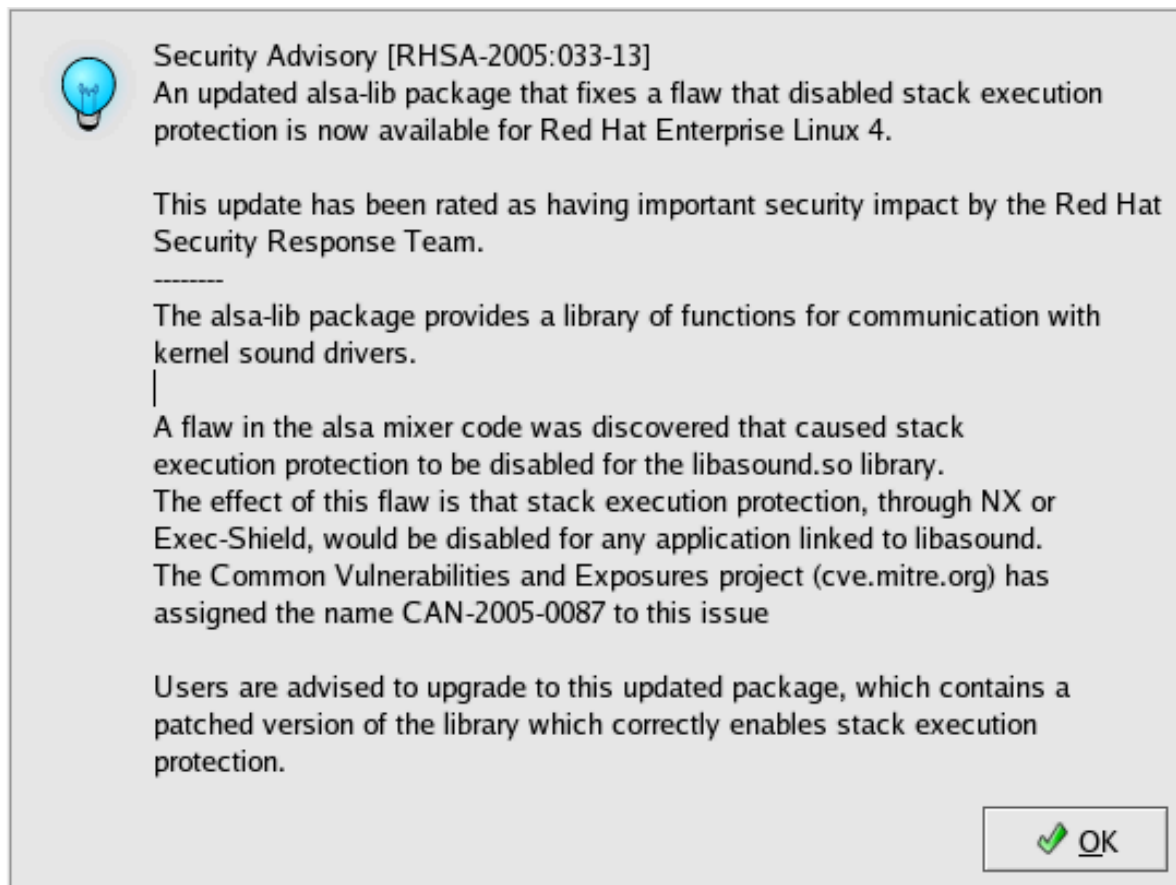


Figure 4.12. Example Errata Advisory

4.2.6. Retrieving Packages

The Red Hat Update Agent tests the packages you selected to be certain that the requirements of each RPM are met. If any additional packages are required, Red Hat Update Agent displays an error message. Click **OK** to continue.

Once all dependencies are met, Red Hat Update Agent retrieves the packages from RHN. As the packages are downloaded, they are temporarily stored in `/var/spool/updates/`.

When all packages have been downloaded, click **Forward** to continue.

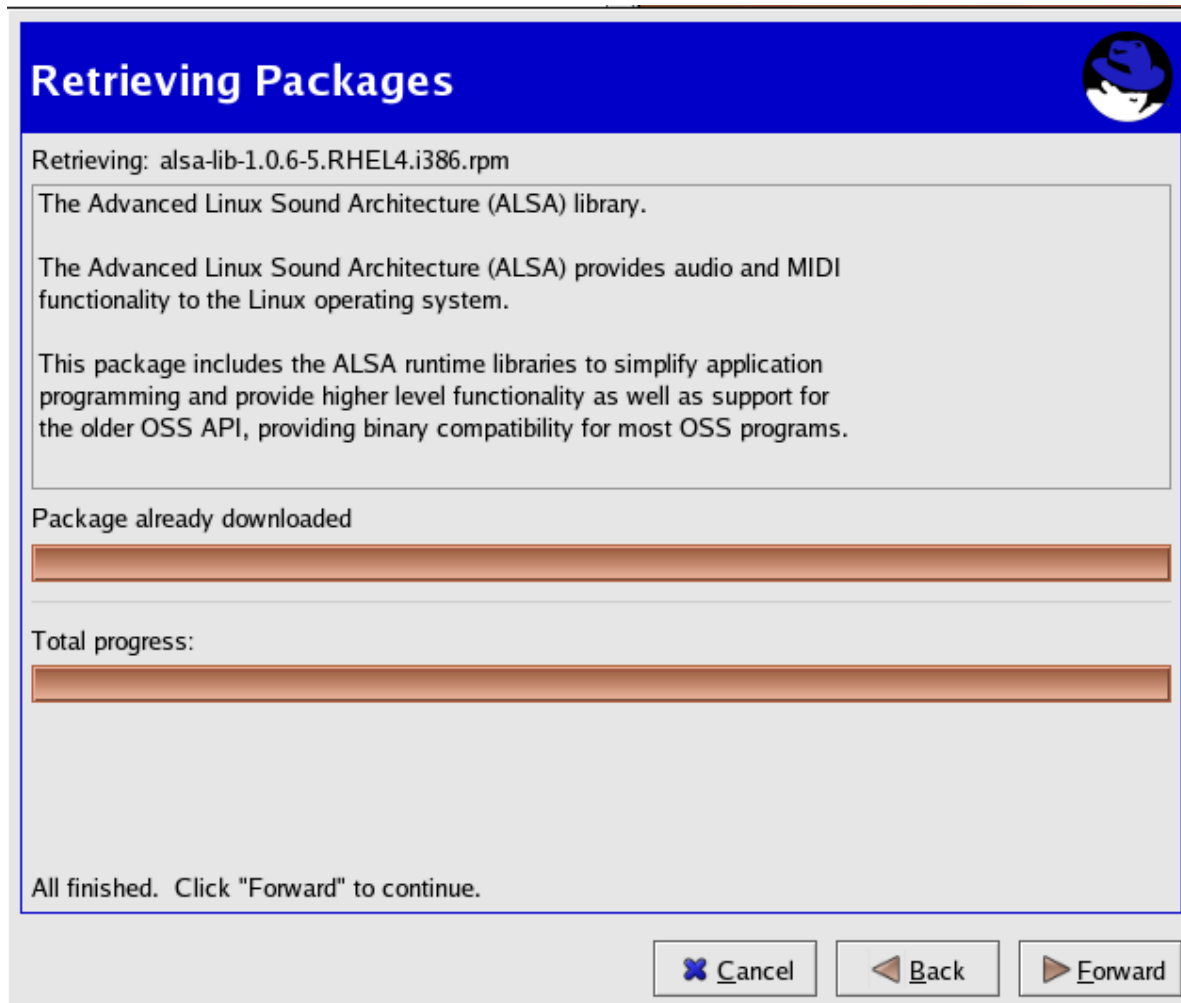


Figure 4.13. Retrieving Packages

4.2.7. Installing Packages

The packages must be installed after downloading them via the **Red Hat Update Agent**. If you chose not to install the packages via the **Red Hat Update Agent**, skip to [Section 4.3.2, “Manual Package Installation”](#) for further instructions. If you configured the Red Hat Update Agent to install the packages (the default setting), the installation process begins. The progress of installing each package, as well as the total progress, is displayed. When the packages have been installed, as seen in [Figure 4.14, “Installing Packages”](#), click **Forward** to continue.

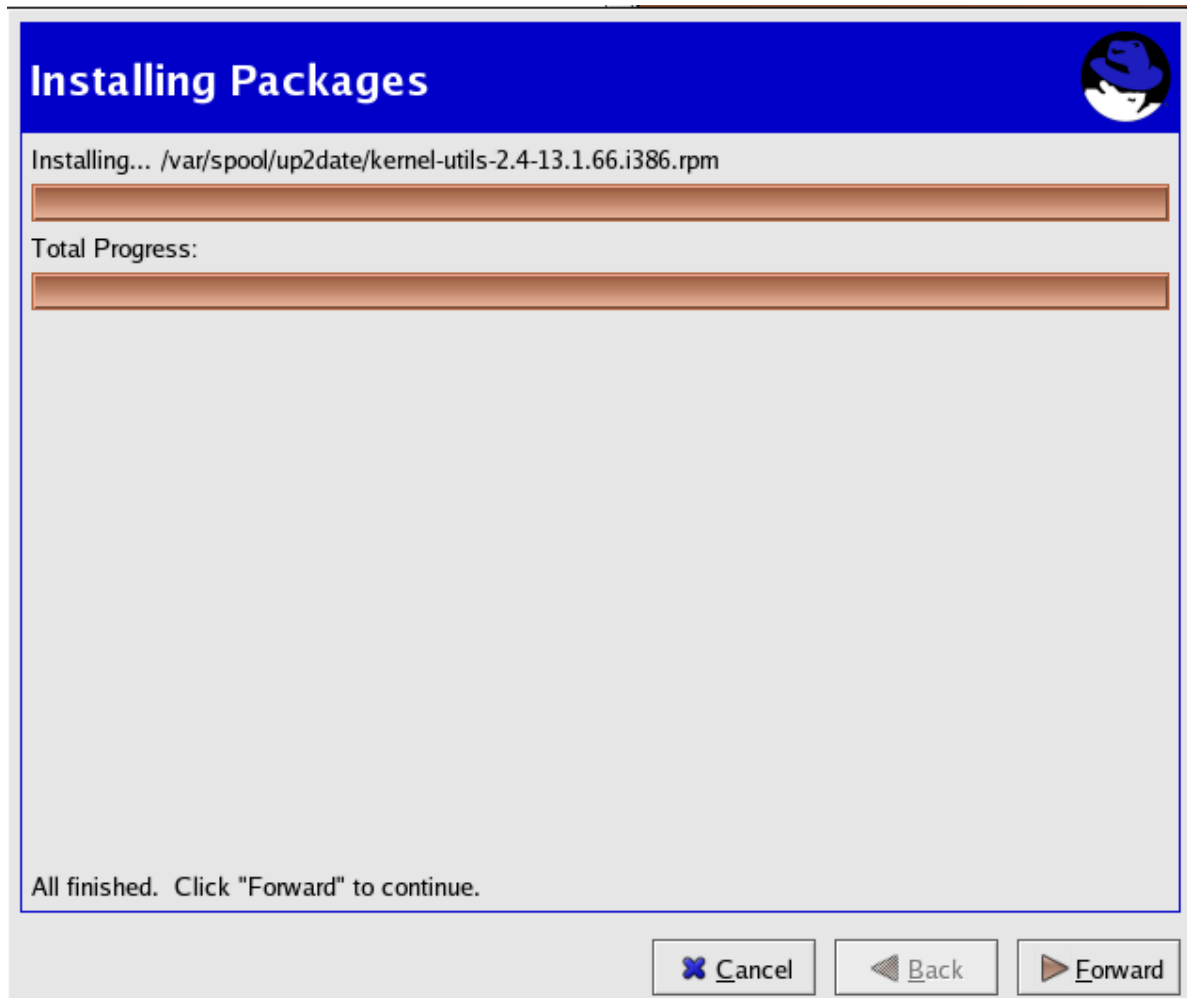


Figure 4.14. Installing Packages

When the **Red Hat Update Agent** has finished downloading the desired packages (and installing them if you chose the install option), it displays the screen in [Figure 4.15, "All Finished"](#). Click **Finish** to exit the **Red Hat Update Agent**.

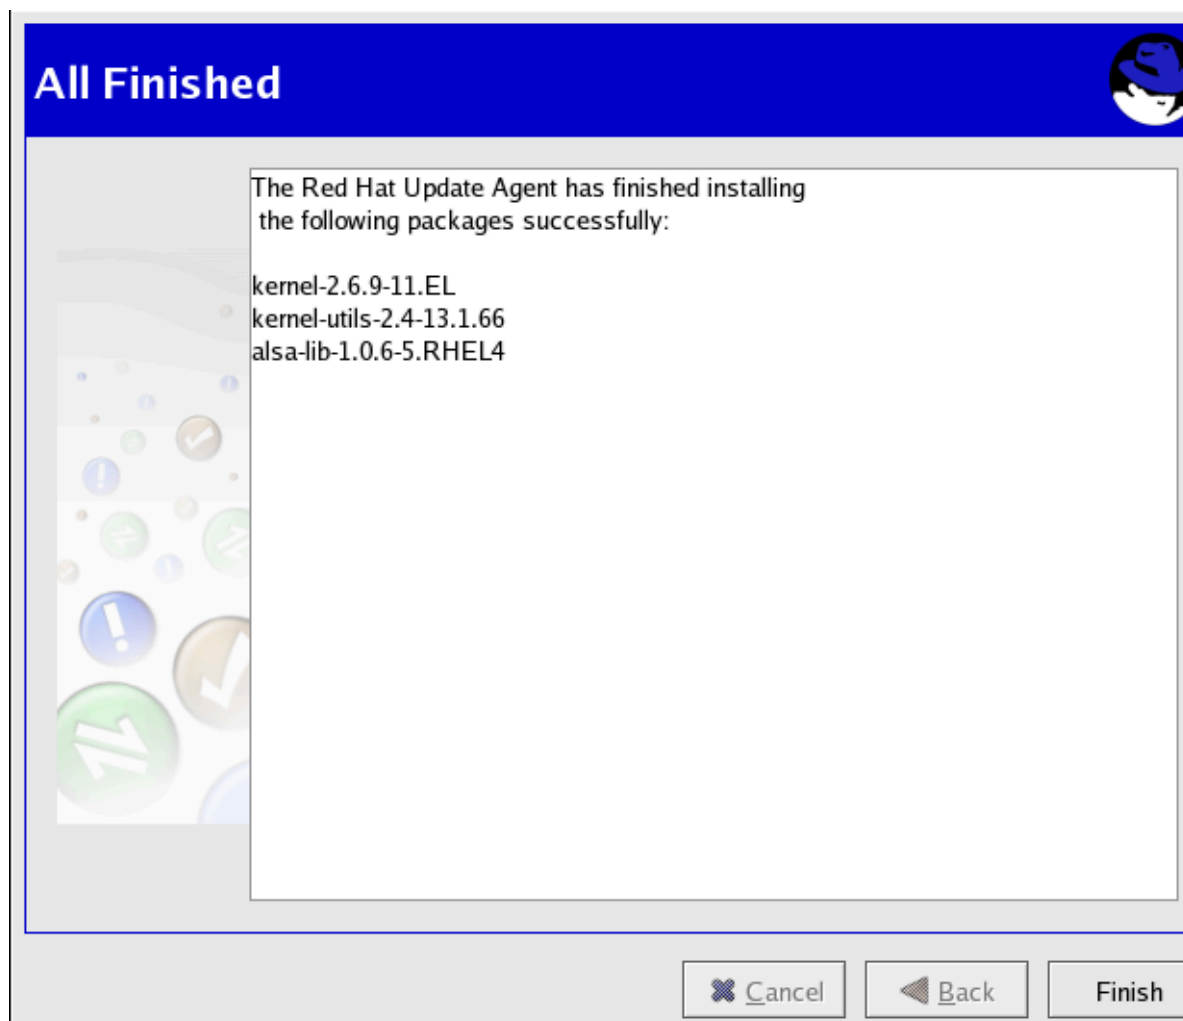


Figure 4.15. All Finished

4.3. Command Line Version

If you are not running X, you can still run the **Red Hat Update Agent** from a virtual console or remote terminal. If you are running X but want to use the command line version, you can force it not to display the graphical interface with the following command:

```
up2date --nox
```

The command line version of the **Red Hat Update Agent** allows you to perform advanced functions or to perform actions with little or no interaction. For example, the following command updates your system with no interaction. It downloads the newer packages and installs them if you configured it to do so.

```
up2date -u
```

The command line version of the **Red Hat Update Agent** accepts the following arguments:

Option	Description
-?, --usage	Briefly describe the available options.
-h, --help	List the available options and exit.
--arch=<i>architecture</i>	Force up2date to install this architecture of the package. Not valid with --update , --list , or --dry-run .
--channel=<i>channel</i>	Specify from which channels to update using channel labels.
--configure	Configure Red Hat Update Agent options. Refer to Section 4.4, "Configuration" for detailed instructions.
-d, --download	Download packages only; do not install them. This argument temporarily overrides the configuration option Do not install packages after retrieval . Use this option if you prefer to install the packages manually.
--dbpath=<i>dir</i>	Specify an alternate RPM database to use temporarily.
--dry-run	Do everything but download and install packages. This is useful in checking dependencies and other requirements prior to actual installation.
-f, --force	Force package installation. This option temporarily overrides the file, package, and configuration skip lists.
--firstboot	Pop up in the center of the screen for Firstboot.
--get	Fetch the package specified without resolving dependencies.
--get-source	Fetch the source package specified without resolving dependencies.
--gpg-flags	Show the flags with which GPG is invoked, such as the keyring.
--hardware	Update this system's hardware profile on RHN.
-i, --install	Install packages after they are downloaded. This argument temporarily overrides the configuration option Do not install packages after retrieval .
--installall=<<i>channel-label</i>>	Install all available packages from a given channel
--justdb	Only add packages to the database and do not install them.
-k, --packagedir	Specify a colon-separated path of directories in which to look for packages before trying to download them.
-l, --list	List packages relevant to the system.
--list-rollback	Show the package rollbacks available.
--nodownload	Do not download packages at all. This is useful in testing.
--nosig	Do not use GPG to check package signatures. This option temporarily overrides the saved configuration option.
--nosrc	Do not download source packages (SRPMs).
--nox	Do not attempt to run in X. This launches the command line version of the Red Hat Update Agent .

Option	Description
-p, --packages	Update packages associated with this System Profile.
--proxy=proxy URL	Specify an HTTP proxy to use.
--proxyPassword=proxy password	Specify a password to use with an authenticated HTTP proxy.
--proxyUser=proxy user ID	Specify a username to use with an authenticated HTTP proxy.
--register	Register (or re-register) this system with RHN. Refer to Section 4.2, “Registration” for detailed instructions.
--serverUrl=server URL	Specify an alternate server from which to retrieve packages.
--showall	List all packages available for download.
--show-available	List all packages available that are not currently installed.
--show-channels	Show the channel name associated with each package.
--show-orphans	List all packages currently installed that are not in channels to which the system is subscribed.
--show-package-dialog	Show the package installation dialog in GUI mode.
--solvedeps=dependencies	Find, download, and install the packages necessary to resolve dependencies.
--src	Download source packages, as well as binary RPMs.
--tmpdir=directory	Temporarily override the configured package directory. The default location is <code>/var/spool/up2date</code> . This option is useful if you do not have enough space in the configured location.
-u, --update	Update system with all relevant packages.
--undo	Reverse the last package set update.
--upgrade-to-release=release version	Upgrade to the channel specified.
--uuid=uuid	Pass in a Unique User ID generated by the Alert Notification tool.
-v, --verbose	Show additional output while updating.
--version	Show <code>up2date</code> version information.
--whatprovides=dependencies	Show the packages that resolve the comma-separated list of dependencies.

Table 4.2. Update Agent Command Line Arguments



Note

The `--solvedeps` and `--whatprovides` options can be used to solve the dependencies for an RPM regardless even if your system does not currently have access to a channel that contains that package.

4.3.1. Installing the Red Hat GPG key

The first time you run the graphical version of the **Red Hat Update Agent**, it prompts you to install the Red Hat GPG key. This key is required to authenticate the packages downloaded from Red Hat Network. If you run the command line version the first time you start **Red Hat Update Agent**, you must install the Red Hat GPG key manually. If you do not have it installed, you will see the following message:

```
Your GPG keyring does not contain the Red Hat, Inc. public key.
Without it, you will be unable to verify that packages Update Agent downloads
are securely signed by Red Hat.
```



Note

GPG keys must be installed for each user. To install the key to use with Red Hat Network, import the key while logged in as root.

The method for installing the key varies depending on your version of RPM. Starting with version 4.1, which shipped with Red Hat Enterprise Linux 3, you may use RPM to import GPG keys. Issue the following command at a shell prompt as root:

```
rpm --import /usr/share/doc/rpm-4.1/RPM-GPG-KEY
```

For older versions of RPM, such as the one that came with Red Hat Enterprise Linux 2.1, use the **gpg** command (as root):

```
/usr/bin/gpg --import /usr/share/rhn/RPM-GPG-KEY
```

To download the Red Hat GPG key first, you may obtain it from <https://www.redhat.com/security/team/key.html>¹. Here's an example:

```
Type bits/keyID Date User ID
pub 1024D/650D5882 2001-11-21 Red Hat, Inc. (Security Response Team)
sub 2048g/7EAB9AFD 2001-11-21
```

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
```

```
Version: GnuPG v1.2.1 (GNU/Linux)
```

```
mQGIBDv70vQRBADh701rf8WuzDg88kq1V/N5KQ1PF0amn0DB/1EeuAD7n6bCBRmV
ekQWJCdfab0Rf1S+VsFg6IAAAmDIarVnacTLQzqCdGJqTpXm/rGVpLv+mCh+0mT9
QRFBjSzB0uPJ0piIvJwSS00D/wJ8XKzHkVNgW3DiJ9Qz2BHYszU2ISI6FwCgxY6d
IVjWT5jblkLNjtd3+fr024ED/i0e2knetTX3S9LjC+HdGvP8Eds92Ti2CnJLaFJK
Rp749PucnK9mzxPc02jSHgdtjWAXst/st+gWFVbFmkjBQDVSd00B/xEwI1T1+LN8
V7R8BE1Bmg99I1JmDVA2BI/seXvafhzly9bxSHScFnceco/Az9umIs3NXwv3/y0m
ZakDBAC6SAGHBmpVvK0deXJDDb4LcbEhErFU3CpRcJz6A0nFuiV1MGdu1ZXvEUgBA
I6/PDE5nBhfZY3zPjyLPZVtgYioJpZqcRIx/g+bX208kPqvJEuZ19tLCdykfZGpy
bsV7QdSGqBk3snN0izmFj543RaHyEbnwKwbnADhujwMeUAXN+7Q8UmVkiEhhdCwg
SW5jLiAoU2VjdXJpdHkgUmVzcG9uc2UgVGvHbSkghPHN1Y2FsZXJ0QHJlZGhhdc5j
b20+ifcEEExECABcFAj3GczYFCwcKAwQDFQMCAxYCAQIXgAAKCRBeVICDZQ1YghAU
AJ0ceQfUMR2dKyLft/1006qUs+MNLQCggJgd08MU02y11TWID3X0YgyQG+2InAQT
```

¹ <https://www.redhat.com/security/team/key.html>

```
AQIABgUCptyYpQAKCRDurUz9SaVj2e97A/0b2s70hhAMljNwMQS4I2UwVgBgtxdu
D+yBcG/3mwL76MJVY7aX+NN/tT9yDGu+FSiQZZCL/40FOHmvjpcDqfJY+zpTlBii
ZMAPJWts2bB+0QaXxUgWlWw84GVf2rA6RSbvMLTbDjTH8t7J1RGP9zAQu8SgraTA
QbQdao6TNxVt+ohGBBMRAGAGBQI+3LjCAAoJECGRgM3bQqYOf5MAoIjJDe+hD0j
9+jlR0qDs91Ii/C2AJ9SBBfd4A8hyR4z3lY7e0LzjWF51LkCDQq7+903EAgA8tMs
xdUmuTfA+X78fMXh7LCvrL4Hi28CqvNM+Au81XJjDLNawZvpVmF1Mmd9h0Xb5Jt2
BZWLR13rcDUByNdw1EWhVAzCz6Bp9Z3MIDhcP00iIBctIHn7YP9fi5vV0G03iryT
XE01mhWoBlC233wr3XHwsqxFfZzaCZqqNKTl0+PNfEAIzJRgtYiW8nzFTPpIR05E
oRn6EvmQfayOF2uYDX9Sk//10D7T7RLtKjM/hPW/9NoCGwwR0aG+VUzVv4ae1h1L
dJGEjpfTdxcrOUMD8xbkuGMznu0mpDI+J2BUDh5n57y0yEMaGrQ0jfy1ZqqDvZg
osY1ZHa6KlmuCWNTnwADBQf/XYhCicp6iLetnPv61YtyRfFRpnK98w3br+fThywC
t81P2nKv8l1io60sRbksGc1gX8Zl6GoHQYfDe7hYsCHZPowErobECFds5E9M7cmzV
TTYNTvrELrs07jyuPb4Q+mHcsYPIlGR3M+rnXKkjloz+05kOPRJaBEBzP6B8SZKy
QNqEfTkyU4Rbhkzz/UxUxZoRZ+ tqVjNbPKFpRraiQRUDsZFbgksBCzkzd0YURvi
Ceg02K7JPKbZJo6eJA10qiBQvAx2EUijZfxIKqZeLx40EKMaL7Wa2CM/xmkQmCgg
Hyu5bmLSMZ7cXFSwyXost78dehCKv9WypXHV3m4iANWFL4hGBBgRAGAGBQI7+903
AAoJEF5UgINlDViCKWcAoMceYstWVKXJTYtzHEL6Wl8rXr8WAKChuapJIA4/eFsf
4ciWtjY8c00v8Q==
=yOVZ
-----END PGP PUBLIC KEY BLOCK-----
```

Save the text file and import it into your keyring using the method applicable to your version of RPM.

4.3.2. Manual Package Installation

If you chose to download, but not install, the software updates with the **Red Hat Update Agent** or from the RHN website, you must install them manually using RPM.

To install them, change to the directory that contains the downloaded packages. The default directory is `/var/spool/up2date`. Type the command `rpm -Uvh *.rpm`. When the packages finish installing, you can delete them if you wish. You do not need them anymore.

After installing the packages, you must update your System Profile so that you are not prompted to download them again. Refer to [Section 4.3.3, “Synchronizing Your System Profile”](#) for details.

4.3.3. Synchronizing Your System Profile

If you configured the **Red Hat Update Agent** to install the latest packages, the System Profile stored by Red Hat Network is updated after the packages are installed. However, if you only download the latest RPM packages using the **Red Hat Update Agent**, download the RPM packages from the website, or upgrade/install/remove RPM packages yourself, your System Profile is not updated automatically. You must send your updated System Profile to the RHN Servers.

To synchronize the RPM package list on your local Red Hat Enterprise Linux 5.3 system and on Red Hat Network, run the command:

```
rhn-profile-sync
```

After running this command, your RHN System Profile reflects the latest software versions installed on your system.

For Red Hat Enterprise Linux 4 systems, use the following command to update the package list, run the command:

```
up2date -p
```

4.3.4. Log File

The **Red Hat Update Agent** keeps a log of all the actions that it performs on your system in the file `/var/log/up2date`. It uses the standard rotating log method. Thus, older logs are in `/var/log/up2date.1`, `/var/log/up2date.2`, and `/var/log/up2date.3`. The log files store actions performed by the **Red Hat Update Agent** such as when your RPM database is opened, when it connects to Red Hat Network to retrieve information from your System Profile, which packages are downloaded, which packages are installed using the **Red Hat Update Agent**, and which packages are deleted from your system after installation. If you choose to install and delete packages yourself, it is not logged in this file. Red Hat Network recommends that you keep a log of actions not performed with the **Red Hat Update Agent**.

4.4. Configuration

The **Red Hat Update Agent** offers various options to configure its settings.

If you are not running the X Window System or prefer the command line version, skip to [Section 4.4.2, “Command Line Version”](#).

4.4.1. Using the Red Hat Update Agent Configuration Tool

You must be root to run the **Red Hat Update Agent Configuration Tool**. If started by a user other than root, the Red Hat Update Agent prompts you for the root password. The **Red Hat Update Agent Configuration Tool** can be started by typing the command `up2date --config` at a shell prompt (for example, an `xterm` or a `gnome-terminal`).

4.4.1.1. General Settings

The **General** tab allows you to enable an HTTP Proxy Server. If your network connection requires you to use an HTTP Proxy Server to make HTTP connections, select the **Enable HTTP Proxy** option and type your proxy server in the text field with the format `http://HOST:PORT`. For example, to use the proxy server `squid.mysite.org` on port 3128, you would enter `squid.mysite.org:3128` in the text field. Additionally, if your proxy server requires a username and password, select the **Use Authentication** option and enter your username and password in the respective text fields.

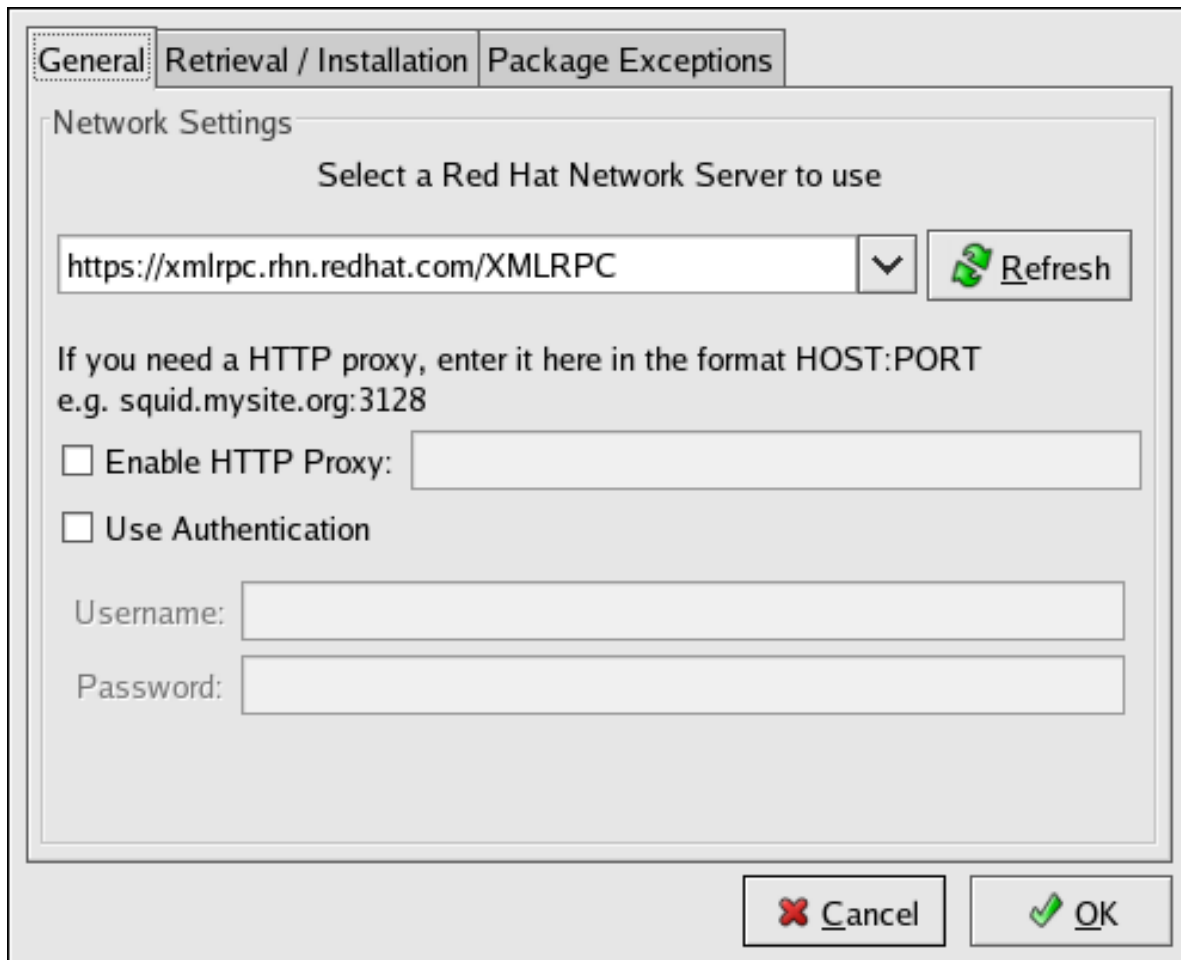


Figure 4.16. General Settings

In addition, RHN Proxy and Satellite customers have the option of selecting Red Hat Network Servers here. These customers should refer to the *RHN Client Configuration Guide* for detailed instructions.

4.4.1.2. Retrieval/Installation Settings

The **Retrieval/Installation** tab allows you to customize your software package retrieval and package w installation preferences.



Warning

You must use **Red Hat Update Agent** Version 2.5.4 or higher to upgrade your kernel automatically. **Red Hat Update Agent** will install the updated kernel and configure LILO or GRUB to boot the new kernel the next time the system is rebooted.

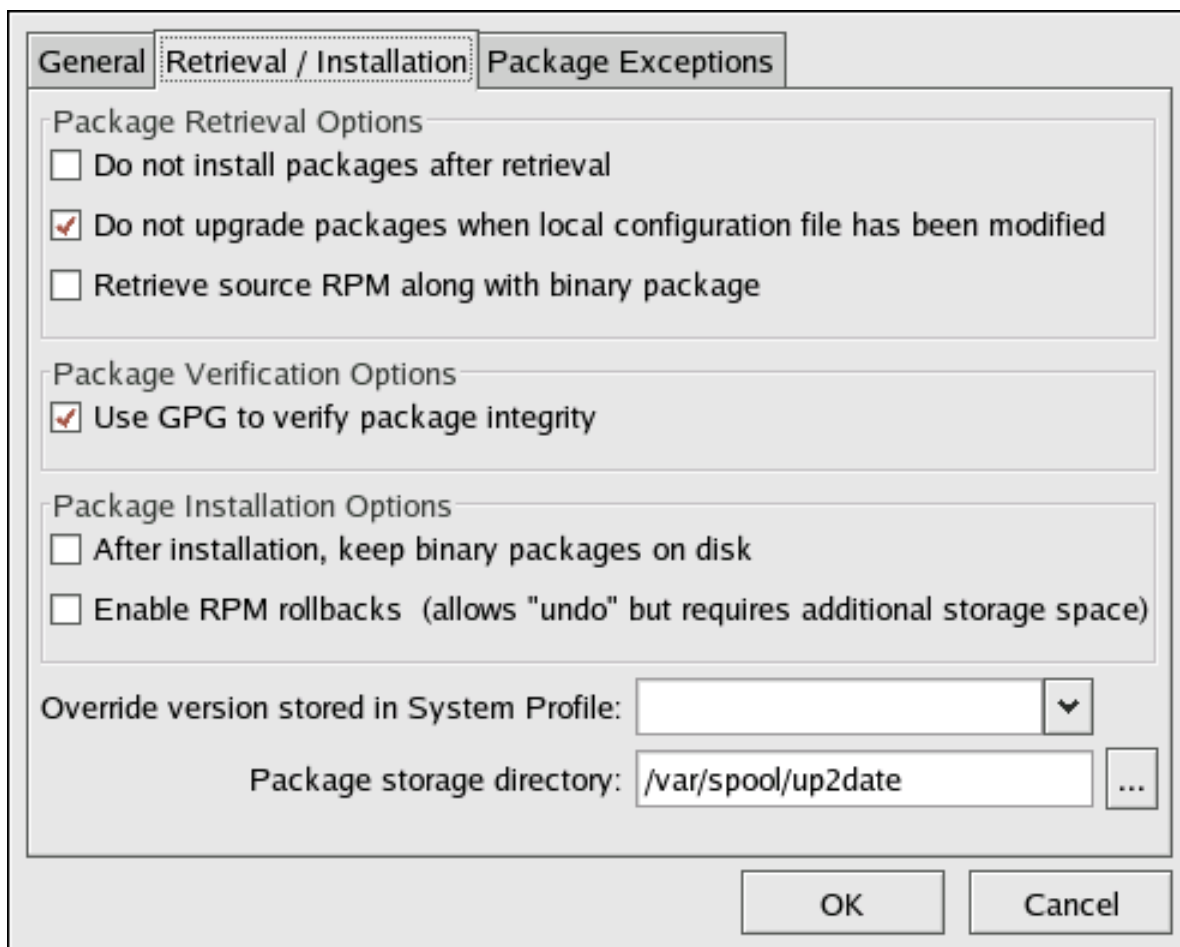


Figure 4.17. Retrieval/Installation Settings

The following package retrieval options can be selected (see [Figure 4.17, “Retrieval/Installation Settings”](#)):

- **Do not install packages after retrieval** — download selected RPM packages to the desired directory and ignore the installation preferences
- **Do not upgrade packages when local configuration file has been modified** — if the configuration file has been modified for a package such as **apache** or **squid**, do not attempt to upgrade it. This option is useful if you are installing custom RPMs on your system and you do not want them updated or reverted to the default Red Hat Enterprise Linux packages.
- **Retrieve source RPM along with binary package** — download both the source (***.src.rpm**) and the binary (***.[architecture].rpm**) files

The following installation options are configurable (see [Figure 4.17, “Retrieval/Installation Settings”](#)):

- **Use GPG to verify package integrity** — before installing packages, verify Red Hat's GPG signature (highly recommended for security reasons)
- **After installation, keep binary packages on disk** — save binary packages in the desired directory instead of deleting them after installation

The following additional options are configurable from this tab:

- **Override version stored in System Profile** — override the Red Hat Linux version in your System Profile
- **Package storage directory** — change the directory where packages are downloaded; the default location is `/var/spool/up2date/`

4.4.1.3. Package Exceptions Settings

The **Package Exceptions** tab allows you to define which packages to exclude from the list of updated RPM packages according to the package name or file name (see [Figure 4.18, “Package Exceptions Settings”](#)).

To define a set of packages to be excluded according to the package name, enter a character string including wild cards (*) in the **Add new** text field under in the **Package Names to Skip** section heading. A wild card at the end of the character string indicates that all packages beginning with the character string are excluded from the list. A wild card at the beginning of the character string indicates that any packages that end with the character string are excluded from the list.

For example, if the string `kernel*` is in the **Package Names to Skip** section, the **Red Hat Update Agent** will not display any packages beginning with kernel.

To exclude packages by file name, apply the same rules to the field below **File Names to Skip** section heading.

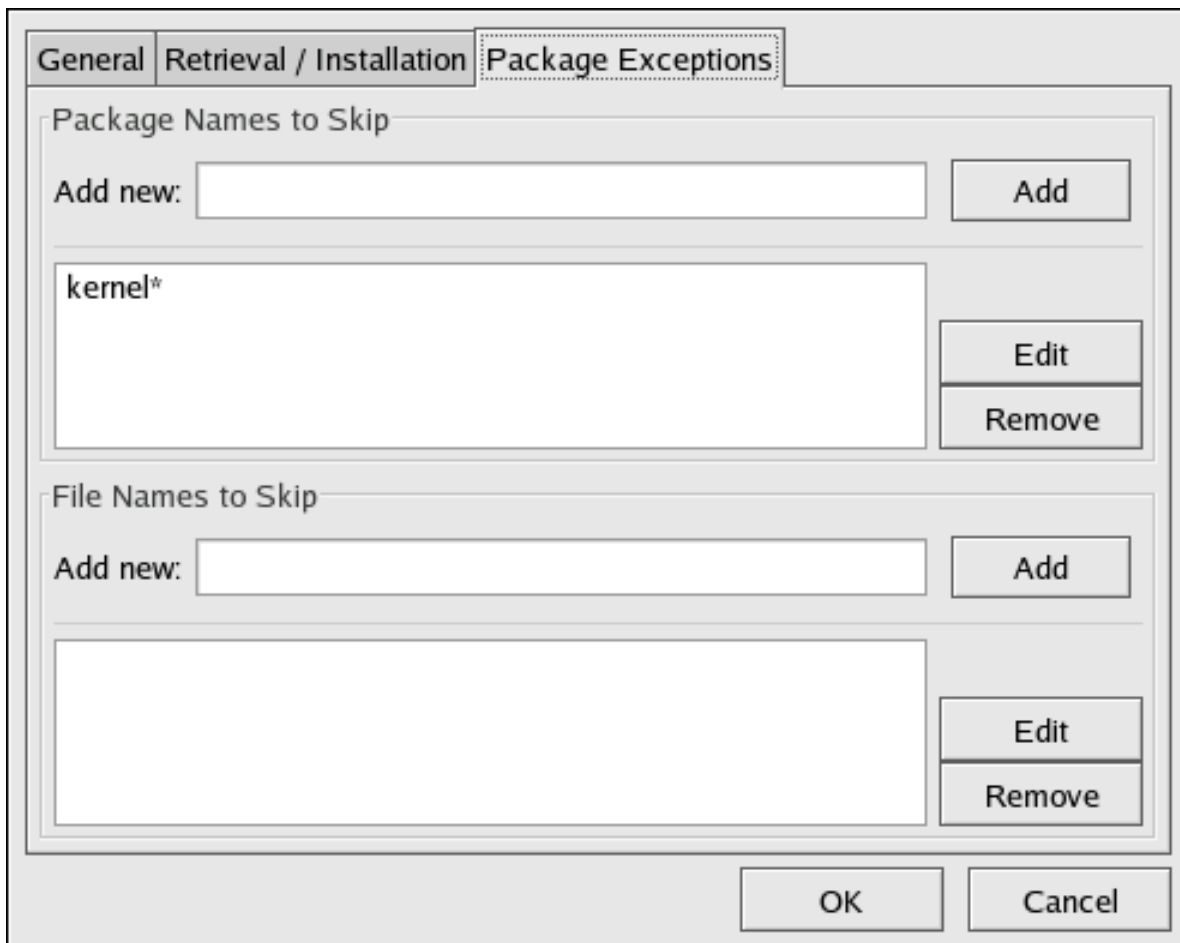


Figure 4.18. Package Exceptions Settings

4.4.2. Command Line Version

The command line version of this tool performs the same function as the graphical version. It allows you to configure the settings used by the **Red Hat Update Agent** and store them in the configuration file `/etc/sysconfig/rhn/up2date`.

To run the command line version of the **Red Hat Update Agent Configuration Tool**, use the following command:

```
up2date --nox --configure
```

You are presented with a list of options and their current values:

```
0. debug          No
1. isatty         Yes
2. depslist       []
3. networkSetup  Yes
4. retrieveOnly   No
5. enableRollbacks No
6. pkgSkipList    ['kernel*']
7. storageDir     /var/spool/up2date
8. adminAddress   ['root@localhost']
9. noBootLoader   No
10. serverURL     https://xmlrpc.rhn.redhat.com/XMLRPC
11. fileSkipList  []
12. sslCACert     /usr/share/rhn/RHNS-CA-CERT
13. noReplaceConfig Yes
14. useNoSSLForPackage No
15. systemIdPath  /etc/sysconfig/rhn/systemid
16. enableProxyAuth No
17. retrieveSource No
18. versionOverride
19. headerFetchCount 10
20. networkRetries 5
21. enableProxy    No
22. proxyPassword
23. noSSLServerURL http://xmlrpc.rhn.redhat.com/XMLRPC
24. keepAfterInstall No
25. proxyUser
26. removeSkipList ['kernel*']
27. useGPG         Yes
28. gpgKeyRing     /etc/sysconfig/rhn/up2date-keyring.gpg
29. httpProxy
30. headerCacheSize 40
31. forceInstall   No
```

Enter number of item to edit <return to exit, q to quit without saving>:

Enter the number of the item to modify and enter a new value for the option. When you finish changing your configuration, press **Enter** to save your changes and exit. Press **q** and then **Enter** to quit without saving your changes.



Important

Although this is not configurable, users should still make note that the port used by the **Red Hat Update Agent** is 443 for SSL (HTTPS) and 80 for non-SSL (HTTP). By default, **up2date** uses SSL only. For this reason, users should ensure that their firewalls allow

connections over port 443. To bypass SSL, change the protocol for serverURL from `https` to `http` in the `/etc/sysconfig/rhn/up2date` configuration file.

4.5. Registering with Activation Keys

In addition to the standard **Red Hat Update Agent** interface, **up2date** offers a utility aimed at batch processing system registrations: activation keys. Each unique key can be used to register Red Hat Enterprise Linux systems, entitle them to an RHN service level, and subscribe them to specific channels and system groups, all in one action. This automation bypasses entitlement and registration via Red Hat Network Registration Client and Red Hat Update Agent.

Alternatively, both the Red Hat Network Registration Client and Red Hat Update Agent offer the activation keys utility **rhnreg_ks** as part of their packages.



Note

Systems running Red Hat Enterprise Linux 2.1 need version 2.9.3-1 or higher of the **rhn_register** package. It is highly recommended that you obtain the latest version before using activation keys.

Before using an activation key you must first generate one through the RHN website. Refer to

[Section 7.4.6, “Activation Keys — !\[\]\(6059a5aa8b4ca7bb793408023d6c6e42_img.jpg\) ”](#) for precise steps.

To use an activation key, run the following command as root from a shell prompt on the system to be registered:

```
rhnreg_ks --activationkey=7202f3b7d218cf59b764f9f6e9fa281b
```

The precise value of the activation key varies.

Systems running Red Hat Enterprise Linux 2.1 substitute the `--serialnumber` option for the `--activationkey` option:

```
rhnreg_ks --serialnumber=7202f3b7d218cf59b764f9f6e9fa281b
```


In addition, Provisioning-entitled systems may use multiple activation keys at once, either at the command line or within kickstart profiles. This allows Administrators to include a variety of values without creating a special key for the desired results. To do this, specify the keys separated by commas, like this:

```
rhnreg_ks --activationkey=7202f3b7d218cf59b764f9f6e9fa281b,\ 39f41081f0329c20798876f37cb9p6a3
```



Note

The trailing backslash (`\`) in this command example is a continuation character; it may safely be omitted.


Refer to [Section 7.4.6.2, “Using Multiple Activation Keys at Once](#) —  ” to understand how differences in activation keys are handled.

The above command performs all the actions of the **Red Hat Network Registration Client** and the registration function of the **Red Hat Update Agent**. Do not run either of these applications for registration after running `rhnreg_ks`.

A System Profile, including software and hardware information, is created for the system and sent to the RHN Servers along with the unique activation key. The system is registered with RHN under the account used to generate the key, entitled to an RHN service offering, and subscribed to the RHN channels and system groups selected during key generation. The system is not subscribed to channels that contain packages unsuitable for the system. For example, a Red Hat Enterprise Linux 2.1 system cannot be subscribed to the Red Hat Enterprise Linux 3 channel.

The unique Digital Certificate for the system is generated on the system in the file `/etc/sysconfig/rhn/systemid`.

When using activation keys to assign channels, consider these rules:

- A key may specify either zero or one base channel. If specified, it must be a custom base channel. If not, the base channel corresponding to the system's Red Hat distribution is chosen. For instance, you may not subscribe a Red Hat Enterprise Linux 2.1 system to the Red Hat Enterprise Linux 3 channel.
- A key may specify any number of child channels. For each child channel, subscription is attempted. If the child channel matches the system's base channel, subscription succeeds. If it does not, the subscription fails silently. Refer to [Section 7.6, “Channels”](#) for more information.
- Keys may be modified by any user with the role of Activation Key Administrator or Satellite Administrator (or both). These permissions are set through the **Users** tab of the RHN website. Refer to [Section 7.9, “Users](#) —  ” for details.
- Systems registered by activation keys are tied to the organization account in which the key was created, not the key itself. After registration, a key can be deleted safely without any effect on the systems it was used to register.

4.6. Registering a System to an Organization

RHN Satellite now supports the Organizations feature, which allows administrators to appropriate software and system entitlements across various organizations, as well as control an organization's access to systems management. Systems can now be registered directly to an organization.

To register a system to an organization on a satellite, you can use the username and password of an account that is created within that organization. For example, if there is an organization called **Sales Team**, with a username `salesadmin` and password `abc123`, using these credentials assures that a system is registered to the proper organization.

For example:

```
rhnreg_ks --user=salesadmin --password=abc123
```



Important

The `--orgid` option (for RHEL 4 and 5) and the `--orgpassword` option (in RHEL 4) in the `rhncfg_ks` command *are not related* to the Organizations feature and should not be used in the context of registering systems to organizations.

For more information about the Organizations feature, refer to [Section 7.11.1, “Admin ¶ Organizations”](#).

Red Hat Network Daemon

The Red Hat Network Daemon (**rhnsd**) periodically connects to Red Hat Network Satellite to check for updates and notifications. The daemon, which runs in the background, is typically started from the initialization scripts in `/etc/init.d/rhnsd` or `/etc/rc.d/init.d/rhnsd`.

To check for updates, **rhnsd** runs an external program called **rhn_check** located in `/usr/sbin/`. This is a small application that makes the network connection to RHN. The Red Hat Network Daemon does not listen on any network ports or talk to the network directly. All network activity is done via the **rhn_check** utility.

5.1. Configuring

The Red Hat Network Daemon can be configured by editing the `/etc/sysconfig/rhn/rhnsd` configuration file. This is actually the configuration file the **rhnsd** initialization script uses. The most important setting offered by the daemon is its check-in frequency. The default interval time is four hours (240 minutes). If you modify the configuration file, you must (as root) restart the daemon with the command **service rhnsd restart** or `/etc/rc.d/init.d/rhnsd restart`.



Important

The minimum time interval allowed is one hour (60 minutes). If you set the interval below one hour, it will default to four hours (240 minutes).

5.2. Viewing Status

You can view the status of the **rhnsd** by typing the command **service rhnsd status** or `/etc/rc.d/init.d/rhnsd status` at a shell prompt.

5.3. Disabling

To disable the daemon, (as root) run the **ntsysv** utility and uncheck **rhnsd**. You can also (as root) execute the command **chkconfig rhnsd off**. Using these two methods only disables the service the next time the system is started. To stop the service immediately, use the command **service rhnsd stop** or `/etc/rc.d/init.d/rhnsd stop`.

5.4. Troubleshooting

If you see messages indicating that checkins are not taking place, the RHN client on your system is not successfully reaching Red Hat Network Satellite. Make certain:

- your client is configured correctly.
- your system can communicate with RHN via SSL (port 443). You may test this by running the following command from a shell prompt:

```
telnet xmlrpc.rhn.redhat.com 443
```

- the Red Hat Network Daemon is activated and running. You may ensure this by running the following commands:

```
chkconfig --level 345 rhnsd on
```

```
service rhnsd start
```

If these are correct and your systems still indicate they are not checking in, please contact our technical support team.

Red Hat Network Alert Notification Tool

The **Red Hat Network Alert Notification Tool** is a notifier that appears on the desktop panel and alerts users when software package updates are available for their Red Hat Enterprise Linux 4 systems. The list of updates is retrieved from the RHN Servers. The system does not have to be registered with Red Hat Network to display a list of updates; however, retrieving the updates with the **Red Hat Update Agent** requires registration with Red Hat Network and a subscription to an RHN service offering. The notifier does not send any identifiable information about the user or the system to the RHN Servers.

To use the **Red Hat Network Alert Notification Tool**, you must install the **rhn-applet** RPM package and use the X Window System.

The **Red Hat Network Alert Notification Tool** appears on the panel by default as shown in [Figure 6.1, “GNOME Panel with Red Hat Network Alert Notification Tool”](#).




Figure 6.1. GNOME Panel with **Red Hat Network Alert Notification Tool**

If it does not appear on the panel, you can add it:

- In Red Hat Enterprise Linux 4 and later, select Applications (the main menu on the panel) => **System Tools** => **Red Hat Network Alert Icon**. To ensure the icon appears on subsequent sessions, select the **Save current setup** checkbox when logging out.
- In Red Hat Enterprise Linux 2.1, select the **Main Menu Button** => **Panel** => **Add to Panel** => **Applet** => **Red Hat Network Monitor**. To move it around the panel, right-click on the applet, select **Move**, move the mouse left and right until it is in the desired location, and click the mouse to place the applet.

6.1. Configuring the Applet

The first time the **Red Hat Network Alert Notification Tool** is run, a configuration wizard starts. It displays the terms of service and allows the user to configure an HTTP proxy as shown in [Figure 6.2, “HTTP Proxy Configuration”](#).

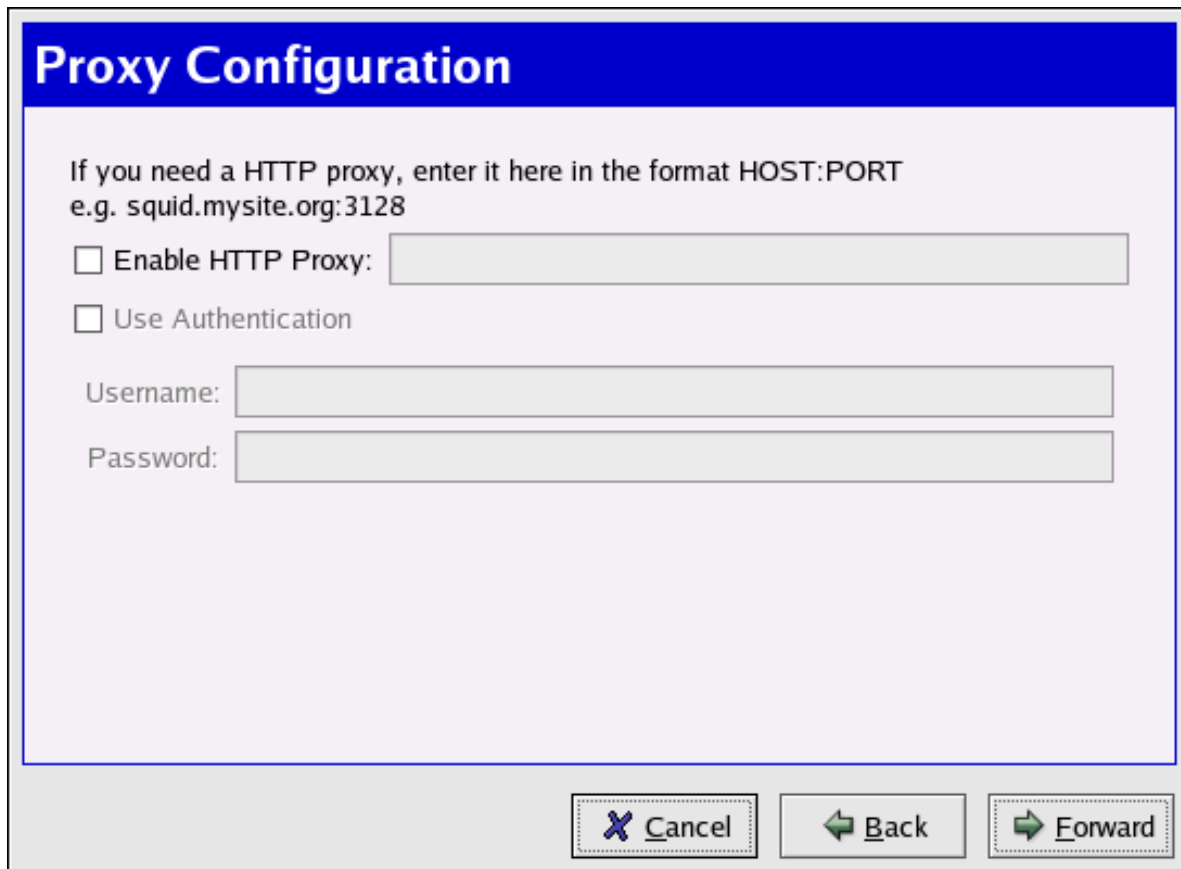



Figure 6.2. HTTP Proxy Configuration

If your network connection requires you to use an HTTP Proxy Server to make HTTP connections, on the **Proxy Configuration** screen, type your proxy server in the text field with the format HOST:PORT. For example, to use the proxy server `http://squid.mysite.org` on port 3128, enter **squid.mysite.org:3128** in the text field. Additionally, if your proxy server requires a username and password, select the **Use Authentication** option and enter your username and password in the respective text fields.

 **Tip**
To run the configuration wizard again, right-click on the applet, and select **Configuration**.

Your preferences are written to the `.rhn-applet.conf` file in your home directory. The **Red Hat Network Alert Notification Tool** also uses the system-wide configuration file `/etc/sysconfig/rhn/rhn-applet`. The setting for `server_url` should be set to your satellite server. For example:

```
server_url=http://YourRHN_Satellite.com/APPLET
```

Or, for SSL:

```
server_url=https://YourRHN_Satellite.com/APPLET
```


You can also configure the **Red Hat Network Alert Notification Tool** to ignore specific packages. To select these packages, click on the applet and select the **Ignored Packages** tab.

6.2. Notification Icons

The applet displays a different icon, depending on the status of the updates. [Table 6.1, “Red Hat Network Alert Notification Tool Icons”](#) shows the possible icons and their meaning.






Icon	Description
	Updates are available
	System is up-to-date
	Checking for updates
	Error has occurred


Table 6.1. **Red Hat Network Alert Notification Tool Icons**

If you see the  icon, it is strongly recommended that you apply the updates. Refer to [Section 6.4, “Applying Updates”](#) for information on applying updates.

If you have scheduled updates to be installed, you can watch the applet icon to determine when

updates are applied. The icon changes to the  icon after the Errata Updates are applied.

If you apply a kernel update (or the kernel update is automatically applied), the applet displays the

 icon until the system is rebooted with the new kernel. If you double-click on the applet, the **Available Updates** tab displays a list of packages that can be updated on your system.

6.3. Viewing Updates

Clicking on the **Red Hat Network Alert Notification Tool** displays a list of available updates. To alter your list of excluded packages, click the **Ignored Packages** tab and make your modifications.

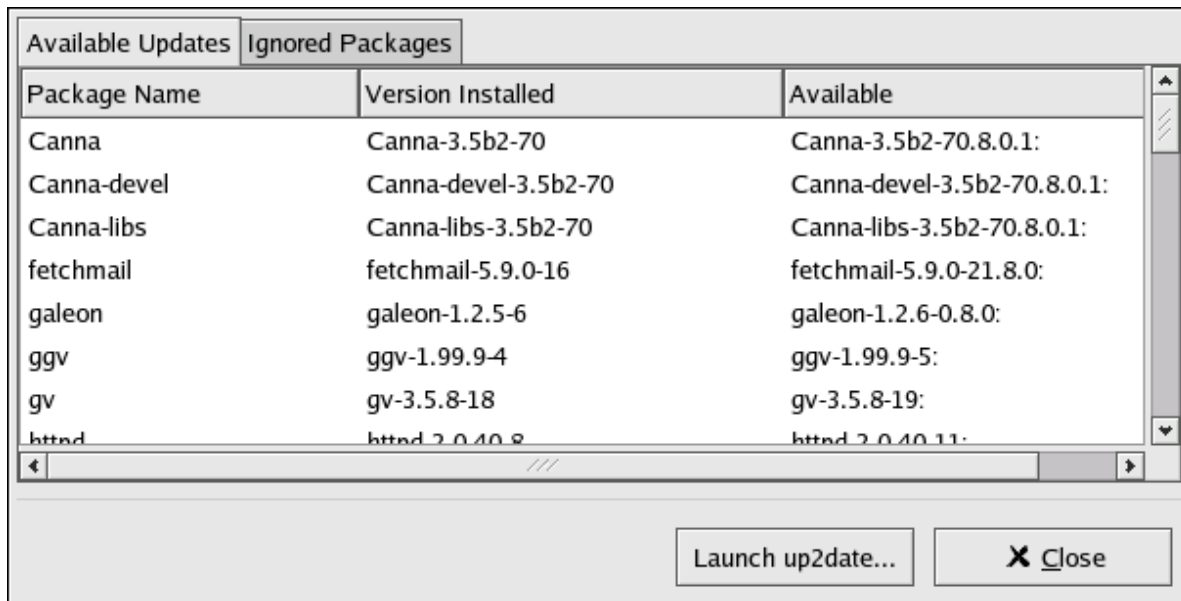


Figure 6.3. Available Updates

6.4. Applying Updates

If the system is registered with RHN and entitled to a service offering, you can apply the Errata Updates with the **Red Hat Update Agent**. To launch the **Red Hat Update Agent**, click on the applet, and then click on the **Launch up2date** button. You can also right-click on the icon and select **Launch up2date**. For more information on the **Red Hat Update Agent**, refer to [Chapter 4, Red Hat Update Agent](#).

6.5. Launching the RHN Website

The simplest way to obtain a comprehensive view of your system's status is to access the RHN website. This can be accomplished through the **Red Hat Network Alert Notification Tool** by right-clicking on it and selecting **RHN Website**. For more information on the RHN website, refer to [Section 7.1, "Navigation"](#).

Red Hat Network Website

You can use the Red Hat Network website to manage multiple Red Hat Enterprise Linux systems simultaneously, including viewing Errata Alerts, applying Errata Updates, and installing packages. This chapter seeks to identify all of categories, pages, and tabs within the website and explain how to use them.

7.1. Navigation

The **Top Navigation Bar** is divided into tabs. Satellite Administrators see the following **Top Navigation Bar**. Note that only RHN Satellite customers see the Monitoring and **Admin** tabs.



Figure 7.1. Top Navigation bar — RHN Satellite

The **Left Navigation Bar** is divided into pages. The links are context-sensitive and may vary slightly between RHN Satellite and non-Satellite web interfaces. The following is an example of the **Left Navigation Bar** for the **Users** tab.



Figure 7.2. Left Navigation Bar — Users

Some pages have sub-tabs. These tabs offer an additional layer of granularity in performing tasks for systems or users. The following is a menu bar for all System Details sub-tabs. This system has Management and Provisioning entitlements, but not Monitoring:




Figure 7.3. Sub-Tabs — System Details

7.1.1. Entitlement Views

Keep in mind, since this guide covers all entitlement levels, some tabs, pages, and even whole categories described here may not be visible to you. For this reason, icons are used here to identify which functions are available to each entitlement level.

Icon	Entitlement
	Management or higher


Icon	Entitlement
	Provisioning
	Monitoring



Table 7.1. Entitlement Icons








If no icon follows a category, page, or tab label within this chapter, the area described is available to all Red Hat Network users. If an icon does follow, the associated entitlement is needed. Remember that Provisioning inherits all of the functions of Management.








If an icon precedes a paragraph, only the specific portion of the page or tab discussed afterward requires the indicated entitlement level. When a page or tab is associated with a particular entitlement level, all of its tabs and subtabs require at least the same entitlement level but may need a higher entitlement. Regardless, each tab is identified separately.

7.1.2. Categories and Pages

This section summarizes all of the categories and primary pages (those linked from the top and left navigation bars) within the RHN website. It does not list the many subpages, tabs and subtabs accessible from the left navigation bar and individual pages. Each area of the website is explained in detail later in this chapter:

- **Overview** — View and manage your primary account information and obtain help.
 - **Overview** — Obtain a quick overview of your account. It notifies you if your systems need attention, provides a quick link to go directly to them, and displays the most recent Errata Alerts for your account.
 - **Your Account** — Update your personal profile and addresses.
 - **Your Preferences** — Indicate if you wish to receive email notifications about Errata Alerts for your systems, set how many items are displayed at one time for lists such as system lists and system group lists, set your time zone, and identify your contact options.
 - **Locale Preferences** — Configure language, timezone, and other customizations for your particular locale.
 - **Subscription Management** — Manage base and add-on system entitlements, such as Management, Provisioning, and Virtualization.
- **Systems** — Manage all of your systems (including virtual guest systems) here.
 - **Overview** —  — View a summary of your systems or system groups showing how many Errata Alerts each system has and which systems are entitled.
 - **Systems** — Select and view subsets of your systems by specific criteria, such as Virtual Systems, Unentitled, Recently Registered, Proxy, and Inactive.
 - **System Groups** —  — List your system groups. Create additional groups.

- **System Set Manager** —  — Perform various actions on collective sets of systems, including scheduling errata updates, package management, listing and creating new groups, and managing channel entitlements.
- **Advanced Search** —  — Quickly search all of your systems by specific criteria, such as name, hardware, devices, system info, networking, packages, and location.
- **Activation Keys** —  — Generate an activation key for an RHN-entitled system. This activation key can be used to grant a specified level of entitlement or group membership to a newly registered system with the **rhnreg_ks** command.
- **Stored Profiles** —  — View system profiles used to provision systems.
- **Custom System Info** —  — Create and edit system information keys containing completely customizable values that can be assigned while provisioning systems.
- **Kickstart** —  — Display and modify various aspects of kickstart profiles used in provisioning systems.
- **Errata** — View and manage Errata Alerts here.
 - **Errata** — List Errata Alerts and download associated RPMs.
 - **Advanced Search** — Search Errata Alerts based on specific criteria, such as synopsis, advisory type, and package name.
 - **Manage Errata** — Manage the errata for an organization's channels.
 - **Clone Errata** — Clone errata for an organization for ease of replication and distribution across an organization.
- **Channels** — View and manage the available RHN channels and the files they contain.
 - **Software Channels** — View a list of all software channels and those applicable to your systems.
 - **Package Search** — Search packages using all or some portion of the package name, description, or summary, with support for limiting searches to supported platforms.
 - **Manage Software Channels** —  — Create and edit channels used to deploy configuration files.
- **Configuration** — Keep track of and manage configuration channels, actions, and individual configuration files.
 - **Overview** — A general dashboard view that shows a configuration summary
 - **Configuration Channels** — List and create configuration channels from which any subscribed system can receive configuration files

- **Configuration Files** — List and create files from which systems receive configuration input
- **Systems** — List the systems that have RHN-managed configuration files.
- **Schedule** — Keep track of your scheduled actions.
 - **Pending Actions** — List scheduled actions that have not been completed.
 - **Failed Actions** — List scheduled actions that have failed.
 - **Completed Actions** — List scheduled actions that have been completed. Completed actions can be archived at any time.
 - **Archived Actions** — List completed actions that have been selected to archive.
- **Users** —  — View and manage users for your organization.
 - **User List** —  — List users for your organization.
- **Monitoring** —  — Run probes and receive notifications regarding systems.
 - **Status** —  — View probes by state.
 - **Notification** —  — View contact methods established for your organization.
 - **Probe Suites** —  — Manage your monitoring infrastructure using suites of monitoring probes that apply to one or more assigned systems.
 - **Scout Config Push** —  — Displays the status of your monitoring infrastructure.
- **Admin** (visible only to Satellite administrators) — List, create, and manage one or more Satellite organizations, from which the Satellite administrator can assign channel entitlements, create and assign administrators for each organization, and other tasks.
 - **Organizations** — List and create new organizations
 - **Subscriptions** — List and manage the software and system entitlements for all organizations across the Satellite.
 - **Users** — List all users on the Satellite, across all organizations. Click individual usernames to change administrative privileges for the user.



Note

Users created for organization administration can only be configured by the organization administrator, *not* the Satellite administrator.

- **Satellite Configuration** — Make general configuration changes to the Satellite, including Proxy settings, Certificate configuration, Bootstrap Script configuration, Organization changes, and Restart the Satellite Server.
- **Task Engine Status** — configures the daemon that runs on the Satellite server itself and performs routine operations, such as database cleanup, Errata mailings, and other tasks that are performed in the background.

7.1.3. Errata Alert Icons

Throughout Red Hat Network you will see three Errata Alert icons.  represents a Security Alert.



represents a Bug Fix Alert.  represents an Enhancement Alert.

In the **Overview** page, click on the Errata advisory to view details about the Erratum or click on the number of affected systems to see which are affected by the Errata Alert. Both links take you to tabs of the **Errata Details** page. Refer to [Section 7.5.2.2, “Errata Details”](#) for more information.

7.1.4. Quick Search

In addition to the Advanced Search functionality for Packages, Errata, Documentation, and Systems offered within some categories, RHN Satellite also offers a Quick Search tool near the top of each page. To use it, select the search item (choose from **Systems**, **Packages**, **Documentation**, and **Errata**) and type a keyword to look for a name match. Click the **Search** button. Your results appear at the bottom of the page.

If you misspell a word during your search query, the Satellite search engine institutes *approximate string* (or *fuzzy string*) matching, giving you results that may be similar in spelling to your misspelled queries.

For example, if you want to search for a certain development system called **test-1.example.com** that is registered to the Satellite, but you misspell your query **tset**, the test-1.example.com system still appears in the search results



Note

If you add a distribution or register a system to a Satellite, it may take several minutes for it to be indexed and appear in search results.

For advanced System searches, refer to [Section 7.4.5, “Advanced Search](#) — ”.

For advanced Errata searches, refer to [Section 7.5.3, “Advanced Search”](#).

For advanced Package searches, refer to [Section 7.6.2, “Package Search”](#).

For advanced Documentation searches, refer to [Section 7.12.8, “Search”](#).

7.1.5. Systems Selected

Also near the top of the page is a tool for keeping track of the systems you have selected for use in the System Set Manager. It identifies the number of selected systems at all times and provides the means

to work with them. Clicking the **Clear** button deselects all systems, while clicking the **Manage** button launches the System Set Manager with your selected systems in place.

These systems can be selected in a number of ways. Only systems with at least a Management entitlement are eligible for selection. On all system and system group lists, a Select column exists for this purpose. Select the checkboxes next to the systems or groups and click the **Update List** button below the column. Each time, the Systems Selected tool at the top of the page changes to reflect the new number of systems ready for use in the System Set Manager. Refer to [Section 7.4.4, “System Set](#)

[Manager](#) —  ” for details.

7.1.6. Lists

The information within most categories is presented as lists. These lists have some common features for navigation. For instance, you can navigate through virtually all lists by clicking the back and next arrows above and below the right side of the table. Some lists also offer the ability to retrieve items alphabetically by clicking the letters above the table.

7.2. Logging into the RHN Website

Use a web browser to navigate to <http://rhn.redhat.com>. RHN displays the login page shown below unless one of two things is true:

- You have recently logged into your account at <http://www.redhat.com>.
- You have recently either logged into RHN or recently visited the new account verification page.

If you have recently logged into <http://rhn.redhat.com> or <http://www.redhat.com>, you are automatically authenticated and redirected to the **Your RHN** page.

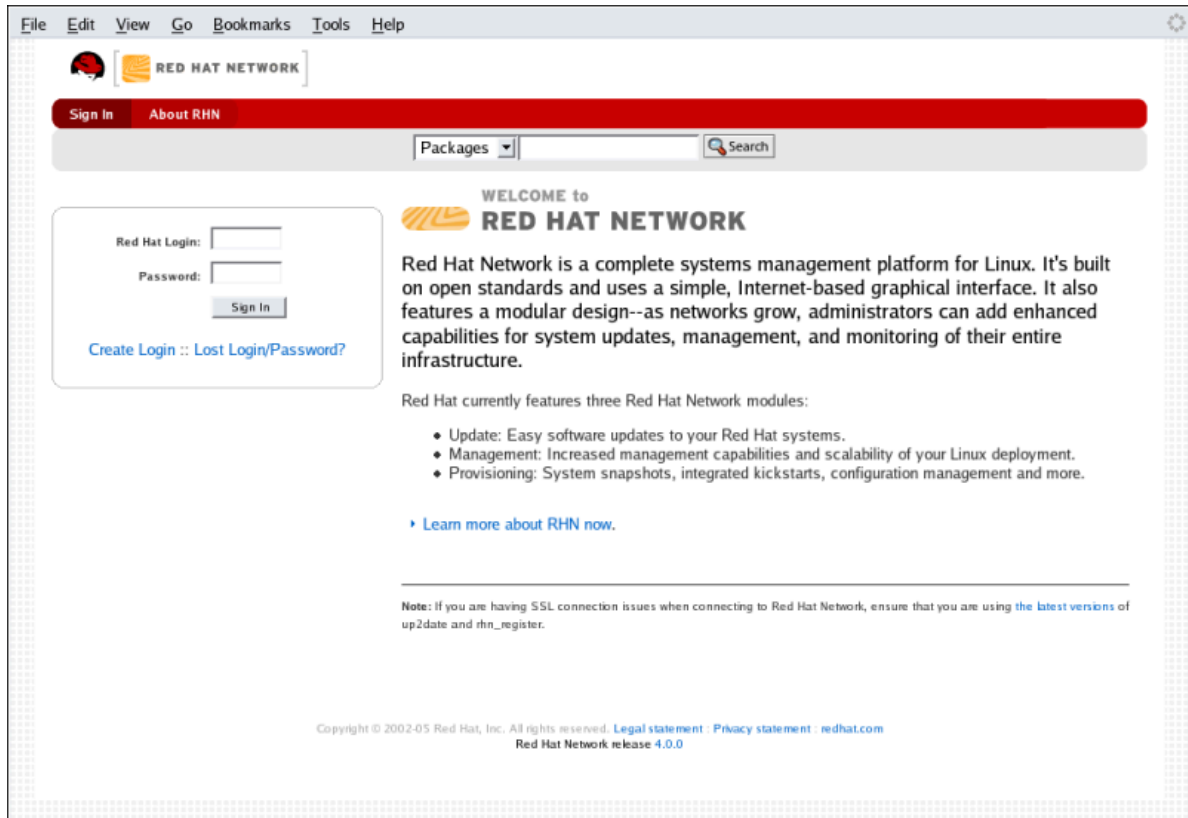


Figure 7.4. RHN Website

If you have not registered a system yet or do not have a redhat.com account, create a new account by following the **Learn More** link, then selecting **Create Login** on the resulting page. After creating a new user account, you must register a system before using RHN. Refer to [Chapter 4, Red Hat Update Agent](#) for step-by-step instructions.

After registering your system with Red Hat Network, go back to <http://rhn.redhat.com> and complete the username and password fields with the same information established during registration. Once complete, press the **Log In** button to continue.



Tip

You may click the **Sign In** tab at the top of the screen to display the fields if they are not already visible.

If you have not previously accepted the **RHN Site Terms** and the **T7** agreement, you will be asked to do so now before proceeding. To read the content of either agreement, click on its title, which will open a new window. When ready to proceed, select the checkbox indicating your acceptance of the agreements and press the **Continue** button.



Note

You must accept both the Site Terms and the T7 agreement in order to use RHN.

Once you have accepted the agreements and pressed the **Continue** button, RHN displays the **Overview** page.

7.3. Overview

After logging into the web interface of Red Hat Network, the first page to appear is **Overview**. This page contains important information about your systems, including summaries of system status, actions, and Errata Alerts.



Tip

If you are new to the RHN web interface, read *Section 7.1, “Navigation”* to become familiar with the layout and symbols used throughout the interface.

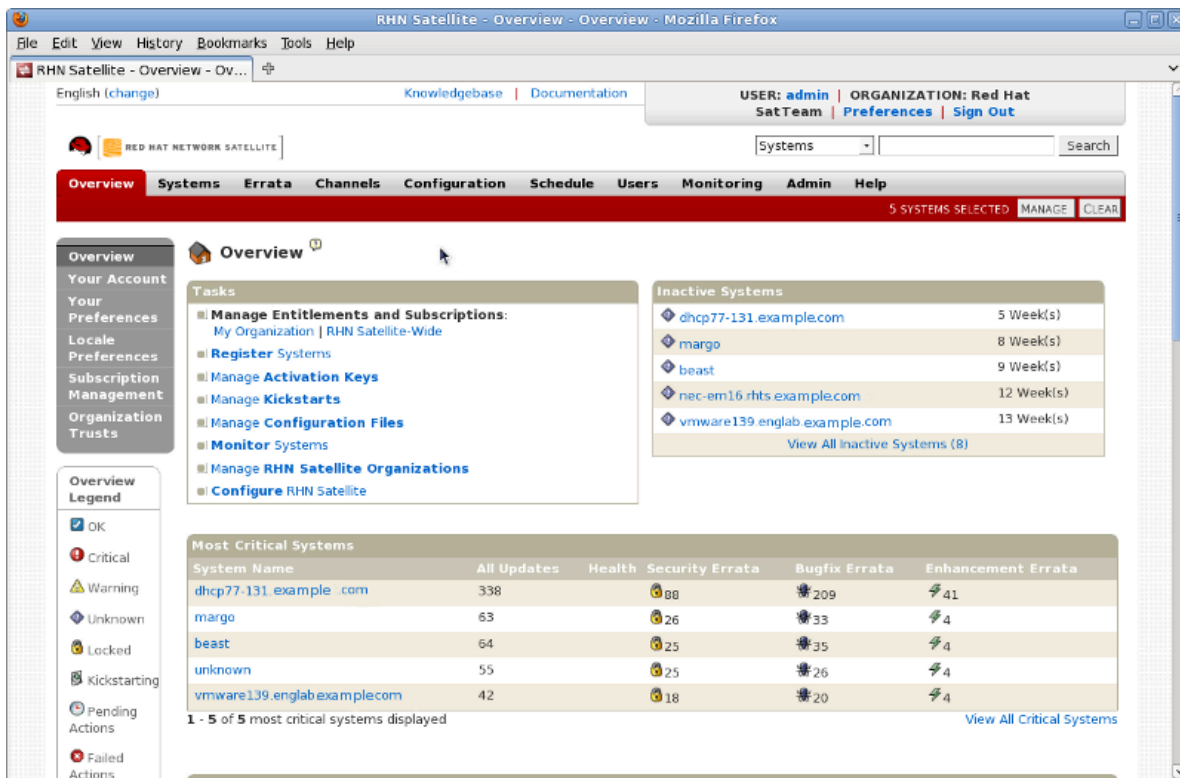




Figure 7.5. Overview

This page is broken into functional areas, with the most critical areas displayed first. Users can control which of the following areas are displayed by making selections on the **Overview** **Your Preferences** page. Refer to *Section 7.3.2, “Your Preferences”* for more information.

- The **Tasks** area lists the most common tasks that an administrator performs via the web. Click on any of the links to be taken to the page within RHN that allows you to accomplish that task.
- To the right is the **Inactive System** listing. If any systems have not been checking in to RHN, they are listed here. Highlighting them in this way allows an administrator to quickly select those systems for troubleshooting.

-  — Customers with Monitoring enabled on their Satellite can also choose to include a list of all probes in the Warning state.
-  — Customers with Monitoring enabled on their Satellite can also choose to include a list of all probes in the Critical state.
- The **Critical Systems** section lists the most critical systems within your organization. It provides a link to quickly view those systems, and displays a summary of the errata updates that have yet to be applied to those systems. Click on the name of the system to be taken to the **System Details** page of that system and apply the errata updates. Below the list is a link to the **Out of Date** systems page.
- Next is the **Recently Scheduled Actions** section. Action that are less than thirty days old are considered recent. This section allows you to see all actions and their status: whether they have failed, completed, or are still pending. Click on the label of any given actions to view the details page for that action. Below the list is a link to the **Pending Actions** page, which lists all actions that have not yet been picked up by your client systems.
- The **Relevant Security Errata** section lists the security errata that are available and have yet to be applied to some or all of your client systems. It is critical that you apply these security errata to keep your systems secure. Below this section are links to all errata and to those errata that apply to your systems.
- The **System Groups** section lists the groups (if any) and indicates whether the systems in those groups are fully updated. Click on the link below this section to be taken to the **System Groups** page, from which you can chose **System Groups** to use with the System Set Manager.
- The **Recently Registered Systems** lists the systems that have been added to the Satellite in the past 30 days. Click the system's name to go the **System Details** page for that particular system.

You can return to this page by clicking **Overview** on the left navigation bar.

7.3.1. Your Account

The **Your Account** page allows you to modify your personal information, such as name, password, and title. To modify any of this information, make the changes in the appropriate text fields and click the **Update** button in the bottom right-hand corner.

Remember, if you change your Red Hat Network password (the one used to log into RHN and redhat.com), you will not see your new one as you type it for security reasons. Also for security, your password is represented by 12 asterisks no matter how many characters it actually contains. Replace the asterisks in the **Password** and **Password Confirmation** text fields with your new password.

7.3.1.1. Addresses

The **Addresses** page allows you to manage your mailing, billing and shipping addresses, as well as the associated phone numbers. Just click **Edit this address** below the address to be modified, make the changes, and click **Update**.

7.3.1.2. Change Email

The email address listed in the **Your Account** page is the address to which Red Hat Network sends email notifications if you select to receive Errata Alerts or daily summaries for your systems on the **Your Preferences** page.

To change your preferred email address, click **Change Email** in the left navigation bar. You are then asked for the new email address. Enter it and click the **Update** button. A confirmation email is sent to the new email address; responding to the confirmation email validates the new email address. Note that false email addresses such as those ending in "@localhost" are filtered and rejected.

7.3.1.3. Account Deactivation

The **Account Deactivation** page provides a means to cancel your Red Hat Network service. Click the **Deactivate Account** button to deactivate your account. The web interface returns you to the login screen. If you attempt to log back in, an error message advises you to contact the Satellite Administrator for your organization. Note that if you are the only Satellite Administrator for your organization, you are unable to deactivate your account.

7.3.2. Your Preferences

The **Your Preferences** page allows you to configure Red Hat Network options, including:

- Email Notifications — Determine whether you want to receive email every time an Errata Alert is applicable to one or more systems in your RHN account.



Important

This setting also enables Management and Provisioning customers to receive a daily summary of system events. These include actions affecting packages, such as scheduled Errata Updates, system reboots, or failures to check in. In addition to selecting this checkbox, you must identify each system to be included in this summary email. (By default, all Management and Provisioning systems are included in the summary.) This can be done either individually through the **System Details** page or for multiple systems at once through the **System Set Manager** interface. Note that RHN sends these summaries only to verified email addresses. To disable all messages, simply deselect this checkbox.

- RHN List Page Size — Maximum number of items that appear in a list on a single page. If more items are in the list, clicking the **Next** button displays the next group of items. This preference applies to system lists, Errata lists, package lists, and so on.
- "Overview" Start Page — select the information areas that are displayed on the **Overview** Start Page. Check the box to the left of the information area you would like to include.

After making changes to any of these options, click the **Save Preferences** button in the bottom right-hand corner.

7.3.3. Locale Preferences

The **Overview** ▯ **Locale Preferences** page allows each user to tailor their RHN interface to the local time and their preferred language. Select the appropriate timezone from the **Time Zone** dropdown box, then click the **Save Preferences** button to apply the selection.

When the language preference is set to **Use Browser Settings**, RHN uses the language preference from the user's browser (such as **Firefox**) to determine which language to use for the web interface. When one of the listed languages is selected, the user sees the web interface in that language each time they log in, regardless of their browser's settings. Choosing a preferred language may be helpful for users traveling abroad. To select a default language, click the radio button to the left of the appropriate language and click the **Save Preferences** button to apply the change.

7.3.4. Subscription Management

To use all of the features of RHN, your systems must be *entitled* — subscribed to an RHN service level. Use the **System Entitlements** page to configure which systems are entitled to which service offerings. There are six primary types of entitlements:

- **Update** — manages a single Red Hat Enterprise Linux system. It includes Errata Alerts, Scheduled Errata Updates, Package Installation, and the **Red Hat Update Agent**.
- **Management** — manages multiple systems with multiple system administrators. In addition to the features of the Update offering, it includes system group management, user management, and the **System Set Manager** interface to quickly perform actions on multiple systems.
- **Provisioning** — offers the highest level of functionality. It should be used to provision multiple systems that will need to be re-installed and reconfigured regularly. The Provisioning offering provides tools for kickstarting machines, managing their configuration files, conducting snapshot rollbacks, and inputting searchable custom system information, as well as all of the functionality included in the Management service level.
- **Monitoring** — monitors the health of multiple systems. The Monitoring offering provides probes that watch system metrics and notify Administrators when changes occur. Such notifications alert Administrators to system performance degradation before it becomes critical.
- **Virtualization** — applies to virtual host systems. Virtual hosts with this entitlement may register as many as four guest systems without violating RHN's Service Level Agreement. Guest systems may be subscribed to any channel with the **virtualization-free** channel group label without consuming channel entitlements. Subscribing a guest to any channel that does not belong to **virtualization-free**, such as a Directory Server or RHN Satellite channel, consumes an additional channel entitlement.
- **Virtualization Platform** — also applies to virtual host systems. Host systems to which this entitlement apply may register an unlimited number of virtual guests without invalidating your Service Level Agreement. Guests of a host with this entitlement may subscribe to any channel that has the **virtualization-platform-free** content group label without consuming any channel entitlements. Subscribing a guest to any channel that does not belong to **virtualization-platform-free**, such as a Directory Server or RHN Satellite channel, consumes an additional channel entitlement.



Tip

The two virtualization entitlements specifically apply to host systems.

Guest systems that exist on unregistered hosts are treated the same as any physical system — each guest consumes a channel and a system entitlement.

7.3.4.1. System Entitlements

The **System Entitlements** page allows you to view, add, and remove the entitlements for your registered systems. Red Hat Network Satellite allows you to apply and remove entitlements at will, allowing you to adjust your Red Hat Network infrastructure as your organization grows and changes.

To enable the base entitlement, select the checkbox to the left of the system, then click the **Set to Management Entitled** button. For add-on entitlements, select the system's checkbox, followed by the desired entitlement from the drop-down box, and finally press the **Add Entitlement** button.

If clicking on an entitlement fails to update the information in the table, you may need to purchase additional entitlements. Check the number of available subscriptions, listed in bold below the table. Non-RHN Satellite customers may purchase more entitlements; click the **Buy Now** link at the left of the page to do so.

When an entitlement expires, the last system entitled to the same service level (such as Management) will be unentitled. For instance, if you have 10 Red Hat Enterprise Linux AS systems entitled to Management and either one of the RHN entitlements or one of the operating system subscriptions expire, the last system subscribed or entitled will have their subscription or entitlement removed.

7.3.4.2. Virtualization Entitlements

This page only appears if you have applied Virtualization or Virtualization Platform entitlements. It allows you to quickly assess whether you have used these entitlements in the most effective manner.

The first table on this page displays any Virtualization-entitled hosts that have more guest systems than are allowed in the Red Hat Network service level agreement. If you would like to upgrade these systems to any available Virtualization Platform entitlements, click the profile name of that system. This displays the **System Details** page for the system. Click the **Edit Properties** link on the page to edit that system's add-on entitlements.

The second table displays any Virtualization Platform-entitled hosts that have fewer than four guests. It may be advisable to downgrade these systems' entitlements to the Virtualization entitlement. To do so, click the profile name of the system you would like to downgrade, then edit the add-on entitlements from the resulting **System Details** page.

7.3.4.3. Software Channel Entitlements

The software channels listed on this page are the subscription-based channels to which your organization has paid access. The table lists each of the supported operating systems that can be managed via RHN, the number of such systems you have registered with RHN, and finally the remaining number of entitlements for that operating system. Clicking on the name of the channel opens a page that displays information about the channels associated with that channel entitlement. Clicking on the number of entitled systems displays a list of the systems so entitled.

7.3.5. Organization Trusts

The **Organization Trusts** page displays the trusts established with your organization (that is, the organization with which you, the logged-in user, are associated). The page also lists **Channels Shared**: that is channels available to your organisation via others in the established trusts.

You can filter the list of trusts by keyword using the **Filter by Organization** text box and clicking **Go**.

For more information about Organizational Trusts, refer to [Section 9.6, “Organizational Trusts”](#).

7.4. Systems

If you click the **Systems** tab on the top navigation bar, the **Systems** category and links appear. The pages in the **Systems** category allow you to select systems so that you can perform actions on them and create System Profiles.

7.4.1. Overview —

The **Overview** page provides a summary of your systems, including their status, number of associated Errata and packages, and entitlement level. Clicking on the name of a system takes you to its **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.

Clicking the **View System Groups** link at the top of the **Overview** page takes you to a similar summary of your system groups. It identifies group status and displays the number of systems contained. Clicking on the number of systems takes you to the **Systems** tab of the **System Group Details** page, while clicking on the system name takes you to the **Details** tab for that system. Refer to

[Section 7.4.3.3, “System Group Details — !\[\]\(d0262bbe9d2356661a2e89321dfcc781_img.jpg\) ”](#) for more information.

You can also click the **Use Group** button in the **System Groups** section of the **Overview** page to go directly to the **System Set Manager**. Refer to [Section 7.4.4, “System Set Manager — !\[\]\(51514032c8ca341817228f39f1307b05_img.jpg\) ”](#) for more information.

7.4.2. Systems








The **Systems** page displays a list of all of your registered systems. The **Systems** list contains several columns of information for each system:

- **Select** — Update or unentitled systems cannot be selected. To select systems, mark the appropriate checkboxes. Selected systems are added to the **System Set Manager**. After adding systems to the **System Set Manager**, you can use it to perform actions on them simultaneously. Refer to

[Section 7.4.4, “System Set Manager — !\[\]\(f219cfc00b8db0cd1a81ae1fc9afaf28_img.jpg\) ”](#) for details.

- **Status** — Shows which type of Errata Alerts are applicable to the system or confirms that it is up-to-date. Some icons are linked to pages providing resolution. For instance, the standard Updates icon is linked to the **Upgrade** subtab of the packages list, while the Critical Updates icon links directly to the **Update Confirmation** page. Also, the Not Checking In icon is linked to instructions for resolving the issue.

-  — System is up-to-date

-  — Critical Errata available, update *strongly* recommended
 -  — Updates available and recommended
 -  — System is locked; Actions prohibited
 -  — System is being kickstarted
 -  — Updates have been scheduled
 -  — System not checking in properly (for 24 hours or more)
 -  — System not entitled to any update service
- **Errata** — Total number of Errata Alerts applicable to the system.
 - **Packages** — Total number of package updates for the system. Includes packages from Errata Alerts as well as newer packages that are not from Errata Alerts. For example, imagine a client system that has an early version of a package installed. If this client is then subscribed to the appropriate base channel of RHN (such as Red Hat Enterprise Linux 5), that channel may have an updated version of the package. If so, the package appears in the list of available package updates.



Important

If the RHN website identifies package updates for the system, yet the **Red Hat Update Agent** responds with "Your system is fully updated" when run, a conflict likely exists in the system's package profile or in the **up2date** configuration file. To resolve the conflict, either schedule a package list update or remove the packages from the Package Exceptions list for the **Red Hat Update Agent**. Refer to *Section 7.4.2.9, "System Details"* or *Section 4.4.1.3, "Package Exceptions Settings"*, respectively, for instructions.

- **System** — The name of the system as configured when registering it. The default name is the hostname of the system. Clicking on the name of a system takes you to the **System Details** page for the system. Refer to *Section 7.4.2.9, "System Details"* for more information.
- **Base Channel** — The primary channel for the system, based upon its operating system distribution. Refer to *Section 7.6.1, "Software Channels"* for more information.
- **Entitlement** — Whether or not the system is entitled and at what service level.

Links in the left navigation bar below **Systems** enable you to select and view predefined sets of your systems. All of the options described above can be applied within these pages.

7.4.2.1. All

The **All** page contains the default set of your systems. It displays every system you have permission to manage. A user has permission to manage a system if he is the only user in his organization, if he is an Satellite Administrator, or if the system is a member of a group to which he has admin rights.

7.4.2.2. Virtual Systems

To reach this page, select the **Systems** tab, followed by the **Systems** subtab from the left navigation bar, and finally select **Virtual Systems** from the left navigation bar. This page lists each virtual host of which the RHN Satellite is aware and the guest systems on those hosts.

System

This column displays the name of each guest system.

Updates

This column indicates whether the guest systems have any errata that have not yet been applied to them.

Status

This column indicates whether a guest is running, paused, or stopped.

Base Channel

This column indicates the base channel to which the guest is currently subscribed.

Only those guests that are registered with RHN are displayed in blue text. Clicking on the hostname of such a guest system displays that system's **System Details** page.

7.4.2.3. Out of Date

The **Out of Date** page displays the systems that have applicable Errata Alerts that have not been applied.

7.4.2.4. Unentitled —

The **Unentitled** page displays the systems that have not yet been entitled for Red Hat Network service.

7.4.2.5. Ungrouped

The **Ungrouped** page displays the systems that have not yet been assigned to a specific system group.

7.4.2.6. Inactive

The **Inactive** page displays the systems that have not checked into RHN for 24 hours or more. When the **Red Hat Update Agent** connects to RHN to see if there are any updates available or if any actions have been scheduled, this is considered a check-in. If you are seeing a message indicating checkins are not taking place, the RHN client on your system is not successfully reaching Red Hat Network for some reason. This indicates:

- The system is not entitled to any RHN service. System Profiles that remain unentitled for 180 days (6 months) are removed.

- The system is entitled, but the Red Hat Network Daemon has been disabled on the system. Refer to [Chapter 5, Red Hat Network Daemon](#) for instructions on restarting and troubleshooting.
- The system is behind a firewall that does not allow connections over https (port 443).
- The system is behind an HTTP proxy server that has not been properly configured.
- The system is connected to an RHN Proxy Server or RHN Satellite that has not been properly configured.
- The system itself has not been properly configured, perhaps pointing at the wrong RHN Server.
- The system is not on the network.
- Some other barrier exists between the system and the RHN Servers.

7.4.2.7. Recently Registered

The **Recently Registered** page displays any new systems that have been registered in a given period of time. Use the drop-down menu to specify new systems registered in days, weeks, 30- and 180-day increments, and yearly.

7.4.2.8. Proxy

The **Proxy** page displays the RHN Proxy Server systems registered to your RHN account.

7.4.2.9. System Details

Click on the name of a system on any page and RHN displays the **System Details** page for that client. From here, you may modify the displayed information or remove the system altogether by clicking the **delete system** link on the top-right corner.





Note

The **delete system** link in the upper right of this screen refers to the system profile only. Deleting a host system profile will not destroy or remove the registration of guest systems. Deleting a guest system profile does not remove it from the list of guests for its host, nor does it stop or pause the guest. It does, however, remove your ability to manage it via RHN.

If you mistakenly delete a system profile from RHN, you may re-register the system.

The **System Details** page is further divided into the following tabs:

- Details
- Software
- Configuration
- Provisioning — 

- Monitoring — 
- Groups
- Events

The following sections discuss these tabs and their sub-tabs in detail.

7.4.2.9.1. System Details ▢ Details

This page is not accessible from any of the standard navigation bars. However, clicking on the name of a system anywhere in the web interface brings you to this page. The default tab displayed on this page is the **Details ▢ Overview** subtab. Other tabs are available, depending on the current entitlement level of the system.

7.4.2.9.1.1. System Details ▢ Details ▢ Overview

This system summary page displays the system status message and the following key information about the system:

System Info

System Status Message

This message indicates the current state of your system in relation to RHN.



Note

If updates are available for any entitled system, the message **Critical updates available** appears. To apply these updates, click the **update now** link.

system ID

A unique identifier generated each time a system registers with RHN.



Note

The system ID can be used to eliminate duplicate profiles from RHN. Compare the system ID listed on this page with the information stored on the client system in the `/etc/sysconfig/rhn/systemid` file. In that file, the system's current ID is listed under "system_id". The value starts after the characters "ID-". If the value stored in the file does not match the value listed in the profile, the profile is not the most recent one and may be removed.

Hostname

The hostname as defined by the client system. This information is often found in `/etc/hostname` for Red Hat Enterprise Linux systems.

IP Address

The IP address of the client.

Kernel

The kernel that is installed and operating on the client system.

Registered

The date and time at which the system registered with RHN and created this profile.

Checked In

The date and time at which the system last checked in with RHN.

Last Booted

The date and time at which the system was last started or restarted.



Note

Systems with a Management entitlement can be rebooted from this screen.

- Select **Schedule system reboot**
- Provide the earliest date and time at which the reboot may take place.
- Click the **Schedule Reboot** button in the lower right.

When the client checks in after the scheduled start time, RHN will instruct the system to restart itself.

Locked

Indicates whether a system has been locked.

Actions cannot be scheduled for locked systems through the web interface until the lock is removed manually. This does not include preventing auto-errata updates scheduled through the web interface. To prevent the application of auto-errata updates, de-select **Auto Errata Update** from the **System Details** ▢ **Details** ▢ **Properties** subtab.


Locking a system can help to prevent you from accidentally making any changes to a system until you are ready to do so. For example, the system may be a production system that you do not wish to receive updates or new packages until you decide to unlock it.



Important

Locking a system in the web interface *will not* prevent any actions that originate from the client system. For example, if a user logs into the client directly and runs **up2date**, **up2date** will install available errata whether or not the system is locked in the web interface.

Further, locking a system *does not* restrict the number of users who can access the system via the web interface. If you wish to restrict access to the system, associate that system with a System Group and assign it a System Group Administrator. Refer

to [Section 7.4.3, "System Groups](#) —  " for more information about System Groups.

It is also possible to lock multiple systems via the System Set Manager. Refer to

[Section 7.4.4.12.4, "System Set Manager ▢ Misc ▢ Lock Systems](#) —  " to learn how to do so.



— OSA status is also displayed for client systems registered to a Satellite that have a Provisioning entitlement and have enabled OSA.

Push enables Satellite customers to immediately initiate tasks on Provisioning-entitled system rather than wait for those systems to check in with RHN. Scheduling actions through push is identical to the process of scheduling any other action, except that the task begins immediately instead of waiting the set interval.

In addition to the configuration of the Satellite, each client system to receive pushed actions must have the **osad** package installed and its service started. Refer to the *Enabling Push to Clients* section of the *RHN Satellite 5.2.0 Installation Guide* for details.

Subscribed Channels

Base Channel

The first line indicates the base channel to which this client is subscribed. The base channel should match the operating system of the system.

Child Channels

The subsequent lines of text, which depend from the base channel, are child channels. Examples are the **Red Hat Network Tools** channel and the **RHEL AS Extras** channel.



Note

The final link under **Subscribed Channels** is the **Alter Channel subscriptions** link. Click on this link to select from the available base and child channels for this system. When finished making selections, click the **Change Subscriptions** button to confirm the changes.

System Properties

Profile Name

This editable name for the system profile is set to the system's hostname by default. It serves to distinguish this system profile from others.

Entitlement

The base entitlement currently applied to this system.

Notifications

Indicates the the notification options for this system. You can choose whether you wish to receive email notifying you of available errata updates for this system. In addition, you may choose to include Management-entitled systems in the daily summary email.

Auto Errata Update

Indicates whether this system is configured to accept updates automatically.

Description

This information is automatically generated at registration. You can edit this to include any information you wish.

Location

If entered, this field displays the physical address of the system.

The final link on the page is **Edit these properties**. Clicking this link opens the **System Details** ▢ **Properties** subtab. On this page, edit any text you choose, then click the **Update Properties** button to confirm.

7.4.2.9.1.2. System Details ▢ Details ▢ Properties

This subtab allows you to alter the following basic properties of your system:

Profile Name

By default, this is the hostname of the system. You can however alter the profile name to anything that allows you to distinguish this profile from others.

Base Entitlement

Select a base channel for the system from the available base entitlements.

Add-on entitlements

If available, apply a Monitoring, Provisioning, Virtualization, or Virtualization Platform entitlement to the system.

Notifications

Toggle whether notifications about this system are sent and whether this system is included in the daily summary. (By default, all Management and Provisioning systems are included in the summary.) This setting keeps you abreast of all advisories pertaining to the system. Anytime an update is produced and released for the system, a notification is sent via email.

The daily summary reports system events that affect packages, such as scheduled Errata Updates, system reboots, or failures to check in. In addition to including the system here, you must choose to receive email notification in the **Your Preferences** page of the **Overview** category.

Auto-errata update

If this box is checked, available errata are automatically applied to the system when it checks in. This action takes place without user intervention. Customers should note that Red Hat does not recommend the use of the auto-update feature for production systems because conflicts between packages and environments can cause system failures. The Red Hat Network Daemon must be enabled on the system for this feature to work.


Description

By default, this text box records the operating system, release, and architecture of the system when it first registers. You may edit this information to include anything you like.

The remaining fields record the physical address at which the system is stored. To confirm any changes to these fields, click the **Update Properties** button.



Note

Many of these properties can be set for multiple systems at once through the System Set Manager interface. Refer to *Section 7.4.4, “System Set Manager* —  *” for details.*

7.4.2.9.1.3. System Details ▢ Details ▢ Remote Command —

This subtab allows you to run a remote command on the system if the system possesses a Provisioning entitlement. Before doing so, you must first configure the system to accept such commands.

- First, subscribe the system to the RHN Tools channel and use **up2date** to install the **rhncfg**, **rhncfg-client**, and **rhncfg-actions** packages.

```
up2date rhncfg rhncfg-client rhncfg-actions
```

- Log into the system as root and add the following file to the local RHN configuration directory: **allowed-actions/scripts/run**.
 - Create the necessary directory on the target system:

```
mkdir -p /etc/sysconfig/rhn/allowed-actions/script
```

- Create an empty **run** file in that directory to act as a flag to RHN signaling permission to allow remote commands:

```
touch /etc/sysconfig/rhn/allowed-actions/script/run
```

Once the setup is complete, refresh the page in order to view the text fields for remote commands. You may then identify a specific user, group, and timeout period, as well as the script itself on this page. Select a date and time to begin attempting the command, and click **Schedule Remote Command**.

7.4.2.9.1.4. System Details ▢ Details ▢ Reactivation —

An activation key specific to this System Profile. Reactivation keys, available only for systems that have a Provisioning entitlement, include this system's ID, history, groups, and channels. This key can then be used only once with the **rhnreg_ks** command line utility to re-register this system and regain all Red Hat Network settings. Refer to [Section 4.5, "Registering with Activation Keys"](#) for instructions. Unlike typical activation keys, which are not associated with a specific system ID, keys created here do not show up within the **Activation Keys** page.

Reactivation keys can be combined with activation keys to aggregate the settings of multiple keys for a single system profile. For example:

```
rhnreg_ks --server=<server-url> --activationkey=<reactivation-key>,<activationkey> --force
```



Warning

When kickstarting a system with its existing RHN profile, the kickstart profile uses the system-specific activation key created here to re-register the system and return its other

RHN settings. For this reason, you should not regenerate, delete, or use this key (with `rhnreg_ks`) while a profile-based kickstart is in progress. If you do, the kickstart will fail.

7.4.2.9.1.5. System Details ▢ Details ▢ Hardware

This subtab provides detailed information about the system, including networking, BIOS, storage, and other devices. This appears only if you selected to include the hardware profile for this machine during registration. If the hardware profile looks incomplete or outdated, click the **Schedule Hardware Refresh** button to schedule a Hardware Profile update for your system. The next time the RHN Daemon connects to RHN, it will update your System Profile with the latest list of hardware.

7.4.2.9.1.6. System Details ▢ Details ▢ Notes

This subtab provides a place to create notes about the system. To add a new note, click the **create new note** link, type a subject and details, and click the **Create** button. To modify a note, click on its subject in the list of notes, make your changes, and click the **Update** button. To remove a note, click on its subject in the list of notes and then click the **delete note** link.

7.4.2.9.1.7. System Details ▢ Details ▢ Custom Info —

This subtab, available for systems with a Provisioning entitlement, provides completely customizable information about the system. Unlike **Notes**, **Custom Info** is structured, formalized, and can be searched upon. Before you can provide custom information about a system, you must first have **Custom Information Keys**. This is done via the **Custom System Info** page, available from the left navigation bar. Refer to [Section 7.4.8, “Custom System Info — !\[\]\(17acf1afa8cdf0b67c53d4865a5ed469_img.jpg\)”](#) for instructions.

Once you have created one or more Keys, you may assign a value for this system by select the **create new value** link. Click the name of the key in the resulting list and enter a value for it in the **Description** field, then click the **Update Key** button.

7.4.2.9.1.8. System Details ▢ Details ▢ Proxy

Activates an RHN Proxy Server. This tab is only available for Provisioning-entitled systems. Select a version of RHN Proxy Server and click the **Activate Proxy** button to begin the installation and activation process. For detailed information, refer to the *RHN Proxy Server Guide* and the *Client Configuration Guide*.

7.4.2.9.1.9. System Details ▢ Details ▢ Satellite

Displays the certificate of an active Red Hat Network. You can deactivate an old certificate here and upload a new one if necessary. This tab requires a Provisioning entitlement. For detailed information on activating a Satellite, refer to the *RHN Satellite Installation Guide*.

7.4.2.9.2. System Details ▢ Software

This tab and its accompanying subtabs allow you to manage the software of the system: errata, packages and package profiles, and software channel memberships.

7.4.2.9.2.1. System Details ▢ Software ▢ Errata

This subtab contains a list of Errata Alerts applicable to the system. Refer to [Section 7.1.3, “Errata Alert Icons”](#) for meanings of the icons on this tab. To apply updates, select them and click the **Apply Errata** button. Double-check the updates to be applied on the confirmation page, then click the **Confirm** button. After confirming, the action is added to the **Pending Actions** list under **Schedule**. Errata that have been scheduled cannot be selected for update. In the place of a checkbox is a clock icon that, when clicked, takes you to the **Action Details** page.

To help users determine whether an update has been scheduled, a **Status** column exists within the Errata table. Possible values are: None, Pending, Picked Up, Completed, and Failed. This column identifies only the latest action related to an Erratum. For instance, if an action fails and you reschedule it, this column shows the status of the Erratum as Pending only (with no mention of the previous failure). Clicking a status other than None takes you to the **Action Details** page. This column corresponds to the one on the **Affected Systems** tab of the **Errata Details** page.

7.4.2.9.2.2. System Details ▢ Software ▢ Packages

This subtab allows you to manage the packages on the system.



— When selecting packages to install, upgrade, or remove, Provisioning customers have the option of running a remote command automatically before or after the package installation. Refer to

[Section 7.4.2.9.1.3, “System Details ▢ Details ▢ Remote Command — !\[\]\(17413706fd4997a1a4bdf85c6864eee1_img.jpg\) ”](#) for more information.

Packages

The default display of the **Packages** tab describes the options available to you and provides the means to update your package list. To update or complete a potentially outdated list, possibly due to the manual installation of packages, click the **Update Package List** button on the bottom right-hand corner of this page. The next time the RHN Daemon connects to RHN, it updates your System Profile with the latest list of installed packages.

List/Remove

Lists installed packages from the system's software System Profile and enables you to remove them. Click on a package name to view its **Package Details** page. To delete packages from the system, select their checkboxes and click the **Remove Packages** button on the bottom right-hand corner of the page. A confirmation page appears with the packages listed. Click the **Confirm** button to remove the packages.

Upgrade

Displays a list of packages that have a new version available based on the package versions in the channels for the system. Click on the latest package name to view its **Package Details** page. To upgrade packages immediately, select them and click the **Upgrade Packages** button. To download the packages as a .tar file, select them and click the **Download Packages** button.

Install

Enables you to install new packages on the system from the available channels. Click on the package name to view its **Package Details** page. To install packages, select them and click the **Install Selected Packages** button.

Verify

Validates the packages installed on the system against its RPM database. This is the equivalent of running `rpm -V`. Specifically, this tab allows you to compare the metadata of the system's

packages with information from the database, such as MD5 sum, file size, permissions, owner, group and type. To verify a package or packages, select them, click the **Verify Selected Packages** button, and confirm this action. Once finished, you can view the results by selecting this action within the **History** subtab under **Events**.

Profiles

Gives you the ability to compare the packages on this system with the packages of stored profiles and other Management and Provisioning systems. To make the comparison with a stored profile, select that profile from the pulldown menu and click the **Compare** button. To make the comparison with another system, select it from the associated pulldown menu and click the **Compare** button. To create a stored profile based upon the existing system, click the **Create System Profile** button, enter any additional information you desire, and click the **Create Profile** button. These profiles are kept within the **Stored Profiles** page linked from the left navigation bar.



— Once package profiles have been compared, Provisioning customers have the ability to synchronize the packages of the selected system with the package manifest of the compared profile. Note that this action may delete packages on the system not in the profile, as well as install packages from the profile. To install specific packages, select the checkboxes of packages from the profile. To remove specific packages already installed on the system itself, select the checkboxes of packages showing a difference of **This system only**. To synchronize fully the system's packages with the compared profile, select the master checkbox at the top of the column. Then click the **Sync Packages to** button. On the confirmation screen, review the changes, select a time frame for the action, and click the **Schedule Sync** button.

7.4.2.9.2.3. System Details ▯ Software ▯ Software Channels

Software channels provide a well-defined method to determine which packages should be available to a system for installation or upgrade based upon its operating systems, packages, and functionality. Click a channel name to view its **Channel Details** page. To modify the child channels associated with this system, use the checkboxes next to the channels and click the **Change Subscriptions** button. You will receive a success message or be notified of any errors. To change the system's base channel, select the new one from the pulldown menu and click the **Modify Base Channel** button. Refer to [Section 7.6.1, "Software Channels"](#) for more information.

7.4.2.9.3. System Details ▯ Configuration —

This tab and its subtabs, which do not appear without a Provisioning entitlement, assist in managing the configuration files associated with the system. These configuration files may be managed solely for the current system, or may be distributed widely via a Configuration Channel. The following section describe these and other available options on the **System Details ▯ Configuration** subtabs.



Note

To manage the configuration of a system, it must have the latest **rhncfg*** packages installed. Refer to [Section 7.7.1, "Preparing Systems for Config Management"](#) for instructions on enabling and disabling scheduled actions for a system.

This section is available to normal users with access to systems that have configuration management enabled. Like software channels, configuration channels store files to be installed on systems. While software updates are provided by RHN, configuration files are managed solely by you. Also unlike

software packages, various versions of configuration files may prove useful to a system at any given time. Remember, only the latest version can be deployed.

7.4.2.9.3.1. System Details ▯ Configuration ▯ Overview

This subtab provides access to the configuration statistics of your system and to the most common tasks used to manage configuration files. You may change the settings listed under Configuration Stats by clicking on the blue text for that setting. Alternatively, you may perform any of the common configuration management tasks listed on the right of the screen by clicking one of the links.

7.4.2.9.3.2. System Details ▯ Configuration ▯ Managed Files

This subtab lists all configuration files currently associated with the system.

Filename

This column shows both the name and the deployment path for this file.

Revision

This column increments any time you make a change to the managed file.

From Config Channel

This column indicates the name of the channel that contains the file, or displays **(system override)** for files available to this system only.

Overrides

If this configuration file overrides another, the overridden file is listed in this column along with its host channel.

If you wish to deploy any of these files to the client system, overwriting any changes that have been made locally, check the box to the left of the file and click the **Deploy Configuration** button. On the following screen, choose a deployment time and click the **Schedule Deploy** button to confirm.



Note

If you click on the **Filename** of a **(system override)** file, you can edit its contents.

The **Overrides** column identifies the configuration file in an unsubscribed channel that would replace the same file in a currently subscribed channel. For example, if a system has '/etc/foo' from channel 'bar' and '/etc/foo' from channel 'baz' is in the Overrides column, then unsubscribing from channel 'bar' will mean that the file from channel 'baz' will be applicable. Also, if nothing is in the 'Overrides' column for a given file path, then unsubscribing from the channel providing the file will mean that the file is no longer managed (though it will *not* remove the file from the system).

7.4.2.9.3.3. System Details ▯ Configuration ▯ Compare Files

This subtab compares a configuration file as stored on the Satellite with the file as it exists on the client. (It does not, for example, compare versions of the same file stored in different channels.) Select the files to be diffed, click the **Compare Files** button, select a time to perform the diff, and click the **Schedule Compare** button to confirm. After the diff has been performed, you may return to this page to view the results.

7.4.2.9.3.4. System Details ▯ Configuration ▯ Manage Configuration Channels

This subtab allows you to subscribe to and rank configuration channels that may be associated with the system, lowest first.

The **List/Unsubscribe from Channels** subtab contains a list of the system's configuration channel subscriptions. Click the checkbox next to the Channel and click **Unsubscribe** to remove the subscription to the channel.

The **Subscribe to Channels** subtab lists all available configuration channels. To subscribe to a channel, select the checkbox next to it and press **Continue**. To subscribe to all configuration channels, click **Select All** and press **Continue**. The **View/Modify Rankings** page automatically loads.

The **View/Modify Rankings** subtab allows users rank the priority in which files from a particular configuration channel are weighted. The higher the channel is on the list, the more its files take precedence over files on lower-ranked channels (for example, the higher-ranked channel may have an `httpd.conf` file that will take precedence over the file on lower-ranked channel)

7.4.2.9.3.5. System Details ▯ Configuration ▯ Local Overrides

This subtab displays the default configuration files for the system and allows you to manage them. If no files exist, you may use the **add files**, **upload files**, and **add directories** links within the page description to associate files with this system. These tabs correspond to those within the **Configuration Channel Details** page, affecting your entire organization and available only to Configuration Administrators. Refer to [Section 7.7.3.1, "Configuration ▯ Configuration Channels ▯ Configuration Channel Details"](#) for more information.

If a file exists, click its name to go to the **Configuration File Details** page. Refer to [Section 7.7.4, "Configuration Files"](#) for instructions. To replicate the file within a config channel, select its checkbox, click the **Copy to Config Channel** button, and select the destination channel. To remove a file, select it and click **Delete Selected Files**.

7.4.2.9.3.6. System Details ▯ Configuration ▯ Sandbox

This subtab allows you to manipulate configuration files without deploying them. This sandbox provides you with an area in which to experiment with files without affecting your systems. To add files, click the **import new files** link, enter the path to the file on your local system, and click the **Add** button. Select the **Import Files** button to confirm.

7.4.2.9.4. System Details ▯ Provisioning —

This tab and its subtabs allow you to schedule and monitor kickstarts and to return your system to a previous state. Kickstart is a Red Hat utility that allows you to automate the reinstallation of a system. Snapshots keep a record of every change to a Provisioning system and allow you to "undo" those changes at will. Both features are described in the sections that follow.

7.4.2.9.4.1. System Details ▯ Provisioning ▯ Kickstart —

This subtab is further divided into **Session Status**, which tracks the progress of previously scheduled kickstarts, and **Schedule**, which allows you to configure and schedule a kickstart for this system.

The **Schedule** subtab allows you to schedule the selected system for kickstart. Choose from the list of available kickstart profiles, select a time for the kickstart to begin, and click the **Schedule Kickstart**

and **Finish** button to begin the kickstart. You may first alter kickstart settings by clicking the **Advanced Configuration** button.



Note

You must first create a kickstart profile before it appears on this subtab. If you have not created any profiles, refer to [Section 7.4.9.3, “Create a New Kickstart Profile”](#) before scheduling a kickstart for a system.

The **Variables** subtab can be used to create Kickstart variables, which substitute values into kickstart files. To define a variable, create a name-value pair (*name/value*) in the text box.

For example, if you wanted to kickstart a system that joins the network for specified department (for example the Engineering organization) you can create a profile variable to set the IP address and the gateway server address to a variable that any system using that profile will use. Add the following line to the **Variables** text box.

```
IPADDR=192.168.0.28
GATEWAY=192.168.0.1
```

To use the system variable, you can use the name of the variable within the profile to substitute in the value. For example, the **network** portion of a kickstart file could look like the following:

```
network --bootproto=static --device=eth0 --onboot=on --ip=$IPADDR
--gateway=$GATEWAY
```

The **\$IPADDR** will be **192.168.0.28**, and the **\$GATEWAY** will be **192.168.0.1**



Note

There is a hierarchy when creating and using variables in kickstart files. System kickstart variables take precedence over Profile variables, which in turn take precedence over Distribution variables. Understanding this hierarchy can alleviate confusion when using variables in kickstarts.

Using variables are just one part of the larger Cobbler infrastructure for creating templates that can be shared between multiple profiles and systems. For more information about Cobbler and kickstart templates, refer to [Chapter 11, Cobbler](#).

7.4.2.9.4.2. System Details ▯ Provisioning ▯ Snapshots —

Snapshots enable you to roll back the system's package profile, configuration files, and RHN settings. Snapshots are captured whenever an action takes place on a Provisioning-entitled system. The **Snapshots** subtab lists all snapshots for the system, including the reason the snapshot was taken, the time it was taken, and the number of tags applied to each snapshot. To revert to a previous configuration, click the **Reason** of the snapshot taken and review the potential changes on the provided subtabs, starting with **Rollback**.

Each subtab provides the specific changes that will be made to the system during the rollback:

- group memberships
- channel subscriptions
- installed packages
- configuration channel subscriptions
- configuration files
- snapshot tags

When satisfied with the reversion, return to the **Rollback** subtab and click the **Rollback to Snapshot** button. To see the list again, click **Return to snapshot list**.

7.4.2.9.4.3. System Details ▯ Provisioning ▯ Snapshot Tags —

Provides a means to add meaningful descriptions to your most recent system snapshot. This can be used to indicate milestones, such as a known working configuration or a successful upgrade. To tag the most recent snapshot, click **create new system tag**, enter a descriptive term in the **Tag name** field, and click the **Tag Current Snapshot** button. You may then revert using this tag directly by clicking its name in the Snapshot Tags list. To delete tags, select their checkboxes, click **Remove Tags**, and confirm the action.

7.4.2.9.5. System Details ▯ Virtualization

This is tab allows you to create a new virtual guest on a host system or allows you to change the status of virtual guests.

The **Virtualization** tab has two subtabs, **Details** and **Kickstart**. These tabs appear the same for both virtual hosts and guests, but the functionality only makes sense for virtual hosts. It is not possible to create a guest system that runs on another guest system.

7.4.2.9.5.1. System Details ▯ Virtualization ▯ Details

Details is the default tab. For host systems, it presents a table of the host system's virtual guests. For each guest system, the following information is provided:

Status

This field indicates whether the virtual system is running, paused, stopped, or has crashed.

Updates

This field indicates whether errata applicable to the guest have yet to be applied.

Base Software Channel

This field indicates the Base Channel to which the guest is subscribed.



Tip

If a guest system has not registered to the Satellite, this information appears as plain text in the table.

If you have System Group Administrator responsibilities assigned for your guest systems, it is possible that a user could see the message **You do not have permission to access this system** within the table. This is because it is possible to assign virtual guests on a single host to multiple System Group Administrators. Only users that have System Group Administrator privileges on the host system may create new virtual guests.

7.4.2.9.5.2. System Details ▯ Monitoring —

This tab is only visible for systems registered to a RHN Satellite with Monitoring enabled and that are Monitoring entitled. It displays all of the probes monitoring the system. The **State** column shows icons

representing the status of each probe. Refer to [Section 7.10, “Monitoring — !\[\]\(9dfdaff1d86ba3c1f8353b4d1b61b8c5_img.jpg\) ”](#) for descriptions of these states. Clicking the **Probe Description** takes you to its **Current State** page. The **Status String** column displays the last message received from the probe.

To add a probe to the system, click the **create new probe** link at the top-right corner of the page and complete the fields on the following page. Refer to [Section 8.5.1, “Managing Probes”](#) for detailed instructions.

Once the probe has been added, you must reconfigure your Monitoring infrastructure to recognize it.

Refer to [Section 7.10.4, “Scout Config Push — !\[\]\(642aa997563f9a325b310230bb5078b7_img.jpg\) ”](#) for details. After the probe has run, its results

become available on the **Current State** page. Refer to [Section 7.10.1.7, “Current State — !\[\]\(2b376d1a92330ab09dad2665d2f89bf5_img.jpg\) ”](#) for details.

To remove a probe from a system, click on the name of the probe, then click the **delete probe** link in the upper right corner. Finally, click the **Delete Probe** button to complete the process.

7.4.2.9.5.3. System Details ▯ Groups —

This tab and its subtabs allow you to manage the system's group memberships.

7.4.2.9.5.3.1. System Details ▯ Groups ▯ List/Leave —

This subtab lists groups to which the system belongs and enables you to cancel those associations. Only System Group Administrators and Satellite Administrators can remove the system from groups. Non-admins just see a **Review this system's group membership** page. To remove the system from groups, select the groups' checkboxes and click the **Leave Selected Groups** button. Click on a group's name to go to its **System Group Details** page. Refer to [Section 7.4.3.3, “System Group](#)

[Details — !\[\]\(274fd520e03b61c1b9ffc861754cacdc_img.jpg\) ”](#) for more information.

7.4.2.9.5.3.2. System Details ▯ Groups ▯ Join —

Lists groups that the system may be subscribed to. Only System Group Administrators and Satellite Administrators can add the system to groups. Non-admins see a **Review this system's group membership** page. To add the system to groups, select the groups' checkboxes and click the **Join Selected Groups** button.





7.4.2.9.5.4. System Details ▯ Events

Displays past, current, and scheduled actions on the system. You may cancel pending events here. The following sections describe the **Events** sub-tabs and the features they offer.

7.4.2.9.5.4.1. System Details ▯ Events ▯ Pending

Lists events that are scheduled but have not begun. A prerequisite action must complete successfully before a given action is attempted. If an action has a prerequisite, no checkbox is available to cancel that action. Instead, a checkbox appears next to the prerequisite action; canceling the prerequisite action causes the action in question to fail.

Actions can be chained in this manner so that action 'a' requires action 'b' which requires action 'c'. Action 'c' is the first one attempted and has a checkbox next to it until it is completed successfully - if any action in the chain fails, the remaining actions also fail. To unschedule a pending event, select the event and click the **Cancel Events** button at the bottom of the page. The following icons indicate the type of events listed here:

-  — Package Event
-  — Errata Event
-  — Preferences Event
-  — System Event

7.4.2.9.5.4.2. System Details ▯ Events ▯ History

The default display of the **Events** tab lists the type and status of events that have failed, occurred or are occurring. To view details of an event, click its summary in the **System History** list. To again view the table, click **Return to history list** at the bottom of the page.

7.4.3. System Groups —

The **System Groups** page allows all RHN Management and Provisioning users to view the **System Groups** list. Only System Group Administrators and Satellite Administrators may perform the following additional tasks:

1. Create system groups. (Refer to [Section 7.4.3.1, “Creating Groups”](#).)
2. Add systems to system groups. (Refer to [Section 7.4.3.2, “Adding and Removing Systems in Groups”](#).)
3. Remove systems from system groups. (Refer to [Section 7.4.2.9, “System Details”](#).)
4. Assign system group permissions to users. (Refer to [Section 7.9, “Users — !\[\]\(564903337f30b845a5f6979939a95fe6_img.jpg\) ”](#).)

The **System Groups** list displays all of your system groups.




The **System Groups** list contains several columns for each group:

- **Select** — These checkboxes enable you to add systems in groups to the **System Set Manager**. To select groups, mark the appropriate checkboxes and click the **Update** button below the column. All systems in the selected groups are added to the **System Set Manager**. You can then use the **System Set Manager** to perform actions on them simultaneously. It is possible to select only those systems that are members of all of the selected groups, excluding those systems that belong only to one or some of the selected groups. To do so, select them and click the **Work with Intersection** button. To add all systems in all selected groups, select them and click the **Work with Union** button. Each system will show up once, regardless of the number of groups to which it belongs. Refer to

[Section 7.4.4, “System Set Manager — !\[\]\(919a2cb85b99741a73c0c31a427236a8_img.jpg\) ”](#) for details.

- **Updates** — Shows which type of Errata Alerts are applicable to the group or confirms that it is up-to-date. Clicking on a group's status icon takes you to the **Errata** tab of its **System Group Details** page. Refer to [Section 7.4.3.3, “System Group Details — !\[\]\(38441ceaa711016e0bf2ad46ad394ff4_img.jpg\) ”](#) for more information.

The status icons call for differing degrees of attention:

-  — All systems within group are up-to-date
-  — Critical Errata available, update *strongly* recommended
-  — Updates available and recommended
- **Group Name** — The name of the group as configured during its creation. The name should be explicit enough to easily differentiate between it and other groups. Clicking on the name of a group takes you to **Details** tab of its **System Group Details** page. Refer to [Section 7.4.3.3, “System Group Details — !\[\]\(8a17676a8da87a4e59299223a765e613_img.jpg\) ”](#) for more information.
- **Systems** — Total number of systems contained by the group. Clicking on the number takes you to the **Systems** tab of the **System Group Details** page for the group. Refer to [Section 7.4.3.3, “System Group Details — !\[\]\(f7fdc7cc047b770fc5fdd2c2137c07d9_img.jpg\) ”](#) for more information.
- **Use in SSM** — Clicking the **Use Group** button in this column loads the group from that row and launches the **System Set Manager** immediately. Refer to [Section 7.4.4, “System Set Manager — !\[\]\(3ca549f0313858650ddae522dc3cfea6_img.jpg\) ”](#) for more information.

7.4.3.1. Creating Groups

To add a new system group, click the **create new group** link at the top-right corner of the page. Type a name and description and click the **Create Group** button. Make sure you use a name that clearly sets this group apart from others. The new group will appear in the **System Groups** list.

7.4.3.2. Adding and Removing Systems in Groups

Systems can be added and removed from system groups in two places: the **Target Systems** tab of the **System Group Details** page and the **Groups** tab of the **System Details** page. The process is

similar in both instances. Select the systems to be added or removed and click the **Add Systems** or **Remove Systems** button.

7.4.3.3. System Group Details —

At the top of each **System Group Details** page are two links: **work with group** and **delete group**. Clicking **delete group** deletes the System Group and should be used with caution. Clicking **Work with Group** functions similarly to the **Use Group** button from the **System Groups** list in that it loads the group's systems and launches the **System Set Manager** immediately. Refer to [Section 7.4.4, "System](#)

[Set Manager](#) —  " for more information.


The **System Group Details** page is broken down into tabs:

7.4.3.3.1. System Group Details ▯ Details —

Provides the group name and group description. To change this information, click **Edit Group Properties**, make your changes in the appropriate fields, and click the **Modify Details** button.

7.4.3.3.2. System Group Details ▯ Systems —

Lists systems that are members of the system group. Clicking links within the table takes you to corresponding tabs within the **System Details** page for the associated system. To remove systems from the group, select the appropriate checkboxes and click the **Remove from group** button on the bottom of the page. Clicking it does not delete systems from RHN entirely. This is done through the

System Set Manager or **System Details** pages. Refer to [Section 7.4.4, "System Set Manager](#) —  " or [Section 7.4.2.9, "System Details"](#), respectively.

7.4.3.3.3. System Group Details ▯ Target Systems —

Target Systems — Lists all systems in your organization. This tab enables you to add systems to the specified system group. Select the systems using the checkboxes to the left and click the **Add Systems** button on the bottom right-hand corner of the page.

7.4.3.3.4. System Group Details ▯ Errata —

List of relevant Errata for systems in the system group. Clicking the Advisory takes you to the **Details** tab of the **Errata Details** page. (Refer to [Section 7.5.2.2, "Errata Details"](#) for more information.)

Clicking the Affected Systems number lists all of the systems addressed by the Errata. To apply the Errata Updates in this list, select the systems and click the **Apply Errata** button.

7.4.3.3.5. System Group Details ▯ Admins —

List of all organization users that have the ability to manage the system group. Satellite Administrators are clearly identified. System Group Administrators are marked with an asterisk (*). To change the system group's users, select and unselect the appropriate checkboxes and click the **Update** button.

7.4.3.3.6. System Group Details ▯ Probes —

List all probes assigned to systems in the system group. The **State** shows the status of the probe. Click the individual **System** for details on the probe and to make changes to the probe configuration. Click the **Probe** to generate a customizable report on the monitoring.

7.4.4. System Set Manager —

Many actions performed for individual systems through the System Details page may be performed for multiple systems via the System Set Manager, including:

- Apply Errata updates
- Upgrade packages to the most recent versions available
- Add/remove systems to/from system groups
- Subscribe/unsubscribe systems to/from channels
- Update system profiles
- Modify system preferences such as scheduled download and installation of packages
- Kickstart several Provisioning-entitled systems at once
- Set the subscription and rank of configuration channels for Provisioning-entitled systems
- Tag the most recent snapshots of your selected Provisioning-entitled systems
- Revert Provisioning-entitled systems to previous snapshots
- Run remote commands on Provisioning-entitled systems

Before performing actions on multiple systems, select the systems you wish to modify. To do so, click the **List the systems** link, check the boxes to the left of the systems you wish to select, and click the **Update List** button.

You can access the System Set Manager in three ways:

1. Click the **System Set Manager** link in the left gray navigation area.
2. Click the **Use Group** button in the **System Groups** list.
3. Check the **Work with Group** link on the **System Group Details** page.

7.4.4.1. System Set Manager ▯ Overview —

Description of the various options available to you in the remaining tabs.

7.4.4.2. System Set Manager ▯ Systems —

List of systems now selected. To remove systems from this set, select them and click the **Remove** button.

7.4.4.3. System Set Manager ▯ Errata —

List of Errata Updates applicable to the current system set. Click the number in the Systems column to see to which systems in the System Set Manager the given Errata applies. To apply updates, select the Errata and click the **Apply Errata** button.

7.4.4.4. System Set Manager ▯ Packages —

Options to modify packages on the system within the following subtabs (Click the number in the Systems column to see to which systems in the System Set Manager the given package applies):



— When selecting packages to install, upgrade, or remove, Provisioning customers have the option of running a remote command automatically before or after the package installation. Refer to

[Section 7.4.2.9.1.3, “System Details ▯ Details ▯ Remote Command — !\[\]\(ec9132f1d27c8919987d92907322654d_img.jpg\) ”](#) for more information.

7.4.4.4.1. System Set Manager ▯ Packages ▯ Upgrade —

A list of all the packages installed on the selected systems that might be upgraded. Systems must be subscribed to a channel providing the package for the system to be able to upgrade the package. If multiple versions of a package appear, note that only the latest version available to each system is upgraded on that system. Select the packages to be upgraded, then click the **Upgrade Packages** button.

7.4.4.4.2. System Set Manager ▯ Packages ▯ Install —

A list of channels from which you may retrieve packages. This list includes all channels to which systems in the set are subscribed; a package is installed on a system only if the system is subscribed to the channel from which the package originates. Click on the channel name and select the packages from the list. Then click the **Install Packages** button.

7.4.4.4.3. System Set Manager ▯ Packages ▯ Remove —

A list of all the packages installed on the selected systems that might be removed. Multiple versions appear if systems in the System Set Manager have more than one version installed. Select the packages to be deleted, then click the **Remove Packages** button.

7.4.4.5. System Set Manager ▯ Verify

A list of all installed package whose contents, MD5 sum, and other details may be verified. At the next check in, the verify event issues the command `rpm --verify` for the specified package. If there are any discrepancies, they are displayed in the System Details page for each system.

Select the checkbox next to all packages to be verified, then click the **Verify Packages** button. On the next page, select either **Schedule actions ASAP** or choose a date and time for the verification, then click the **Schedule Verifications** button.

7.4.4.6. System Set Manager ▫ Patches

Tools to manage patches to Solaris clients. Patches may be installed or removed via the subtabs.

7.4.4.7. System Set Manager ▫ Patch Clusters

Tools to manage patch clusters for Solaris clients. Patches may be installed or removed via the subtabs.

7.4.4.8. System Set Manager ▫ Groups —

Tools to create groups and manage group membership. These functions are limited to Satellite Administrators and System Group Administrators. To add a new group, click **create new group** on the top-right corner. In the resulting page, type its name and description in the identified fields and click the **Create Group** button. To add or remove the selected systems in any of the system groups, toggle the appropriate radio buttons and click the **Alter Membership** button.

7.4.4.9. System Set Manager ▫ Channels —

Options to manage channel associations through the following subtabs:

7.4.4.9.1. System Set Manager ▫ Channels ▫ Channel Subscriptions —

To subscribe or unsubscribe the selected systems in any of the channels, toggle the appropriate checkboxes and click the **Alter Subscriptions** button. Keep in mind that subscribing to a channel uses a channel entitlement for each system in the selected group. If too few entitlements are available, some systems fail to subscribe. Systems must subscribe to a base channel before subscribing to a child channel.

7.4.4.10. System Set Manager ▫ Configuration —

Like the options within the **System Details ▫ Channels ▫ Configuration** tab, the subtabs here can be used to subscribe the selected systems to configuration channels and deploy and compare the configuration files on the systems. The channels are created in the **Manage Config Channels** interface within the **Channels** category. Refer to [Section 7.7.2, “Overview”](#) for channel creation instructions.

To manage the configuration of a system, install the latest **rhncfg*** packages. Refer to [Section 7.7.1, “Preparing Systems for Config Management”](#) for instructions on enabling and disabling scheduled actions for a system.

7.4.4.10.1. System Set Manager ▫ Configuration ▫ Deploy Files —

Use this subtab to distribute configuration files from your central repository on RHN to each of the selected systems. The table lists the configuration files associated with any of the selected systems. Clicking its system count displays the systems already subscribed to the file.

To subscribe the selected systems to the available configuration files, select the checkbox for each desired file. When done, click **Deploy Configuration** and schedule the action. Note that the files deployed are of the latest version at the time of scheduling and do not account for versions that may appear before the action takes place.

7.4.4.10.2. System Set Manager ▯ Configuration ▯ Compare Files —

Use this subtab to validate configuration files on the selected systems against copies in your central repository on RHN. The table lists the configuration files associated with any of the selected systems. Clicking its system count displays the systems already subscribed to the file.

To compare the configuration files deployed on the systems with those in RHN, select the checkbox for each file to be validated. Then click **Analyze Differences** and schedule the action. Note that the files compared are of the latest version at the time of scheduling and do not account for versions that may appear before the action takes place. Find the results within the main **Schedule** category or within the **System Details ▯ Events** tab.

7.4.4.10.3. System Set Manager ▯ Configuration ▯ Subscribe to Channels —

Subscribe systems to configuration channels according to order of preference. This tab is available only to Satellite Administrators and Configuration Administrators. Enter a number in the **Rank** column to subscribe to a channel. Channels are accessed in the order of their rank, starting from the number 1. Channels not assigned a numeric value are not associated with the selected systems. Your local configuration channel always overrides all other channels. Once you have established the rank of the config channels, you must decide how they are applied to the selected systems.

The three buttons below the channels reflect your options. Clicking **Subscribe with Highest Priority** places all the ranked channels before any other channels to which the selected systems are currently subscribed. Clicking **Subscribe With Lowest Priority** places the ranked channels after those channels to which the selected systems are currently subscribed. Clicking **Replace Existing Subscriptions** removes any existing association and starts cleanly with the ranked channels, leaving every system with the same config channels in the same order.

In the first two cases, if any of the newly ranked config channels is already in a system's existing config channel list, the duplicate channel is removed and replaced according to the new rank, effectively reordering the system's existing channels. When such conflicts exist, you are presented with a confirmation page to ensure the intended action is correct. When the change has taken place, a message appears at the top of the page indicating the update was successful.

7.4.4.10.4. System Set Manager ▯ Configuration ▯ Unsubscribe from Channels —



Administrators may unsubscribe from configuration channels by clicking the checkbox by the name of the channel and clicking **Unsubscribe Systems** button.

7.4.4.10.5. System Set Manager ▯ Configuration ▯ Enable Configuration —

Administrators may enable configuration channel management by clicking the checkbox by the name of the channel and clicking **Enable RHN Configuration Management** button. You can also

schedule the action by clicking the **Schedule package installs for no sooner than** radio button and using the drop-down menus to configure date and time, then clicking **Enable RHN Configuration Management**.

7.4.4.11. System Set Manager ▯ Provisioning —

Options for provisioning systems through the following subtabs:

7.4.4.11.1. System Set Manager ▯ Provisioning ▯ Kickstart —

Use this subtab to re-install Red Hat Enterprise Linux on the selected Provisioning-entitled systems. To schedule kickstarts for these systems, select a distribution, identify the type (IP address or manual), and click **Continue**. Finish choosing from the options available on the subsequent screen. If any of the systems connect to RHN via a RHN Proxy Server, choose either the **Preserve Existing Configuration** radio button or the **Use RHN Proxy** radio button. If you choose to kickstart through a RHN Proxy Server, select from the available Proxies listed in the drop-down box beside the **Use RHN Proxy** radio button. All of the selected systems will kickstart through the selected Proxy. Click the **Schedule Kickstart** button to confirm your selections. When the kickstarts for the selected systems are successfully scheduled, the web interface returns you to the System Set Manager page.

7.4.4.11.2. System Set Manager ▯ Provisioning ▯ Tag Systems —

Use this subtab to add meaningful descriptions to the most recent snapshots of your selected systems. To tag the most recent system snapshots, enter a descriptive term in the **Tag name** field and click the **Tag Current Snapshots** button.

7.4.4.11.3. System Set Manager ▯ Provisioning ▯ Rollback —

Use this subtab to rollback selected Provisioning-entitled systems to previous snapshots marked with a tag. Click the name of the tag, verify the systems to be reverted, and click the **Rollback Systems** button.

7.4.4.11.4. System Set Manager ▯ Provisioning ▯ Remote Command —

Use this subtab to issue remote commands on selected Provisioning-entitled systems. First create a **run** file on the client systems to allow this function to operate. Refer to the description of the **Configuration** subtab of the **Channels** tab for instructions. You may then identify a specific user, group, timeout period, and the script on this page. Select a date and time to perform the command, and click **Schedule Remote Command**.

7.4.4.12. System Set Manager ▯ Misc —

Misc — Update System Profiles and preferences for the system set through the following links:

7.4.4.12.1. System Set Manager ▯ Misc ▯ System Profile Updates —

Click **Update Hardware Profile** followed by the **Confirm Refresh** button to schedule a hardware profile update. Clicking **Update Package Profile**, followed by the **Confirm Refresh** button schedules a package profile update.

7.4.4.12.2. System Set Manager ▯ Misc ▯ Custom System Information —

Click **Set a custom value for selected systems** followed by the name of a key to allow you to provide values for all selected systems. Enter the information and click the **Set Values** button. Click **Remove a custom value from selected systems** followed by the name of a key to allow you to remove values for all selected systems. Click the **Remove Values** button to finalize the deletion.

7.4.4.12.3. System Set Manager ▯ Misc ▯ Reboot Systems —

Select the appropriate systems and click the **Reboot Systems** link to set those systems for reboot. To immediately cancel this action, click the **list of systems** link that appears within the confirmation message at the top of the page, select the systems, and click **Unschedule Action**.

7.4.4.12.4. System Set Manager ▯ Misc ▯ Lock Systems —

Select the appropriate systems and click the **Lock Systems** link to prevent the scheduling of any action through RHN that affects the selected systems. This can be reversed by clicking the **Unlock Systems** link.

7.4.4.12.5. System Set Manager ▯ Misc ▯ Delete Systems —

Click **Delete System Profiles**, then click the **Confirm Deletions** button to remove the selected profiles permanently.

7.4.4.12.6. System Set Manager ▯ Misc ▯ Add or Remove Add-On Entitlements —



Select, via the radio button, whether to **Add**, **Remove**, or make **No Change** in the entitlements of the selected systems. Click the **Change Entitlements** button to confirm your selection.

7.4.4.12.7. System Set Manager ▯ Misc ▯ System Preferences —

Toggle the **Yes** and **No** radio buttons and click the **Change Preferences** button to alter your notification preferences for the selected systems. You may apply these preferences to individual systems through the **Properties** subtab of the **System Details** page. Refer to [Section 7.4.2.9.1.2, “System Details ▯ Details ▯ Properties”](#) for instructions.

- **Receive Notifications of Updates/Errata** — This setting keeps you abreast of all advisories pertaining to your systems. Any time an update is produced and released for a system under your supervision, a notification is sent via email.

- **Include system in Daily Summary** — This setting includes the selected systems in a daily summary of system events. (By default, all Management and Provisioning systems are included in the summary.) These system events are actions that affect packages, such as scheduled Errata Updates, system reboots, or failures to check in. In addition to including the systems here, you must choose to receive email notifications in the **Your Preferences** page of **Your RHN**. Refer to [Section 7.3.2, “Your Preferences”](#) for instructions. Note that RHN sends these summaries only to verified email addresses.
- **Automatic application of relevant Errata** — This setting enables the automatic application of Errata Updates to the selected systems. This means packages associated with Errata are updated without any user intervention. Customers should note that Red Hat does not recommend the use of the auto-update feature for production systems because conflicts between packages and environments can cause system failures.

7.4.5. Advanced Search —

The **System Search** page allows you to search through your systems according to specific criteria. These criteria include custom system information, system details, hardware, devices, interface, networking, packages, and location.

Searches can be refined using the **Fields to Search** drop-down menu, which is set to **Name/Description** by default.

The following list details the **Fields to Search** drop-down menu.

- **DMI Info** — The *Desktop Management Interface* (DMI) is a standard for management of components on computer system. You can search for RHN Satellite systems using the following DMI retrieval methods:
 - **System** — Product names or numbers, Manufacturer names, Serial numbers, and other information that may be unique to a system
 - **BIOS** — BIOS support information such as BIOS vendor name and version, hardware support enabled in the BIOS, and more
 - **Asset Tag** — A unique identifier assigned by an IT department (or vendor) to a system for better tracking, management and inventory
- **Location** — The physical location of a system, which includes the following:
 - **Address** — The address of the system or system set
 - **Building** — The building or site in an address
 - **Room** — The server or system room within a building
 - **Rack** — The designated location within a server room where a system is situated.
- **Details** — The unique identifiers assigned to a system by system administrators and particularly Satellite Administrators, including the following:
 - **Name/Description** — The name assigned to a system by the Satellite Administrator upon adding it to the RHN Satellite server.

- ID — An identifier that is unique to a system or system set.
- Custom Info — Information about the system that is unique only to that system.
- Snapshot Tag — The name assigned to a new or previous system snapshot
- Running Kernel — The currently running kernel on a system registered to the Satellite
- **Hardware** — Systems can be searched by particular components in the system, including the following:
 - CPU Model — The CPU model name (such as **Pentium** or **Athlon**)
 - CPU MHz Less Than — Search systems with a processor less than a user-designated speed in Megahertz.
 - CPU MHz More Than — Search systems with a processor more than a user-designated speed in Megahertz.
 - Number of CPUs Less Than — Search systems with a sum of processors less than a user-designated quantity.
 - Number of CPUs Greater Than — Search systems with a sum of processors greater than a user-designated quantity.
 - RAM Less Than — Search systems with a sum of memory less than a user-designated quantity in megabytes.
 - RAM More Than — Search systems with a sum of memory more than a user-designated quantity in megabytes.
- **Packages** — Systems can be searched by the packages installed (and not yet installed) on the system.
 - Installed Packages — Filter systems based on particular installed packages
 - Needed Packages — Filter systems based on particular packages that have yet to be installed
- **Activity** — Systems can be searched by the amount of time since first or last check-in with the RHN Satellite
 - Days Since Last Check-in — The amount of time (in days) that systems have last checked into RHN Satellite.
 - Days Since First Check-in — The amount of time (in days) that have passed since the systems first checked into RHN Satellite
- **Network Info** — Systems can be searched based on specific networking details such as IP address.
 - Hostname — The name associated with a system registered to RHN Satellite
 - IP Address — The network address of the system registered to RHN Satellite

- **Hardware Devices** — Systems can be searched by specific hardware details such as driver names and Device or Vendor IDs
 - Description — Device summary information, such as brand or model name/number (such as **Intel 82801HBM/HEM**)
 - Driver — The kernel driver or module name (such as **tulip.o** or **iw13945**)
 - Device ID — The hexadecimal number corresponding to the device installed in the system.
 - Vendor ID — The hexadecimal number corresponding to the vendor of the device installed in the system.

The Activity selections (**Days Since Last Checkin**, for instance) can be especially useful in finding and removing outdated System Profiles. Type the keyword, select the criterion to search by, use the radio buttons to identify whether you wish to query all systems or only those loaded in the **System Set Manager**, and click the **Search** button. You may also select the **Invert Result** checkbox to list those systems that do *not* match the criteria selected.

The results appear at the bottom of the page. For details about using the resulting system list, refer to [Section 7.4.2, “Systems”](#).

7.4.6. Activation Keys —

RHN Management and Provisioning customers with the Activation Key Administrator role (including Satellite Administrators) can generate activation keys through the RHN website. These keys can then be used to register a Red Hat Enterprise Linux system, entitle the system to an RHN service level and subscribe the system to specific channels and system groups through the command line utility **rhnreg_ks**. Refer to [Section 4.5, “Registering with Activation Keys”](#) for instructions on use.



Note

System-specific activation keys created through the **Reactivation** subtab of the **System Details** page are not part of this list because they are not reusable across systems.

7.4.6.1. Managing Activation Keys


To generate an activation key:

1. Select **Systems** => **Activation Keys** from the top and left navigation bars.
2. Click the **create new key** link at the top-right corner.



Warning

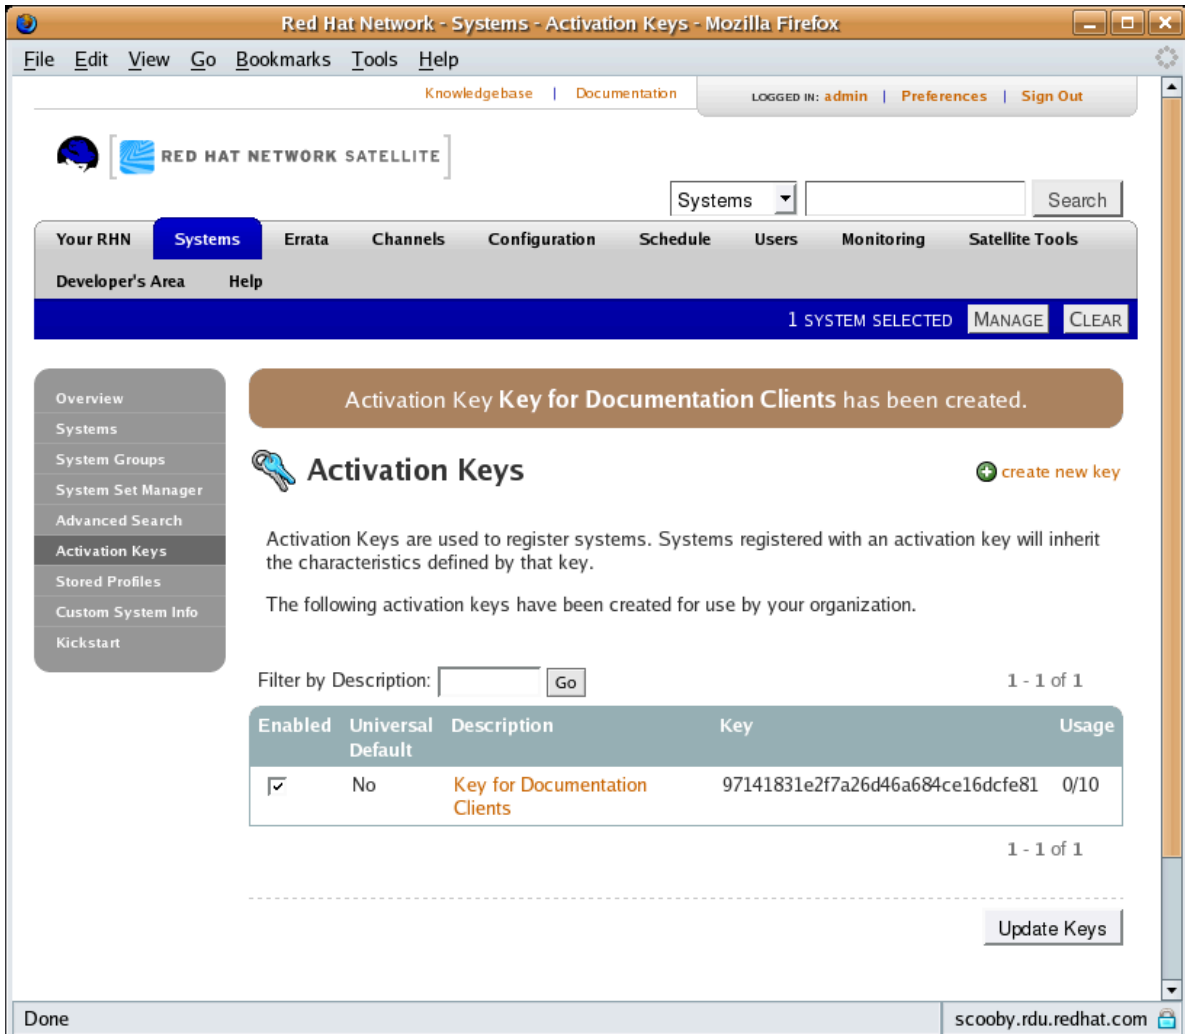
In addition to the fields listed below, RHN Satellite customers may also populate the **Key** field itself. This user-defined string of characters can then be supplied with **rhnreg_ks** to register client systems with the Satellite. *Do not insert commas in the key.* All other characters are accepted. Commas are problematic since they are

the separator used when including two or more activation keys at once. Refer to *Section 7.4.6.2, “Using Multiple Activation Keys at Once —  ”* for details.

3. Provide the following information:

- **Description** — User-defined description to identify the generated activation key.
- **Usage Limit** — The maximum number of registered systems that can be registered to the activation key at any one time. Leave blank for unlimited use. Deleting a system profile reduces the usage count by one and registering a system profile with the key increases the usage count by one.
- **Base Channel** — The primary channel for the key. Selecting nothing will enable you to select from all child channels, although systems can be subscribed to only those that are applicable.
- **Add-on Entitlements** — The supplemental entitlements for the key, which includes Monitoring, Provisioning, Virtualization, and Virtualization Platform. All systems will be given these entitlements with the key.
- **Universal default** — Whether or not this key should be considered the primary activation key for your organization.

Click **Create Key**.



Red Hat Network - Systems - Activation Keys - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

Knowledgebase | Documentation

LOGGED IN: admin | Preferences | Sign Out

RED HAT NETWORK SATELLITE

Systems

Your RHN **Systems** Errata Channels Configuration Schedule Users Monitoring Satellite Tools

Developer's Area Help

1 SYSTEM SELECTED **MANAGE** CLEAR

Activation Key Key for Documentation Clients has been created.

Activation Keys + create new key

Activation Keys are used to register systems. Systems registered with an activation key will inherit the characteristics defined by that key.

The following activation keys have been created for use by your organization.

Filter by Description: Go 1 - 1 of 1

Enabled	Universal Default	Description	Key	Usage
<input checked="" type="checkbox"/>	No	Key for Documentation Clients	97141831e2f7a26d46a684ce16dcfe81	0/10

1 - 1 of 1

Update Keys

Done scooby.rdu.redhat.com

Figure 7.6. Activation Keys

After creating the unique key, it appears in the list of activation keys along with the number of times it has been used. Note that only Activation Key Administrators can see this list. At this point, you may associate child channels and groups with the key so that systems registered with it automatically subscribe to them.

To change information about a key, such as the channels or groups, click its description in the key list, make your modifications in the appropriate tab, and click the **Update Key** button. To disassociate channels and groups from a key, deselect them in their respective menus by **Ctrl**-clicking their highlighted names. To remove a key entirely, click the **delete key** link in the top-right corner of the edit page.

A system may be set to subscribe to a base channel during registration with an activation key. However, if the activation key specifies a base channel that is not compatible with the operating system of the systems, the registration fails. For example, a Red Hat Enterprise Linux AS v.4 for x86 system cannot register with an Activation Key that specifies a Red Hat Enterprise Linux ES v.4 for x86 base channel. A system is always allowed to subscribe to a custom base channel.

To disable system activations with a key, unselect the corresponding checkbox under the **Enabled** column in the key list. The key can be re-enabled by selecting the checkbox. After making these changes, click the **Update Keys** button on the bottom right-hand corner of the page.

7.4.6.2. Using Multiple Activation Keys at Once —

Provisioning customers should note that multiple activation keys can be included at the command line or in a single kickstart profile. This allows you to aggregate the aspects of various keys without recreating a new key specific to the desired systems, simplifying the registration and kickstart processes while slowing the growth of your key list.

Without this stacking ability, your organization would need at least six activation keys to manage four server groups and subscribe a server to any two groups. Factor in two versions of the operating system, such as Red Hat Enterprise Linux 4 and 5, and you need twice the number of activation keys. A larger organization would need keys in the dozens.

Registering with multiple activation keys requires some caution; conflicts between some values cause registration to fail. Conflicts in the following values do not cause registration to fail, a combination of values is applied: software packages, software child channels, and config channels. Conflicts in the remaining properties are resolved in the following manner:

- base software channels — registration fails
- entitlements — registration fails
- enable config flag — configuration management is set

Do not use system-specific activation keys along with other activation keys; registration fails in this event.

You are now ready to use multiple activation keys at once. This is done with comma separation at the command line with `rhncfg_ks` or in a kickstart profile within the **Post** tab of the **Kickstart Details** page. Refer to [Section 4.5, “Registering with Activation Keys”](#) and [Section 7.4.9.3, “Create a New Kickstart Profile”](#), respectively, for instructions.

7.4.7. Stored Profiles —

RHN Provisioning customers can create package profiles through the **Profiles** subtab of the **Packages** tab within the **System Details** page. Those profiles are displayed on the **Stored Profiles** page, where they may be edited and even deleted.

To edit a profile, click its name in the list, alter its name and description, and click the **Update Profile** button. To view software associated with the profile, click the **Packages** subtab. To remove the profile entirely, click **delete stored profile** at the upper-right corner of the page.

7.4.8. Custom System Info —

RHN Provisioning customers may include completely customizable information about their systems. Unlike notes, the information here is more formal and may be searched upon. For instance, you may decide to identify an asset tag for each system. To do this, you must create an **asset** key within the **Custom System Info** page.

Click **create new key** at the upper-right corner of the page. Enter a descriptive label and description, such as **Asset** and **Precise location of each system**, and click the **Create Key**. The key will then show up in the custom info keys list.

Once the key exists, you may assign a value to it through the **Custom Info** tab of the **System Details** page. Refer to [Section 7.4.2.9.1.7, “System Details ▯ Details ▯ Custom Info — !\[\]\(35e4f762fc1cfea5610d92e2d225d5b4_img.jpg\) ”](#) for instructions.

7.4.8.1. **rhn-custom-info**

In addition to the Satellite web interface for creating and listing custom information keys, there is a command-line tool called **rhn-custom-info** that performs the same actions at a shell prompt, for administrators who may not have access to the web interface.

The usage of **rhn-custom-info** is as follows:

```
rhn-custom-info options key1 value1
```

For example:

```
rhn-custom-info --username=admin --password=f00b4rb4z --server-url=satellite.example.com --list-values
```

The command lists the custom keys and their values for the satellite.example.com Satellite server.

For more information, refer to the help file by typing **rhn-custom-info -h**.

7.4.9. Kickstart —

Kickstart configuration files allow administrators to create an environment for automating otherwise time-consuming system installations, such as multiple servers or workstations. Kickstart files can be created, modified, and managed within the RHN Satellite interface, and customized by the RHN Satellite web-based interface.


RHN Satellite also features the **Cobbler** installation server that allows administrators to perform unattended installations using a Pre-Execution Environment (PXE) server, installation and configuration of full and para-virtualized guest systems, and re-installation of running systems. For more information on configuring Cobbler and its associated helper program **Koan**, refer to [Chapter 11, Cobbler](#).

To satisfy the provisioning needs of customers, RHN Satellite provides an interface for developing kickstart profiles that can be used to install Red Hat Enterprise Linux or other operating systems on either new or already-registered systems. This enables systems to be installed automatically to particular specifications.



Important

If your systems are connected to RHN Hosted servers, you will need an external installation tree for each distribution to be kickstarted. This tree can be hosted anywhere that is accessible by the target system via HTTP. If the systems are connected through an RHN Proxy Server, then you may place the installation tree in `/var/www/html/pub/` on the Proxy. RHN Satellites already have a tree for each Red Hat distribution and therefore do not require separate trees. Even if the system connects through an RHN Proxy Server

to get to the Satellite, these trees will be available for kickstart. Refer to *Section 7.4.9.6*, “Kickstart ▯ Distributions —  ” for instructions on setting up installation trees.

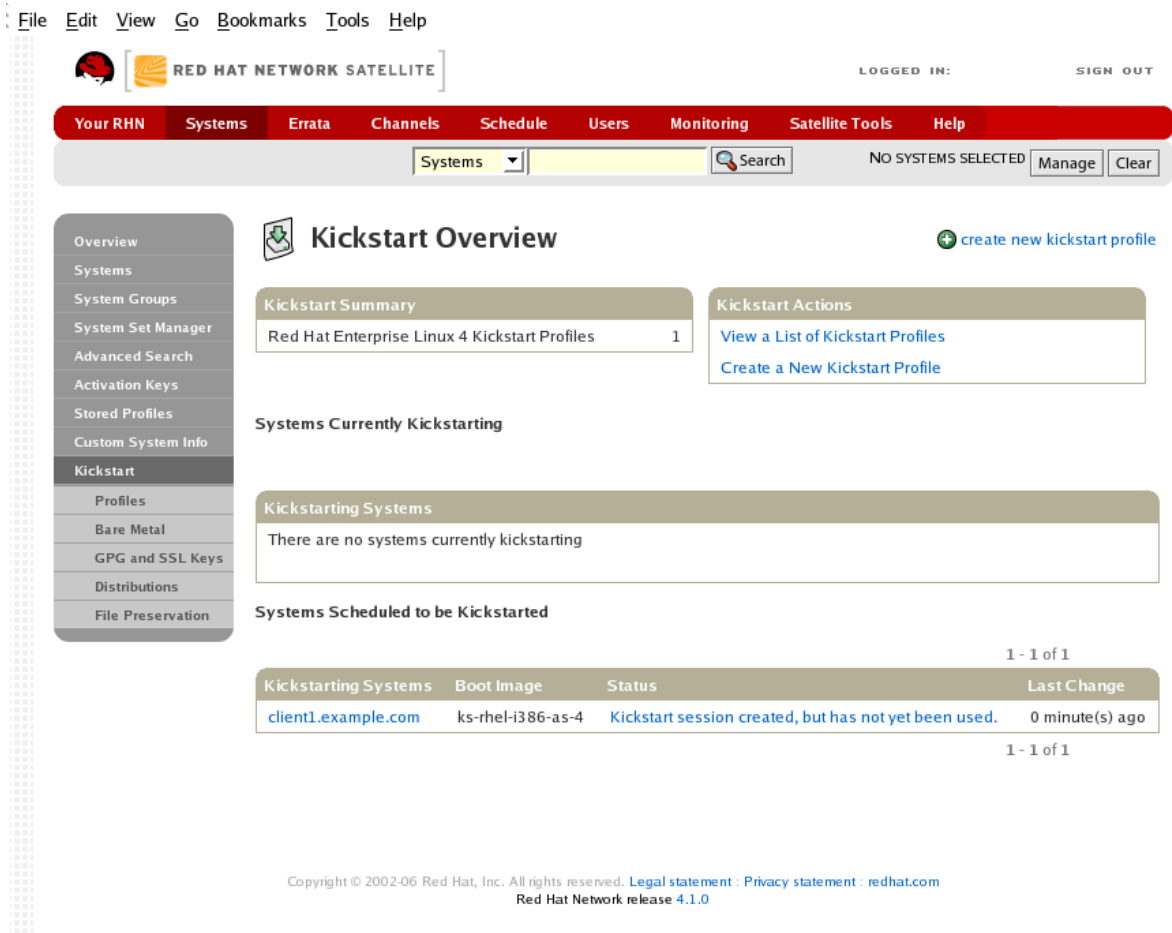


Figure 7.7. Kickstart Overview

This overview page displays the status of kickstart on your client systems: the types and number of profiles you have created and the progress of systems that are scheduled to be kickstarted. In the upper right is the **Kickstart Actions** section, which contains a series of links to management actions for your kickstart profiles. Before explaining the various kickstart options that are available from this page, the next section provides some introduction to the subject of kickstart.

7.4.9.1. Introduction to Kickstart

Many system administrators would prefer to use an automated installation method to install Red Hat Enterprise Linux on their machines. To answer this need, Red Hat created the kickstart installation method. Using kickstart, a system administrator can create a single file containing the answers to all the questions that would normally be asked during a typical installation.

Kickstart files can be kept on a single server system and read by individual computers during the installation. This installation method can support the use of a single kickstart file to install Red Hat Enterprise Linux on multiple machines, making it ideal for network and system administrators.

The *Red Hat Enterprise Linux System Administration Guide* contains an in-depth discussion of kickstart and is available here: <http://www.redhat.com/docs/manuals/enterprise/>.

7.4.9.1.1. Kickstart Explained

When a machine is to receive a network-based kickstart, the following events must occur in this order:

1. After being placed on the network and turned on, the machine's PXE logic broadcasts its MAC address and a request to be discovered.
2. If a static IP address is not being used, the DHCP server recognizes the discovery request and extends an offer of network information needed for the new machine to boot. This includes an IP address, the default gateway to be used, the netmask of the network, the IP address of the TFTP or HTTP server holding the bootloader program, and the full path and file name of that program (relative to the server's root).
3. The machine applies the networking information and initiates a session with the server to request the bootloader program.
4. The bootloader, once loaded, searches for its configuration file on the server from which it was itself loaded. This file dictates which kernel and kernel options, such as the initial RAM disk (initrd) image, should be executed on the booting machine. Assuming the bootloader program is SYSLINUX, this file is located in the **pxelinux.cfg** directory on the server and named the hexadecimal equivalent of the new machine's IP address. For example, a bootloader configuration file for Red Hat Enterprise Linux AS 2.1 should contain:

```
port 0
prompt 0
timeout 1
default My_Label
label My_Label
    kernel vmlinuz
    append ks=http://myrhnsatellite/ initrd=initrd.img network apic
```

5. The machine accepts and uncompresses the init image and kernel, boots the kernel, and initiates a kickstart installation with the options supplied in the bootloader configuration file, including the server containing the kickstart configuration file.
6. This kickstart configuration file in turn directs the machine to the location of the installation files.
7. The new machine is built based upon the parameters established within the kickstart configuration file.

7.4.9.1.2. Kickstart Prerequisites

Although Red Hat Network has taken great pains to ease the provisioning of systems, some preparation is still required for your infrastructure to handle kickstarts. For instance, before creating kickstart profiles, you may consider:

- A DHCP server is not required for kickstarting, but it can make things easier. If you are using static IP addresses, you should select static IP while developing your kickstart profile.
- An FTP server can be used in place of hosting the kickstart distribution trees via HTTP.
- If conducting a bare metal kickstart, you should 1)Configure DHCP to assign required networking parameters and the bootloader program location. 2)Specify within the bootloader configuration file the kernel to be used and appropriate kernel options.

7.4.9.1.3. Building Bootable Kickstart ISOs

While you can schedule a registered system to be kickstarted to a new operating system and package profile, it is also useful to be able to kickstart a system that is not registered with RHN, or does not yet have an operating system installed. One common method of doing this is to create a bootable CD-ROM that is inserted into the target system. When the system is rebooted, it boots from the CD-ROM, loads the kickstart configuration from the RHN Servers or your Satellite, and proceeds to install Red Hat Enterprise Linux according to the kickstart profile you have created.

To do this, copy the contents of `/isolinux` from the first CD-ROM of the target distribution. Then edit the `isolinux.cfg` file to default to 'ks'. Change the 'ks' section to the following template:

```
label ks
kernel vmlinuz
  append text ks={url} initrd=initrd.img lang= devfs=nomount ramdisk_size=16438 \
  {ksdevice}
```

IP addressed-based kickstart URLs will look something like this:

```
http://my.sat.server/kickstart/ks/mode/ip_range
```

The kickstart distribution selected by the IP range should match the distribution from which you are building, or errors will occur. `{ksdevice}` is optional, but looks like:

```
ksdevice=eth0
```

It is possible to change the distribution for a kickstart profile within a family, such as Red Hat Enterprise Linux AS 4 to Red Hat Enterprise Linux ES 4, by specifying the new distribution label. Note that you cannot move between versions (2.1 to 3) or between updates (U1 to U2).

Next, you may customize `isolinux.cfg` further for your needs, such as by adding multiple kickstart options, different boot messages, shorter timeout periods, etc.

Next, create the ISO as described in the *Making an Installation Boot CD-ROM* section of the *Red Hat Enterprise Linux 3 Installation Guide*. Alternatively, issue the command:

```
mkisofs -o file.iso -b isolinux.bin -c boot.cat -no-emul-boot -boot-load-size 4 \
-boot-info-table -R -J -v -T isolinux/
```

Note that `isolinux/` is the relative path to the directory containing the isolinux files from the distribution CD, while `file.iso` is the output ISO file, which is placed into the current directory.


You may then burn the ISO to CD-ROM. To use the disc (assuming you left the label for the kickstart boot as 'ks'), boot the system and type "ks" at the prompt. When you press **Enter**, the kickstart should begin.

7.4.9.1.4. Integrating Kickstart with PXE

In addition to CD-ROM-based installs, RHN supports kickstarts through a Pre-Boot Execution Environment (PXE). This is less error-prone than CDs, enables kickstarting from bare metal, and integrates with existing PXE/DHCP environments.

To use this method, make sure your systems have network interface cards (NIC) that support PXE, install and configure a PXE server, ensure DHCP is running, and then place the appropriate files on an HTTP server for deployment. Once the kickstart profile has been created, use the URL from the **Kickstart Details** page, as for CD-ROM-based installs.

To obtain specific instructions for conducting PXE kickstarts, refer to the *PXE Network Installations* chapter of the *Red Hat Enterprise Linux 4 System Administration Guide*.

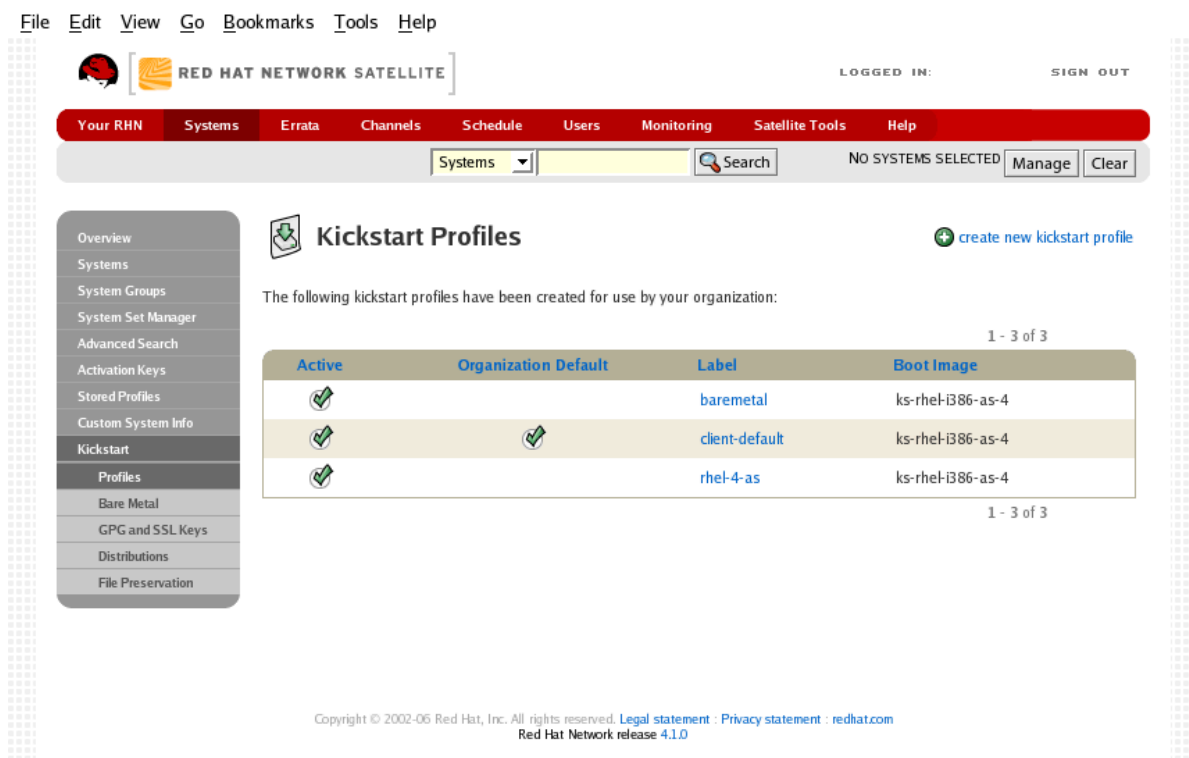


Tip

Upon running the **Network Booting Tool** as described in the Red Hat Enterprise Linux 4: System Administration Guide, ensure that you select "HTTP" as the protocol and include the domain name of the RHN Satellite in the Server field if you intend to use it to distribute the installation files.

The following sections describe the kickstart options available from the **Systems** ▢ **Kickstart** page.

7.4.9.2. Kickstart Profiles



File Edit View Go Bookmarks Tools Help

LOGGED IN: SIGN OUT

Your RHN Systems Errata Channels Schedule Users Monitoring Satellite Tools Help

Systems Search NO SYSTEMS SELECTED Manage Clear

Overview
Systems
System Groups
System Set Manager
Advanced Search
Activation Keys
Stored Profiles
Custom System Info
Kickstart
Profiles
Bare Metal
GPG and SSL Keys
Distributions
File Preservation

Kickstart Profiles [+ create new kickstart profile](#)

The following kickstart profiles have been created for use by your organization:

Active	Organization Default	Label	Boot Image
<input checked="" type="checkbox"/>		baremetal	ks-rhel-B386-as-4
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	client-default	ks-rhel-B386-as-4
<input checked="" type="checkbox"/>		rhel-4-as	ks-rhel-B386-as-4

1 - 3 of 3

1 - 3 of 3

Copyright © 2002-06 Red Hat, Inc. All rights reserved. [Legal statement](#) · [Privacy statement](#) · [redhat.com](#)
Red Hat Network release 4.1.0

Figure 7.8. Kickstart Profiles

This page lists all profiles for your organization, whether those profiles are active, and the distribution tree to which that profile is associated. You can either create a new kickstart profile by clicking the **create new kickstart profile** link, upload or paste the contents of a new kickstart using the **upload new kickstart file**, or edit an existing profile by clicking the name of the profile.

7.4.9.3. Create a New Kickstart Profile

Click on the **Create a New Kickstart Profile** link from the **Systems** ▢ **Kickstart** page to start the brief wizard that populates the base values needed for a kickstart profile.

1. On the first line, enter a kickstart profile label. This label cannot contain spaces, so use dashes (-) or underscores (_) as separators.
2. Select a **Base Channel** for this profile, which consists of packages based on a specific architecture and Red Hat Enterprise Linux release, such as **Red Hat Enterprise Linux (v.5 for 32-bit x86)**.
3. Select a kickstartable tree for this profile. The kickstartable tree drop-down menu is only populated if one or more distributions have been created for the selected base channel.
4. Select the **Virtualization Type** from the drop-down menu. For more information about virtualization, refer to [Chapter 10, Virtualization](#).



Note

If you do not intend to use the kickstart profile to create virtual guest systems, you can leave the drop-down at the default **KVM Virtualized Guest** choice.

5. On the second page, select (or enter) the URL of the kickstart tree.
6. On the third page, select a root password for the system. Be sure to follow the password recommendations from the *Password Security* section of the *Red Hat Enterprise Linux Security Guide*, available at <http://www.redhat.com/docs/manuals/enterprise/>.

Depending on your base channel, your newly created kickstart profile may be subscribed to a channel that is missing required packages. In order for kickstart to work properly, the following packages should be present in this kickstart's base channel: **pyOpenSSL**, **rhnl1b**, **libxml2-python**, and **spacewalk-koan** and associated packages.

To resolve this issue, ensure that the following items are correct:

- Make sure that the **rh-tools** child software channel for the kickstart profile's base channel is available to your organization. If it is not, you must request entitlements for the **rh-tools** software channel from the Satellite administrator.
- Make sure that the **rh-tools** child channel for this kickstart profile's base channel is available to your RHN Satellite. If it is not, contact the Satellite administrator and request a **satellite-sync** of the **rh-tools**.
- Make sure that the **rh-kickstart** and associated packages corresponding to this kickstart are available in the kickstart **rh-tools** child channel. If it is not, you must make them available for this kickstart profile to function properly.

The final stage of the wizard presents the **Kickstart Details** ▢ **Details** tab. On this tab and the other sub-tabs, nearly every option for the new kickstart profile can be customized. The following sections describe the options available on each sub-tab.

7.4.9.3.1. Kickstart Details ▢ Details —




Figure 7.9. Kickstart Details

The figure above shows the sub-tabs that are available from the **Kickstart Details** tab.

From the **Kickstart Details ▢ Details** sub-tab, you can:

- Rename the profile
- Change the operating system it installs by clicking (**Change**)
- Change the **Virtualization Type**



Note

Changing the **Virtualization Type** may require changes to the kickstart profile bootloader and partition options, potentially overwriting user customizations. Consult the **Partitioning** tab to verify any new or changed settings.

- Change the amount of **Virtual Memory** (in Megabytes of RAM) allocated to virtual guests kickstarted with this profile
- Change the number of **Virtual CPUs** for each virtual guest
- Change the the **Virtual Storage Path** from the default in `/var/lib/xen/`
- Change the amount of **Virtual Disk Space** (in Gigabytes) allotted to each virtual guest
- Change the **Virtual Bridge** for networking of the virtual guest
- Deactivate the profile so that it cannot be used to schedule a kickstart by removing the **Active** checkmark
- Check whether to enable logging for custom `%post` scripts to the `/root/ks-post.log` file
- Check whether to enable logging for custom `%pre` scripts to the `/root/ks-pre.log` file
- Check whether to preserve the `ks.cfg` file and all `%include` fragments to the `/root/` directory of all systems kickstarted with this profile.
- Select whether this profile is the default for all of your organization's kickstarts by checking or unchecking the box.
- Add any **Kernel Options** in the corresponding text box.
- Add any **Post Kernel Options** in the corresponding text box.
- Enter comments that are useful to you in distinguishing this profile from others

7.4.9.3.2. Kickstart Details ▯ Operating System —

From this page, you can make the following changes to the operating system that the kickstart profile installs:

Change the base channel

Select from the available base channels, such as **Red Hat Enterprise Linux v.5 for 32-bit x86**. Satellite administrators can see a list of all base channels that are currently synced to the Satellite.

Child Channels

Subscribe to any available child channels of the base channel, such as the **rhn-tools*** channel.

Available Trees

Use the drop-down menu to choose the available trees that are associated with the base channel.

File Location

The exact location from which the kickstart tree is mounted. This value is determined when the profile is created. You can view it on this page but you cannot change it.

7.4.9.3.3. Kickstart Details ▯ Variables

Kickstart variables can be used to substitute values into kickstart profiles. To define a variable, create a name-value pair (*name/value*) in the text box.

For example, if you wanted to kickstart a system that joins the network for specified department (for example the Engineering organization) you can create a profile variable to set the ip address and the gateway server address to a variable that any system using that profile will use. Add the following line to the **Variables** text box.

```
IPADDR=192.168.0.28
GATEWAY=192.168.0.1
```

To use the profile variable, you can use the name of the variable within the profile to substitute in the value. For example, the **network** portion of a kickstart file looks like the following:

```
network --bootproto=static --device=eth0 --onboot=on --ip=$IPADDR --gateway=$GATEWAY
```

The **\$IPADDR** will be **192.168.0.28**, and the **\$GATEWAY** will be **192.168.0.1**



Note

There is a hierarchy when creating and using variables in kickstart files. System kickstart variables take precedence over Profile variables, which in turn take precedence over Distribution variables. Understanding this hierarchy can alleviate confusion when using variables in kickstarts.

Using variables are just one part of the larger Cobbler infrastructure for creating templates that can be shared between multiple profiles and systems. For more information about Cobbler and kickstart templates, refer to [Chapter 11, Cobbler](#).

7.4.9.3.4. Kickstart Details ▯ Advanced Options —

From this page, you can toggle several installation options on and off by checking and unchecking the boxes to the left of the option. For most installations, the default options are correct. The *Red Hat Enterprise Linux System Administration Guide* discusses each of these options in detail.

7.4.9.3.5. Kickstart Details ▯ Bare Metal Kickstart —

This sub-tab provides the information necessary to kickstart systems that are not currently registered with RHN. Using the on-screen instructions, you may either kickstart systems using boot media (CD-ROM) or by IP address.

7.4.9.3.6. System Details ▯ Details —



Figure 7.10. System Details

The figure above shows the sub-tabs that are available from the **System Details** tab.

From the **System Details ▯ Details** sub-tab, you can:

- Select from DHCP and static IP, depending on your network
- Choose the level of SELinux that is configured on kickstarted systems
- Enable configuration management or remote command execution on kickstarted systems
- Change the root password associated with this profile

7.4.9.3.7. System Details ▯ Locale —

From this sub-tab, you can change the timezone associated with kickstarted systems.

7.4.9.3.8. System Details ▯ Partitioning —

From this sub-tab, you can indicate the partitions that you wish to be created during installation. For example:

```
partition /boot --fstype=ext3 --size=200
partition swap --size=2000
partition pv.01 --size=1000 --grow
volgroup myvg pv.01 logvol / --vgname=myvg --name=rootvol --size=1000 --grow
```

7.4.9.3.9. System Details ▢ File Preservation —

If you have previously created a file preservation list, you may include that list as part of the kickstart. This will prevent the files in that list from being over-written during the installation process. Refer

to [Section 7.4.9.7, “Kickstart ▢ File Preservation — !\[\]\(3d8c13c92b853674f749aac6fa869926_img.jpg\) ”](#) for information on how to create a file preservation list.

7.4.9.3.10. System Details ▢ GPG and SSL —

From this sub-tab, select the GPG keys and/or SSL certificates to be imported to the kickstarted system during the %post section of the kickstart. For Satellite customers, this list includes the SSL Certificate used during the installation of the Satellite.



Note

Any GPG key you wish to import to the kickstarted system must be in ASCII rather than binary format.

7.4.9.3.11. System Details ▢ Troubleshooting —

From this sub-tab, you can change information that may help with troubleshooting hardware problems:

Bootloader

For some headless systems, it is better to select the non-graphic LILO bootloader.

Kernel Parameters

Enter kernel parameters here that may help to narrow down the source of hardware issues.

7.4.9.3.12. Software ▢ Package Groups —



Figure 7.11. Software

The figure above shows the sub-tabs that are available from the **Software** tab.

Enter the package groups, such as **@office** or **@admin-tools** you would like to install on the kickstarted system in the large text box on this page. If you would like to know what package groups are available, and what packages they contain, refer to the **RedHat/base/** file of your kickstart tree. Satellite customers will most likely locate this file here: `/var/www/satellite/rhn/kickstart/<kickstart_label>/RedHat/base/comps.xml`.

7.4.9.3.13. Software ▢ Package Profiles —

If you have previously created a Package Profile from one of your registered systems, you can use that profile as a template for the files to be installed on a kickstarted system. Refer to

Section 7.4.2.9.2.2, “System Details ▯ Software ▯ Packages” for more information about package profiles.

7.4.9.3.14. Activation Keys —

Kickstart Details System Details Software **Activation Keys** Scripts Kickstart File

Figure 7.12. Activation Keys

The **Activation Keys** tab, which has no sub-tabs, allows you select Activation Keys to include as part of the kickstart profile. These keys, which must have been created previous to creating the kickstart profile, will be used when re-registering kickstarted systems.

7.4.9.3.15. Scripts —

Kickstart Details System Details Software Activation Keys **Scripts** Kickstart File

Figure 7.13. Scripts

The **Scripts** tab, which has no sub-tabs, is where %pre and %post scripts are created. This page lists any scripts that have already been created for this kickstart profile. To create a new kickstart script:

1. Click the add new kickstart script link in the upper right
2. Enter the path to the scripting language used to create the script, such as /usr/bin/perl
3. Enter the full script in the large text box
4. Indicate whether this script is to be executed in the %pre or %post section of the kickstart process
5. Indicate whether this script is to run outside of the chroot environment. Refer to the *Post-installation Script* section of the *Red Hat Enterprise Linux System Administration Guide* for further explanation of the **nochroot** option



Note

RHN supports the inclusion of separate files within the Partition Details section of the kickstart profile. For instance, you may dynamically generate a partition file based on the machine type and number of disks at kickstart time. This file can be created via %pre script and placed on the system, such as /tmp/part-include. Then you can call for that file by including the following line within the Partition Details field of the **System Details ▯ Partitioning** tab:

```
%include /tmp/part-include
```

7.4.9.3.16. Kickstart File —

[Kickstart Details](#) [System Details](#) [Software](#) [Activation Keys](#) [Scripts](#) **[Kickstart File](#)**

Figure 7.14. Kickstart File

The **Kickstart File** tab, which has no sub-tabs, allows you to view or download the kickstart profile that has been generated from the options chosen in the previous tabs.

7.4.9.4. Kickstart Bare Metal —

Lists the IP addresses that have been associated with kickstart profiles created by your organization. Click either the range or the profile name to access different tabs of the **Kickstart Details** page.

7.4.9.5. Kickstart GPG and SSL Keys —

Lists keys and certificates available for inclusion in kickstart profiles and provides a means to create new ones. This is especially important for customers of RHN Satellite or RHN Proxy Server because systems kickstarted by them must have the server key imported into RHN and associated with the relevant kickstart profiles. Import it by creating a new key here and then make the profile association in the **GPG and SSL keys** subtab of the **Kickstart Details** page.

To develop a new key/certificate, click the **create new stored key/cert** link in the upper-right corner of the page. Enter a description, select the type, upload the file, and click the **Update Key** button. Note that a unique description is required.



Important

The GPG key you upload to RHN must be in ASCII format. Using a GPG key in binary format causes anaconda, and therefore the kickstart process, to fail.

7.4.9.6. Kickstart Distributions —

The **Distributions** page enables you to find and create custom installation trees that may be used for kickstarting.



Note

The **Distributions** page does not display Red Hat distributions already provided. They can be found within the **Distribution** dropdown menu of the **Kickstart Details** page.)

Before creating a distribution, you must make an installation tree available, as described in the *Kickstart Installations* chapter of the *Red Hat Enterprise Linux System Administration Guide*. This tree must be located in a public directory on an HTTP or FTP server.



Important

RHN Satellite users should note that channels imported with `satellite-sync` are made available automatically and do not require the creation of a separate installation tree. These trees are available to client systems that kickstart through the Satellite. While you may be able to access the files from a non-kickstarting client, this functionality is not supported and may be removed at any time in the future.

To create a new distribution, enter a label (without spaces) in the **Distribution Label** field, such as `my-orgs-rhel-as-5`. In the **Tree Path** field, paste the path or URL to the base of the installation tree. (You can test this by appending "README" to the URL in a Web browser, pressing **Enter**, and ensuring that the distribution's readme file appears.)

Select the matching distribution from the **Base Channel** and **Installer Generation** dropdown menus, such as **Red Hat Enterprise Linux (v. 5 for 32-bit x86)** and **Red Hat Enterprise Linux 5**, respectively. When finished, click the **Create Kickstart Distribution** button.

7.4.9.6.1. Kickstart ▯ Distributions ▯ Variables

Kickstart variables can be used to substitute values into kickstart profiles. To define a variable, create a name-value pair (*name/value*) in the text box.

For example, if you wanted to kickstart a system that joins the network for specified department (for example the Engineering organization) you can create a profile variable to set the ip address and the gateway server address to a variable that any system using that profile will use. Add the following line to the **Variables** text box.

```
IPADDR=192.168.0.28
GATEWAY=192.168.0.1
```

To use the distribution variable, you can use the name of the variable within the profile to substitute in the value. For example, the **network** portion of a kickstart file looks like the following:

```
network --bootproto=static --device=eth0 --onboot=on --ip=$IPADDR --gateway=$GATEWAY
```

The **\$IPADDR** will be **192.168.0.28**, and the **\$GATEWAY** will be **192.168.0.1**



Note

There is a hierarchy when creating and using variables in kickstart files. System kickstart variables take precedence over Profile variables, which in turn take precedence over Distribution variables. Understanding this hierarchy can alleviate confusion when using variables in kickstarts.

Using variables are just one part of the larger Cobbler infrastructure for creating templates that can be shared between multiple profiles and systems. For more information about Cobbler and kickstart templates, refer to [Chapter 11, Cobbler](#).

7.4.9.7. Kickstart File Preservation

Collects lists of files to be protected and re-deployed on systems during kickstart. For instance, if you have many custom configuration files located on a system to be kickstarted, enter them here as a list and associate that list with the kickstart profile to be used.

To use this feature, click the **create new file preservation list** link at the top and enter a relevant label and all files and directories to be preserved on the resulting page. Enter absolute paths to all files and directories. Then click **Create List**.



Important

Although file preservation is useful, it does have limitations. First, each list is limited to a total size of 1 MB. Further, special devices like `/dev/hda1` and `/dev/sda1` are not supported. Finally, only file and directory names may be entered. No regular expression wildcards can be included.

When finished, you may include the file preservation list in the kickstart profile to be used on systems containing those files. Refer to [Section 7.4.9.3, “Create a New Kickstart Profile”](#) for precise steps.


7.5. Errata

Select the **Errata** tab from the top navigation bar to track the availability and application of errata to your managed systems.

The first page to appear here is the **Errata Overview** page. This page displays relevant errata, which are errata that apply to at least one system to which you have administrative access and that have not yet been applied.






Tip

To receive an email when Errata Updates are issued for your system, go to **Overview ** **Your Preferences** and select **Receive email notifications**.

Red Hat releases Errata Updates in three categories, or types: Security Updates, Bug Fix Updates, and Enhancement Updates. Each Errata Update is comprised of a summary of the problem and the solution, including the RPM packages required to fix the problem.

Icons are used to identify the three types of Errata Updates:

-  — Security Updates available, update *strongly* recommended
-  — Bug Fix Updates available and recommended
-  — Enhancement Updates available

A summary of each erratum is provided in list form. This view instantly informs you of the type, severity (for Security Updates), and subject of the erratum, as well as the number of affected systems.

In addition to the pages described within this chapter, you may view Errata by product line from the following location: <https://rhn.redhat.com/errata>.

7.5.1. Relevant Errata

As shown in [Figure 7.15, “Errata List”](#), the **Relevant Errata** page displays a customized list of Errata Updates that applies to your registered systems. The list provides a summary of each Errata Update, including its type, severity (for Security Updates), advisory number, synopsis, systems affected, and date updated.

The screenshot shows the Red Hat Network Errata page in a Mozilla Firefox browser window. The page title is "Red Hat Network - Errata - Errata - All". The browser's address bar shows "https://rhn.redhat.com/errata". The page has a navigation menu with "Your RHN", "Systems", "Errata", "Channels", "Configuration", "Schedule", "Users", "Satellite Tools", and "Help". The "Errata" tab is selected. Below the navigation menu, there is a search bar and a "Systems" dropdown menu. The main content area displays "All Errata" with a filter by synopsis and a "Go" button. A table lists the errata updates with columns for Type, Advisory, Synopsis, Systems, and Updated. The table shows 20 items, with the first 19 being security updates (RHSAs) and the 20th being an enhancement update (RHEA).

Type	Advisory	Synopsis	Systems	Updated
Security	RHSA-2007:0964	Important: openssl security update	0	10/12/07
Security	RHSA-2007:0960	Important: hplip security update	0	10/11/07
Security	RHSA-2007:0905	Moderate: kdbase security update	0	10/8/07
Security	RHSA-2007:0909	Moderate: kdelibs security update	0	10/8/07
Enhancement	RHEA-2007:0928	tzdata enhancement update	1	10/4/07
Security	RHSA-2007:0933	Moderate: elinks security update	0	10/3/07
Security	RHSA-2007:0323	Important: xen security update	0	10/2/07
Security	RHSA-2007:0951	Important: nfs-utils-lib security update	0	10/2/07
Security	RHSA-2007:0936	Important: kernel security update	0	9/27/07
Security	RHSA-2007:0937	Important: kernel security update	0	9/27/07
Security	RHSA-2007:0513	Moderate: gimp security update	0	9/26/07
Security	RHSA-2007:0871	Moderate: tomcat security update	0	9/26/07
Security	RHSA-2007:0890	Moderate: php security update	0	9/20/07
Security	RHSA-2007:0845	Important: libvorbis security update	0	9/19/07
Security	RHSA-2007:0013	Important: nfs-utils-lib security update	0	0/10/07

Figure 7.15. Errata List

Clicking on the Advisory takes you to the **Details** tab of the **Errata Details** page. Clicking on the number of associated systems takes you to the **Affected Systems** tab of the **Errata Details** page. Refer to [Section 7.5.2.2, “Errata Details”](#) for more information.

7.5.2. All Errata

The **All Errata** page displays a list of all Errata Updates released by Red Hat. It works much the same as the **Relevant Errata** page in that clicking either the Advisory or the number of systems affected takes you to related tabs of the **Errata Details** page. Refer to [Section 7.5.2.2, “Errata Details”](#) for more information.

7.5.2.1. Apply Errata Updates

Errata Updates include a list of updated packages that are required to apply the Errata Update. To apply Errata Updates to a system, the system must be entitled.

Apply all applicable Errata Updates to a system by clicking on **Systems** ▢ **Systems** in the top and left navigation bars. Click on the name of an entitled system, and click the **Errata** tab of the resulting **System Details** page. When the Relevant Errata list appears, click **Select All** then the **Apply Errata** button on the bottom right-hand corner of the page. Only those Errata that have not been scheduled or were scheduled and failed or canceled are listed. Updates already pending are excluded from the list.

In addition, Management users can apply Errata Updates using two other methods:

- To apply a specific Errata Update to one or more systems, find the update within the Errata lists. In the table, click on the number of systems affected, which takes you to the **Affected Systems** tab of the **Errata Details** page. Select the individual systems to be updated and click the **Apply Errata** button. Double-check the systems to be updated on the confirmation page, then click the **Confirm** button.
- To apply more than one Errata Update to one or more systems, select the systems from a **Systems** list and click the **Update List** button. Click the **System Set Manager** link in the left navigation bar, then click the **Systems** tab. After ensuring the appropriate systems are selected, click the **Errata** tab, select the Errata Updates to apply, and click the **Apply Errata** button. You can select to apply the Errata as soon as possible (the next time the Red Hat Network Daemon on the client systems connect to RHN) or schedule a date and time for the Errata Updates to occur. Then click the **Schedule Updates** button. You can follow the progress of the Errata Updates through the **Pending Actions** list. Refer to [Section 7.8, “Schedule”](#) for more details.



Important

If you use scheduled package installation, the packages are installed via the RHN Daemon. You must enable the RHN Daemon on your systems. Refer to [Chapter 5, Red Hat Network Daemon](#) for more details.

The following rules apply to Errata Updates:

- Each package is a member of one or more channels. If a selected system is not subscribed to a channel containing the package, the package will not be installed on that system.
- If a newer version of the package is already on the system, the package will not be installed on that system.
- If an older version of the package is installed, the package will be upgraded.

7.5.2.2. Errata Details

If you click on the Advisory of an Errata Update in the **Relevant** or **All** pages, its **Errata Details** page appears. This page is further divided into the following tabs:

7.5.2.2.1. Errata Details ▢ Details

This subtab displays the Erratum Report issued by Red Hat. It provides a synopsis of the erratum first, including the severity (for security updates), issue date, and any update dates. This is followed by brief and detailed descriptions of the erratum and the steps required to resolve the issue.

Below the **Affected Channels** label, all channels that contain the affected package are listed. Clicking on a channel name displays the **Packages** subtab of the **Channel Details** page for that channel. Refer to [Section 7.6.1.9, “Software Channel Details”](#) for more information.

Below **Fixes**, the specific Bugzilla entries resolved by this erratum are listed. Clicking on any summary text opens that Bugzilla entry at <http://bugzilla.redhat.com>. Note that you must have a Bugzilla account to view the entry.

Security updates list the specific vulnerability as tracked by <http://cve.mitre.org>. This information is listed below the **CVEs** label.

Red Hat provides security update information in OVAL format. OVAL is an open vulnerability and assessment language promoted by Mitre, <http://oval.mitre.org>. Clicking on the link below the **Oval** label downloads this information to your system.

7.5.2.2. Errata Details ▢ Packages

Provides links to each of the updated RPMs broken down by channel. Clicking on the name of a package displays its **Package Details** page.

7.5.2.2.3. Errata Details ▢ Affected Systems

Lists systems affected by the Errata Update. You can apply updates here. (See [Section 7.5.2.1, “Apply Errata Updates”](#).) Clicking on the name of a system takes you to its **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.

To help users determine whether an update has been scheduled, a Status column exists within the affected systems table. Possible values are: None, Pending, Picked Up, Completed, and Failed. This column identifies only the latest action related to an Erratum. For instance, if an action fails and you reschedule it, this column shows the status of the Erratum as Pending (with no mention of the previous failure). Clicking a status other than None takes you to the **Action Details** page. This column corresponds to one on the **Errata** tab of the **System Details** page.

7.5.3. Advanced Search

The **Erratum Search** page allows you to search through Errata according to specific criteria.

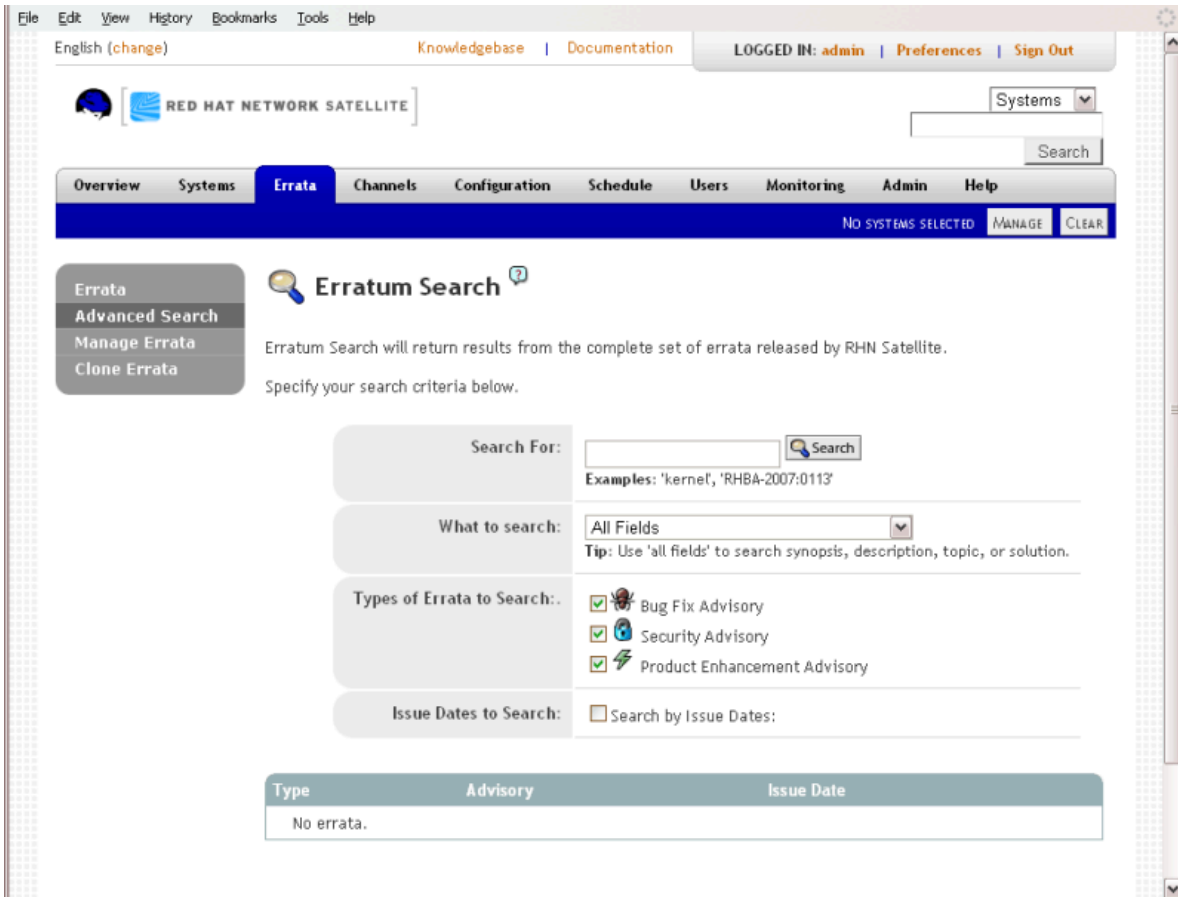


Figure 7.16. Erratum Search

- **All Fields** — Search errata by synopsis, description, topic, or solution.
- **Erratum Advisory** — The way Red Hat Security Response Team codifies Advisories, such as:

RHBA-2007:0530

Searches can be by done year (such as 2007), by type of Advisory (RHBA for Bug fixes, RHEA for Enhancements, and RHSA for Security advisories), or full Advisory name, such as the example above.

- **Package Name** — Users concerned with particular packages can search by package name, such as:

kernel

Package search can be beneficial because search results will be grouped by advisory. For example, searching for kernel-related bugs return results where all packages with the term **kernel** appear grouped by the advisory for which the bug is related.

- **CVE Name** — The name assigned to the Security advisory (RHSA) by the Common Vulnerabilities and Exposures project at <http://cve.mitre.org>. For example:

CVE-2006-4535

You may also filter errata search results by the type of errata issued. Check or uncheck the boxes next to the type of advisory to search.

- Bug Fix Advisory — Errata that contains fixes to issues that were reported by users or discovered during development or testing
- Security Advisory — Errata that fixes a security issue found during development, testing, or reported by users or a software security clearing house. A security advisory usually has one or more CVE names associated with each vulnerability found in each erratum.
- Product Enhancement Advisory — Errata that contains new features, improved functionality, or enhanced performance in the package's software.

7.6. Channels

If you click the **Channels** tab on the top navigation bar, the **Channels** category and links appear. The pages in the **Channels** category enable you to view and manage the channels and packages associated with your systems. In addition, you can obtain ISO images here.

7.6.1. Software Channels

The **Software Channels** page is the first to appear in the **Channels** category. A software channel is a list of Red Hat Enterprise Linux packages grouped by use. Channels are used to choose packages to be installed on a system.

There are two types of software channels: *base channels* and *child channels*.

7.6.1.1. Base Channels

A base channel consists of a list of packages based on a specific architecture and Red Hat Enterprise Linux release. For example, all of the packages in Red Hat Enterprise Linux 5 for the x86 architecture make up a base channel. The list of packages in Red Hat Enterprise Linux 5 for the Itanium architecture make up a different base channel.

A system must be subscribed to one base channel only. This base channel is assigned automatically during registration based upon the Red Hat Enterprise Linux release and system architecture selected. In the case of public free channels, the action will succeed. In the case of paid base channels, this action will fail if an associated entitlement does not exist.

7.6.1.1.1. Extended Update Support (EUS)

In addition to base channels for major versions of Red Hat Enterprise Linux, there are channels for update versions of Red Hat Enterprise Linux, which are also separated by architecture and which can have child channels. These *Extended Update Support* (EUS) channels are for administrators who want to stay with one major or update version of Red Hat Enterprise Linux and customize their package updates for their particular version, rather than upgrade their systems to a new update version that installs new software, hardware drivers, and features on production systems.

For example, administrators can standardize their desktop systems to Red Hat Enterprise Linux 5.1 for x86, while managing servers on Red Hat Enterprise Linux 4.5 for AMD64 and EM64T. Administrators can stay on their version for the duration of the EUS support lifecycle, assured of the behavior of their software version. Additionally, administrators can install critical software updates without introducing bugs from untested new features or software.

7.6.1.2. Child Channels

A child channel is a channel associated with a base channel that contains extra packages. For instance, an organization can create a child channel associated with Red Hat Enterprise Linux 3 for the x86 architecture that contains extra packages needed only for the organization, such as a custom engineering application.

A system can be subscribed to multiple child channels of its base channel. Only packages included in a system's subscribed channels can be installed or updated on that system. Further, RHN Satellite and RHN Proxy Server customers have channel management authority. This authority gives them the ability to create and manage their own custom channels. Refer to the *RHN Channel Management Guide* for details.

Channels can be further broken down by their relevance to your systems, including All Channels, Red Hat Channels, Popular Channels, My Channels, Shared Channels, and Retired channels.

7.6.1.3. All Channels

As shown in [Figure 7.17, “All Channels”](#), the **All Channels** page is shown by default when you click **Software Channels** in the navigation bar. It displays a list of all channels available to your organization. Links within this list go to different tabs of the **Software Channel Details** page. Clicking on a channel name takes you to the **Details** tab. Clicking on the number of packages takes you to the **Packages** tab. Clicking on the number of systems takes you to the **Subscribed Systems** tab. Refer to [Section 7.6.1.9, “Software Channel Details”](#) for details.

Red Hat Network - Channels - Software Channels - Relevant - Relevant channels - Mozilla Firefox

English (change) Knowledgebase | Documentation LOGGED IN: admin | Preferences | Sign Out

RED HAT NETWORK SATELLITE Systems [] Search

Your RHN Systems Errata **Channels** Configuration Schedule Users Satellite Tools Help

NO SYSTEMS SELECTED MANAGE CLEAR

Software Channels Overview

Relevant channels All channels Retired channels

The software channels listed below are those **most relevant** to your organization. **Relevant Channels** include the 32-bit version of the channels your organization is entitled to, plus versions of those channels for any additional architectures that systems in your account have registered to. Channels no longer supported by Red Hat are not listed here.

Alternatively, you may also view a list of **all channels** or a list of 'end-of-life' **retired channels**.

Filter by Channel Name: [] Go

Show All Child Channels | Hide All Child Channels

Channel Name	Packages	Systems
<input type="checkbox"/> Red Hat Enterprise Linux (v. 5 for 32-bit x86)	3136	1
└ Red Hat Network Tools for RHEL Server (v.5 32-bit x86)	378	1
└ RHEL Virtualization (v. 5 for 32-bit x86)	66	1
<input type="checkbox"/> Clone of Clone of Clone of Red Hat Enterprise Linux (v. 5 for 32-bit x86)	2161	0
<input type="checkbox"/> Clone of Clone of Red Hat Enterprise Linux (v. 5 for 32-bit x86)	3136	0
<input type="checkbox"/> pt-custom1	2	0
<input type="checkbox"/> custom01	0	0
<input type="checkbox"/> Clone of Red Hat Enterprise Linux (v. 5 for 32-bit x86)	3136	0

Figure 7.17. All Channels

7.6.1.4. Red Hat Channels

The **Red Hat Channels** page displays the Red Hat channels and their available child channels. Versions of Red Hat Enterprise Linux synced directly from RHN Hosted, for example, are listed in this channel.

7.6.1.5. Popular Channels

The **Popular Channels** page displays the software channels most subscribed by systems registered to your organization. You can refine the search further by using the drop-down menu to list only the channels with at least a certain number of systems subscribed.

7.6.1.6. My Channels

The **My Channels** page displays all of the software channels that belong to your organization, which includes both Red Hat channels and custom channels. You can refine the search further by using the text box to filter by the channel name.

7.6.1.7. Shared Channels

The **Shared Channels** page displays the channels in your organization that you have shared with others in your organizational trust. For more information about organizational trust and channel sharing, refer to [Section 9.6.2, "Sharing Content Channels between Organizations in a Trust"](#).

7.6.1.8. Retired Channels

The Retired Channels page displays channels available to your organization that have reached their end-of-life dates. These channels do not receive updates.

7.6.1.9. Software Channel Details

If you click on the name of a channel, the **Software Channel Details** page appears. This page is broken down into the following tabs:

7.6.1.9.1. Software Channel Details ▢ Details

General information about the channel and the parent channel, if it is a child channel. This is the first tab displayed when you click on a channel. It displays essential information about the channel, such as summary, description, and architecture.



—In addition, a Globally Subscribable checkbox can be seen by Satellite Administrators and Channel Administrators. This signifies the default behavior of every channel allowing any user to subscribe systems to it. Unchecking this box and clicking **Update** causes the appearance of a **Subscribers** tab, which may then be used to grant certain users subscription permissions to the channel. Satellite Administrators and Channel Administrators can always subscribe systems to any channel.



— Only customers with custom base channels may change their systems' base channel assignment. They may do this through the website in two ways:

- Customers with a custom base channel may assign the system to that base channel.
- Customers may revert system subscriptions from a custom base channel to the appropriate distribution-based base channel.



Note

The system base channel's distribution variant must match the variant installed on the system. For example, a system that has Red Hat Enterprise Linux AS v.4 for x86 cannot be registered to a Red Hat Enterprise Linux ES v.4 for x86 base channel.

7.6.1.9.2. Software Channel Details ▢ Errata

List of Errata affecting the channel. The list displays advisory types, names, summaries, and the dates issued. Clicking on an advisory name takes you to its **Errata Details** page. Refer to [Section 7.5.2.2, "Errata Details"](#) for more information.

7.6.1.9.3. Software Channel Details ▢ Packages

List of packages in the channel. To download packages as a .tar file, select them and click the **Download Packages** button at the bottom-left corner of the page. Clicking on a package name takes you to the **Package Details** page. This page displays a set of tabs with information about the package, including which architectures it runs on, the package size, build date, package dependencies, the change log, list of files in the package, newer versions, and which systems have the package installed. From here, you can download the packages as RPMs or SRPMs.

To search for a specific package or a subset of packages, use the package filter at the top of the list. Enter a substring to search all packages in the list for package names that contain the string. For example, typing **ks** in the filter might return: **ksconfig**, **krb5-workstation**, and **links**. The filter is case-insensitive.

7.6.1.9.4. Software Channel Details ▯ Subscribed Systems

List of entitled systems subscribed to the channel. The list displays system names, base channels, and their levels of entitlement. Clicking on a system name takes you to its **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.



—If it is a child channel, you also have the option of unsubscribing systems from the channel. Use the checkboxes to select the systems, then click the **Unsubscribe** button on the bottom right-hand corner.

7.6.1.9.5. Software Channel Details ▯ Target Systems

List of entitled systems that are eligible for subscription to the channel. This tab appears only for child channels. Use the checkboxes to select the systems, then click the **Subscribe** button on the bottom right-hand corner. You will receive a success message or be notified of any errors. This can also be accomplished through the **Channels** tab of the **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.

7.6.2. Package Search

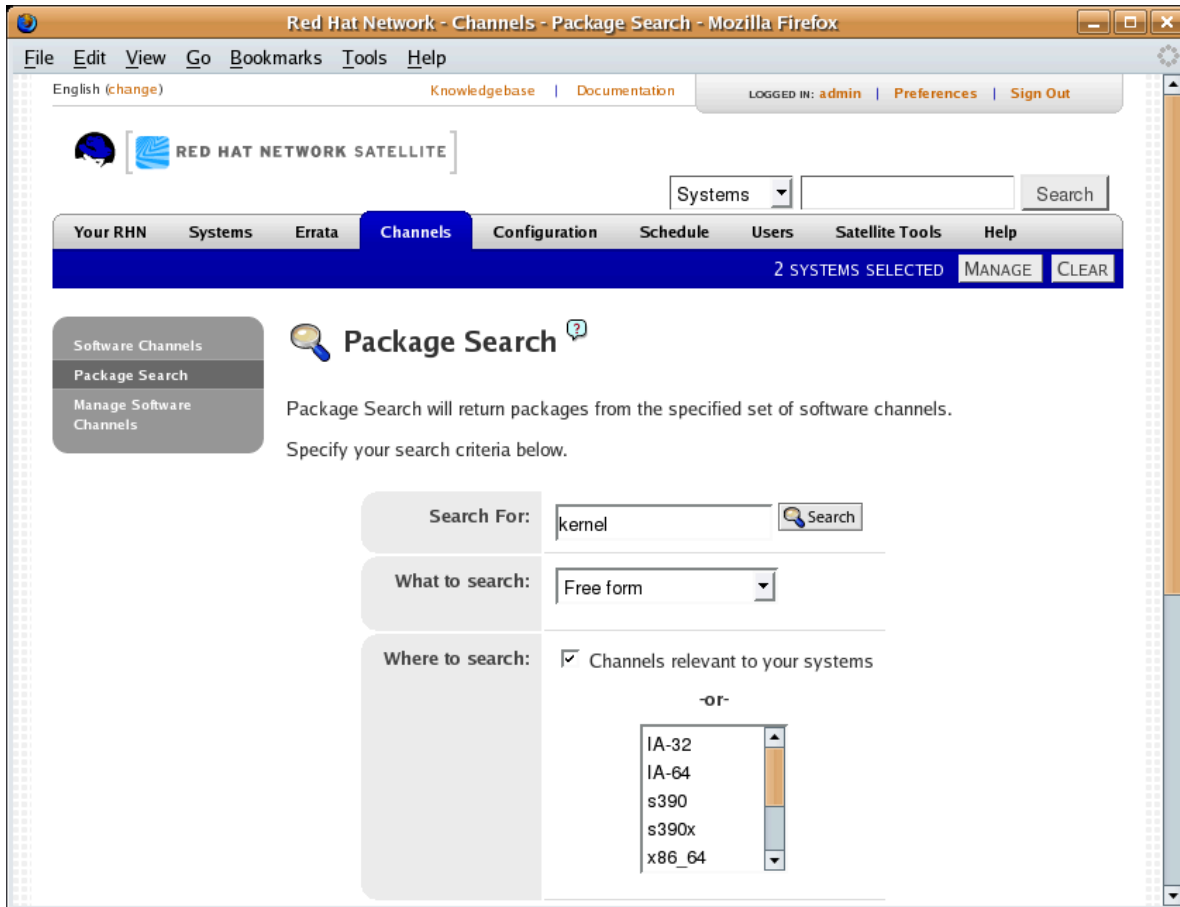


Figure 7.18. Package Search

The **Package Search** page allows you to search through packages using various criteria:

- **Free Form** — a general keyword search for users that are unsure of the details of particular package and its contents.
- **Name Only** — Targeted search for users that need to find a specific packages and do not want to sift through more generalized search results.
- **Name and Description** — Specified searches for a certain package name or program that, while not in the name of the package, may be in the one-line description of the package (for example, searching for the Apache HTTP Server when the actual Red Hat Enterprise Linux package name is `httpd`).
- **Name and Summary** — Similar to a **Name and Description** search, this search criteria searches package names and the longer Summary for the package. So, a search for "web browser" could result in several results that includes both graphical and text-based browsers.

The **Free Form** field additionally allows you to search using *field names* that you prepend to search queries and filter results by that field keyword.

For example, if you wanted to search all of the Red Hat Enterprise Linux v.5 packages for the word **java** in the description and summary, type the following using the **Free Form** field:

```
summary:java and description:java
```

Other supported field names for Documentation search include:

- **name** — Search the package names for a particular keyword
- **version** — Search for a particular package's version
- **filename** — Search the package filenames for a particular keyword
- **description** — Search the packages' detailed description field for a particular keyword
- **summary** — Search the packages' brief summary for a particular keyword
- **arch** — Search the packages by their architecture (such as x86, x86_64, or s390)

Along with search criteria, you can also limit searches to **Channels relevant to your systems** by clicking the checkbox.

Additionally, you can restrict your search by platform or architecture. Choices include **IA-32**, **IA-64**, **s390**, **s390x**, **x86_64**, **PPC**, **Sparc Solaris**, and **i386 Solaris**.

7.6.3. Manage Software Channels

This tab allows Administrators to create, clone, and delete custom channels. These channels may contain altered versions of distribution-based channels or custom packages.

7.6.3.1. Manage Software Channels ▢ Channel Details

The default screen of the Manage Software Channels tab is a listing of all available channels. This includes custom, distribution-based, and child channels.

To clone an existing channel, click the **clone channels** link in the upper right of the screen, select the channel to be cloned from the dropdown menu, and click the **Create Channel** button. The next screen presents various options for the new channel, including base architecture and GPG options. Make your selections and click the **Create Channel** button to complete the process.

To create a new channel, click the **create new channel** link in the upper right of the screen. Select the various options for the new channel, including base architecture and GPG options. Make your selections and click the **Create Channel** button. Note that a channel created in this manner is blank, containing no packages. You must either upload software packages or add packages from other channels. You may also choose to include Errata Updates in your custom channel.

7.6.3.1.1. Manage Software Channels ▢ Channel Details ▢ Channel Details

This screen lists the selections you made during the channel creation process. This page includes the **Globally Subscribable** checkbox that permits all users to subscribe to the channel.

7.6.3.1.2. Manage Software Channels ▢ Channel Details ▢ Managers

This subtab allows you to select which users may alter or delete this channel. Satellite Administrators and Channel Administrators may alter or delete any channel.

To allow a user to alter the channel, select the checkbox next to the user's name and click the **Update** button. To allow all users to manage the channel, click the **Select All** button at the bottom of the list followed by the **Update** button. To remove a user's ability to manage the channel, uncheck the box next to their name and click the **Update** button.

7.6.3.1.3. Manage Software Channels ▯ Channel Details ▯ Errata

This subtab allows channel managers to list, remove, clone, and add Errata to their custom channel. Custom channels not cloned from a distribution may not add Errata until there are packages in the channel. Only Errata that match the base architecture of the channel and apply to a package in that channel may be added to the channel. Finally, only cloned or custom Errata may be added to custom channels. Errata may be included in a cloned channel if they are selected during channel creation.

7.6.3.1.4. Manage Software Channels ▯ Channel Details ▯ Packages

This subtab is similar to the Errata subtab. It allows Channel and Organization Administrators to list, remove, compare, and add packages to the custom channel.

To list all packages in the channel, click the **List / Remove Packages** link. Check the box to the left of any package you wish to remove, then click the **Remove Packages** button in the lower right of the page.

To add packages, click the **Add Packages** link. Choose a channel from which to select packages from the drop-down menu and click the **View** button to continue. Check the box to the left of any package you wish to add to the channel, then click the **Add Packages** button in the bottom right of the screen.

To compare packages within the current channel with those of another channel, select the other channel from the drop-down menu and click the **Compare** button. All packages present in either channel are compared, and the results displayed on the next screen. This information includes the architecture and version of each package.

To make the two channels identical, click the **Merge Differences** button in the lower right. The following screen allows you to select how conflicts are resolved. Click the **Preview Merge** button to view the results of the merging without making any changes to the channels. Finally, select those packages that you wish to merge and click the **Merge Packages** button followed by the **Confirm** button to perform the merge.

7.6.3.2. Manage Software Channels ▯ Manage Software Packages

This tab allows you to manage custom software packages owned by your organization. You may view a list of all custom software or view only those packages in a selected custom channel. To select the channel whose custom packages you wish to view, select the channel from the drop-down menu and click the **View** button.

7.7. Configuration

This tab is the portal to managing your configuration channels and files, whether they are centrally managed or limited to a single system. You must be a Configuration Administrator or an Satellite Administrator to see the **Configuration** tab. In addition, you must have at least one Provisioning entitlement, or the tab does not appear.

Centrally-managed files are those that are available to multiple systems; changes to a single file in a central configuration channel can affect many systems. In addition, there are local configuration channels. Each system with a Provisioning entitlement has a local configuration channel (also referred

to as an override channel) and a Sandbox channel. Both central and local configuration management are discussed in detail later in this chapter.

7.7.1. Preparing Systems for Config Management

For a system to have its configuration managed through RHN, it must have the appropriate tools and **config-enable** file installed. These tools may already be installed on your system, especially if you kickstarted the system with configuration management functionality. If not, they can be found within the RHN Tools child channel for your distribution. Download and install the latest **rhncfg*** packages. They are:

- **rhncfg** — The base libraries and functions needed by all **rhncfg-*** packages.
- **rhncfg-actions** — The code required to run configuration actions scheduled via the RHN website.
- **rhncfg-client** — A command line interface to the client features of the RHN Configuration Management system.
- **rhncfg-management** — A command line interface used to manage RHN configuration.

Next, you must enable your system to schedule configuration actions. This is done using the **rhn-actions-control** command on the client system. This command is included in the **rhncfg-actions** RPM. The RHN Actions Control (**rhn-actions-control**) enables or disables specific modes of allowable actions. Refer to [Section B.1, “Red Hat Network Actions Control”](#) for instructions.

7.7.2. Overview

The **Configuration Overview** page allows you to assess at a glance the status of your configuration files and the systems that use them.

Configuration Summary

This panel provides quick reference information about your configuration files. Clicking on any of the blue text to the right displays an appropriate list of either relevant systems, channel details, or configuration files.

Configuration Actions

This panel offers direct access to the most common configuration management tasks. You can view or create files or channels, or enable configuration management on your systems.

Recently Modified Configuration Files

The list displayed here indicates which files have changed, to which channel they belong, and when they were changed. If no files have been recently changed, no list appears. Click on the name of the file to be taken to that file's **Details** page. Click on the channel name to be taken to the **Channel Details** page for that channel.

Recently Scheduled Configuration Deployments

Each action that has been scheduled is listed here along with the status of the action. Any configuration task that is scheduled, from enabling configuration management on a system to deploying a specific configuration file, is displayed here. This allows you to quickly assess if your tasks have succeeded, and to take action to correct any issues. Clicking on any blue text displays the **System Details** ▢ **Schedule** page for the specified system.

7.7.3. Configuration Channels

As mentioned above, RHN manages both central and local configuration channels and files. Central configuration management allows you to deploy configuration files to multiple systems. Local configuration management allows you to specify overrides, or configuration files that are not changed by subscribing the system to a central channel.

Central configuration channels must be created via the link on this page. Local configuration channels are not created here; they automatically exist for each system to which a Provisioning entitlement has been applied.

Click on the name of the configuration channel to be taken to the details page for that channel. If you click on the number of files in the channel, you are taken to the **List/Remove Files** page of that channel. If you click on the number of systems subscribed to the configuration channel, you are taken to the **Systems ▾ Subscribed Systems** page for that channel.

To create a new central configuration channel:

1. Click the **create new config channel** link in the upper right of this screen.
2. Enter a name for the channel.
3. Enter a label for the channel. This field must contain only alphanumeric characters, "-", "_", and "."
4. Enter a description for the channel. You must enter a description, though there is no character restriction. This field can contain any brief information that allows you to distinguish this channel from others.
5. Press the **Create Config Channel** button to create the new channel.
6. The following page is a subset of the **Channel Details** page, and has three sub-tabs: **Overview**, **Add Files**, and **Systems**. The Channel Details page is discussed fully in [Section 7.7.3.1, "Configuration ▾ Configuration Channels ▾ Configuration Channel Details"](#).

7.7.3.1. Configuration ▾ Configuration Channels ▾ Configuration Channel Details

Overview

This sub-tab is very similar to the **Configuration Overview** page. The **Channel Information** panel provides status information for the contents of the channel. The **Configuration Actions** panel provides access to the most common configuration tasks. The main difference is the **Channel Properties** panel. By clicking on the **Edit Properties** link, you can edit the name, label, and description of the channel.

List/Remove Files

This tab, which only appears if there are files in the configuration channel, lists the files that this configuration channel contains. You can remove a file or files, or copy the latest version into a set of local overrides or into other central configuration channels. Check the box next to any files you wish to manipulate and click the button corresponding to the desired action at the bottom of the screen.

Add Files

The **Add Files** sub-tab has three sub-tabs of its own, which allow you to **Upload**, **Import**, or **Create** configuration files to be included in the channel.

Upload File

To upload a file into the configuration channel, browse for the file on your local system, populate all fields, and click the **Upload Configuration File** button. The **Filename/Path** field is the absolute path where the file will be deployed. You can also indicate the ownership and permissions to be attached to the file when it is deployed. Finally, if the configuration file includes a macro, enter the symbol that marks the beginning and end of the macro.

Import Files

From this page you can import files from other configuration channels, including any locally-managed channels. Check the box to the left of any file you wish to import and press the **Import Configuration File(s)** button.



Note

A sandbox icon indicates that the listed file is currently located in a local sandbox channel. Files in a system's sandbox channel are considered experimental and could be unstable. Use caution when selecting them for a central configuration channel.

Create File

From this page you can create a configuration file from scratch to be included in the configuration channel. Indicate the absolute path along which the file should be deployed, enter the ownership and permissions for the file, and enter the configuration file content in the appropriate fields. Finally, press the **Create Configuration File** button to create the new file.

Deploy Files

This sub-tab only appears when there are files present in the channel. You can deploy all files by pressing the **Deploy All Files** button, or you can check selected files and press the **Deploy Selected Files** button. You will then be asked to select to which systems the file(s) should be applied. The listed systems are those that are subscribed to this channel. If you wish to apply the file to a system not listed here, first subscribe that system to the channel. When ready, press the **Confirm and Deploy to Selected Systems** button to deploy the files.

Systems

This tab, which consists of two sub-tabs, allows you to manage the systems that are subscribed to the configuration channel.

Subscribed Systems

This sub-tab displays a list of all systems that are subscribed to the current channel. Clicking on the name of the system takes you to the **System Details** page for that system.

Target Systems

This sub-tab displays a list of systems that have been enabled for configuration management and that are not yet subscribed to the channel. To add a system to the configuration channel, check the box to the left of the system's name and press the **Subscribe System** button.

7.7.4. Configuration Files

This tab allows you to manage your configuration files independently. Both centrally-managed and locally-managed files can be reached from sub-tabs.



Note

By default, the maximum file size for configuration files is 128KB. If you need to change that value, find and modify the following line in the `/etc/rhn/default/rhn_web.conf` file:

```
web.maximum_config_file_size=128
```

You must also find and change the following line in the `/etc/rhn/default/rhn_server.conf` file to the same value:

```
maximum_config_file_size=128
```

Change the value in both files from **128** to the desired value in kilobytes.

7.7.4.1. Centrally-Managed Files

Centrally-managed files are those that are available to multiple systems. Changing a file within a centrally-managed channel may result in changes to several systems.

This page lists all files that are currently stored in your central configuration channels. Click on the **Path** of a file to be taken to the **Configuration File Details** page for that file. Select the name of the configuration channel to be taken to the **Channel Details** page of the channel that contains the file. Clicking on the number of systems takes you to a listing of systems currently subscribed to the channel containing that file. Finally, clicking on the number of overriding systems displays a list of systems that have a local (or override) version of the configuration files (which means that the centrally-managed file will not be deployed to those systems.)

7.7.5. Locally-Managed Files

Locally-managed configuration files are those files that apply to only one system. They may be files in the system's sandbox or they may be files that can be deployed to the system at any time. Local files have higher priority than centrally-managed files - that is, if a system is subscribed to a configuration channel with a given file, and also has a locally-managed version of that same file, the locally-managed version is the one that will be deployed.

This page lists all of the local (override) configuration files for your systems. This includes the local configuration channels and the sandbox channel for each Provisioning-entitled system.

Click the **Path** of the file to go to the **Config File Details** page for the file. Click the name of the system to which it belongs to go to the **System Details** ▢ **Configuration** ▢ **Configuration** ▢ **Overview** page for the system.

7.7.5.1. Including Macros in your Configuration Files

Being able to store and share identical configurations is useful, but what if you have many variations of the same configuration file? What do you do if you have configuration files that differ only in system-specific details, such as hostname and MAC address?

In traditional file management, you would be required to upload and distribute each file separately, even if the distinction is nominal and the number of variations is in the hundreds or thousands. RHN addresses this by allowing the inclusion of macros, or variables, within the configuration files it manages for Provisioning-entitled systems. In addition to variables for custom system information, the following standard macros are supported:

- `rhn.system.sid`
- `rhn.system.profile_name`
- `rhn.system.description`
- `rhn.system.hostname`
- `rhn.system.ip_address`
- `rhn.system.custom_info(key_name)`
- `rhn.system.net_interface.ip_address(eth_device)`
- `rhn.system.net_interface.netmask(eth_device)`
- `rhn.system.net_interface.broadcast(eth_device)`
- `rhn.system.net_interface.hardware_address(eth_device)`
- `rhn.system.net_interface.driver_module(eth_device)`

To use this powerful feature, either upload or create a configuration file through the **Configuration Channel Details** page. Then, open its **Configuration File Details** page and include the supported macros of your choosing. Ensure that the delimiters used to offset your variables match those set in the **Macro Start Delimiter** and **Macro End Delimiter** fields and do not conflict with other characters in the file. We recommend that the delimiters be two characters in length and must not contain the percent (%) symbol.

As an example, you may have a file applicable to all of your servers that differs only in IP address and hostname. Rather than manage a separate configuration file for each server, you may create a single file, such as `server.conf`, with the IP address and hostname macros included, like so:

```
hostname={| rhn.system.hostname |}  
ip_address={| rhn.system.net_interface.ip_address(eth0) |}
```

Upon delivery of the file to individual systems, whether through a scheduled action in the RHN website or at the command line with the **Red Hat Network Configuration Client (rhncfg-client)**, the variables will be replaced with the hostname and IP address of the system, as recorded in RHN's System Profile. In the above configuration file, for example, the deployed version resembles the following:

```
hostname=test.example.domain.com  
ip_address=177.18.54.7
```

To capture custom system information, insert the key label into the custom information macro (`rhn.system.custom_info`). For instance, if you developed a key labeled "asset" you can add it to

the custom information macro in a configuration file to have the value substituted on any system containing it. The macro would look like this:

```
asset={@ rhn.system.custom_info(asset) @}
```

Upon deployment of the file to a system containing a value for that key, the macro gets translated, resulting in a string similar to the following:

```
asset=Example#456
```

To include a default value, for instance if one is required to prevent errors, you can append it to the custom information macro, like so:

```
asset={@ rhn.system.custom_info(asset) = 'Asset #' @}
```

This default is overridden by the value on any system containing it.

Using the **Red Hat Network Configuration Manager** (`rhncfg-manager`) will not translate or alter files, as that tool is system agnostic — `rhncfg-manager` does not depend on system settings. Binary files cannot be interpolated.

7.7.6. Systems

This page displays status information about your system in relation to configuration. There are two sub-tabs: **Managed Systems** and **Target Systems**.

7.7.6.1. Managed Systems

This page is the default display for the **Configuration** ▾ **Systems** page. The systems displayed here have been fully prepared for configuration file deployment. The number of local and centrally-managed files is displayed. Clicking the name of the system takes you to the **System Details** ▾ **Configuration** ▾ **Overview** page for the system. Clicking on the number of local files takes you to the **System Details** ▾ **Configuration** ▾ **View/Modify Files** ▾ **Locally-Managed Files** page, which allows you to manage which local (override) files apply to the system. Clicking on the number of centrally-managed files takes you to the **System Details** ▾ **Configuration** ▾ **Manage Configuration Channels** ▾ **List/Unsubscribe from Channels** page. This allows you to unsubscribe from any channels you wish.

7.7.6.2. Target Systems

This page displays the systems that are either not prepared for configuration file deployment or have not yet been subscribed to a configuration channel. The table has three columns which identify the system name, whether they are prepared for configuration file deployment, and a list of the steps that have yet to be completed before the system is prepared. By selecting the check box to the left of the profile name and then pressing the **Enable RHN Configuration Management** button, all of the preparatory steps that can be automatically performed are scheduled by RHN.



Note

You will still have to perform a few manual steps to enable configuration file deployment, but on-screen instructions are provided to assist with this step.

7.8. Schedule

If you click the **Schedule** tab on the top navigation bar, the **Schedule** category and links appear. These pages enable you to track the actions taking place within your systems. An action is a scheduled RHN task that is to be performed on one or more client systems. For example, an action can be scheduled to apply all Errata Updates to a system.

Red Hat Network keeps track of the following action types:

1. Package Alteration (installation, upgrade, and removal)
2. Rollback Package Actions
3. System Reboots
4. Errata Updates
5. Configuration File Alteration (deploy, upload, and diff)
6. Hardware Profile Updates
7. Package List Profile Updates
8. Kickstart Initiation
9. Remote Commands

Each page in the **Schedule** category represents an action status.

7.8.1. Pending Actions

As shown in [Figure 7.19, “Schedule - Pending Actions”](#), the **Pending Actions** page is shown by default when you click **Schedule** in the top navigation bar. It displays actions that have not started or are in progress.

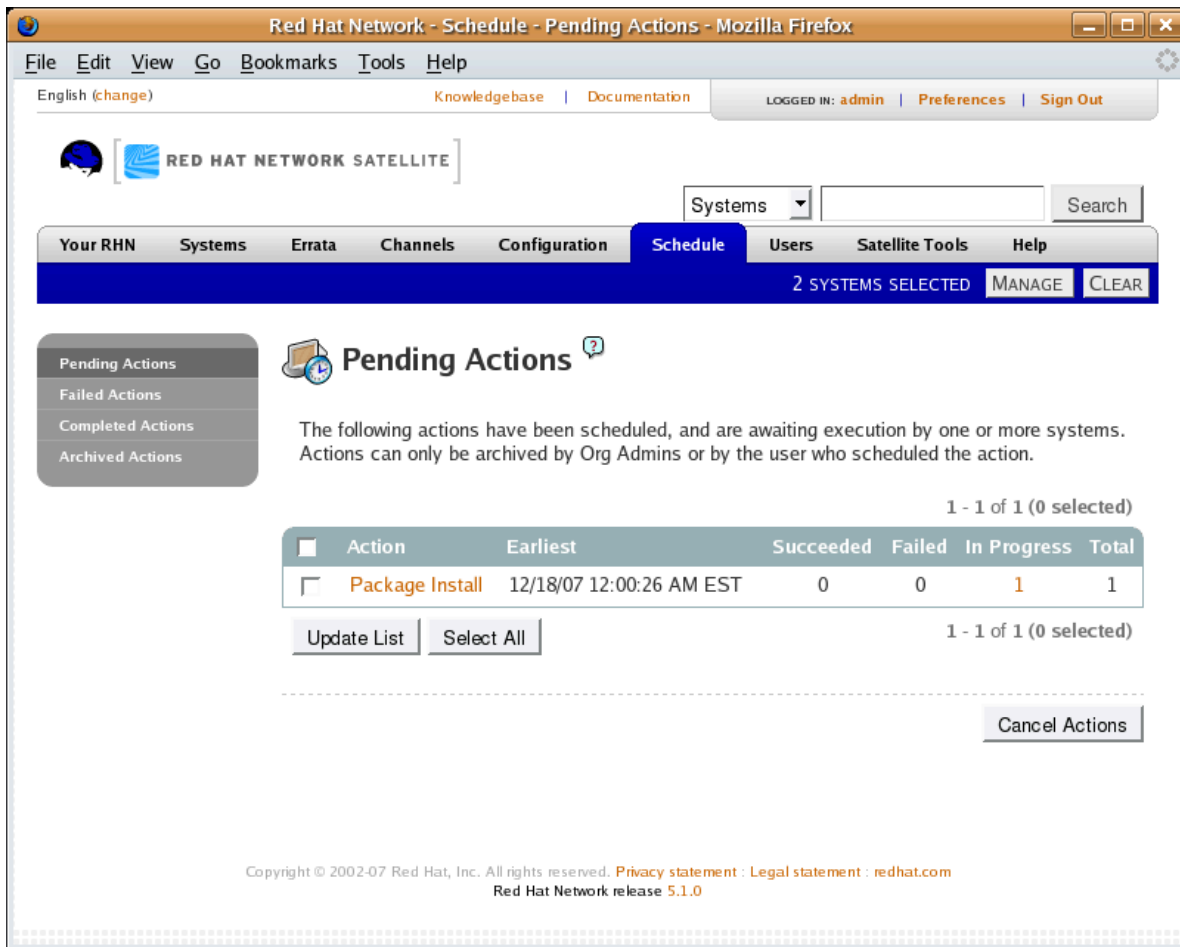


Figure 7.19. Schedule - Pending Actions

7.8.2. Failed Actions

Actions that could not be completed. If the action returns an error, it is displayed here.

7.8.3. Completed Actions

Actions that have succeeded.

7.8.4. Archived Actions

Actions that you have selected to store for review.

7.8.5. Actions List

In each page, each row in the list represents a single scheduled event or action that might affect multiple systems and involve various packages. The list contains several columns of information:

- **Select** — Use the checkboxes in this column to select actions. After selecting actions, you can either add them to your selection list or move them to the **Archived Actions** list. If you archive a pending action, it is not canceled; the action item moves from the **Pending Actions** list to the **Archived Actions** list.

- **Action** — Type of action to perform such as Errata Update or Package Install. Clicking an action name takes you to its **Action Details** page. Refer to [Section 7.8.5.1, “Action Details”](#) for more information.
- **Earliest** — The earliest day and time the action will be performed.
- **Succeeded** — Number of systems on which this action was successful.
- **Failed** — Number of systems on which this action has been tried and failed.
- **In Progress** — Number of systems on which this action is taking place.
- **Total** — Total number of systems on which this action has been scheduled.

7.8.5.1. Action Details

If you click on the name of an action, the **Action Details** page appears. This page is broken down into the following tabs:

7.8.5.1.1. Action Details ▢ Details

General information about the action. This is the first tab you see when you click on an action. It displays the action type, scheduling administrator, earliest execution, and notes. Clicking the Errata Advisory takes you to the **Errata Details** page. The Errata Advisory appears only if the action is an Errata Update. Refer to [Section 7.5.2.2, “Errata Details”](#) for more information.

7.8.5.1.2. Action Details ▢ Completed Systems

List of systems on which the action has been successfully undertaken. Clicking a system name takes you to its **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.

7.8.5.1.3. Action Details ▢ In Progress Systems

List of systems on which the action is now being undertaken. To cancel an action, select the system using the appropriate checkbox and click the **Unschedule Action** button. Clicking a system name takes you to its **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.

7.8.5.1.4. Action Details ▢ Failed Systems

List of systems on which the action has been attempted and failed. The actions can be rescheduled here. Clicking a system name takes you to its **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for more information.

7.9. Users —

Only Satellite Administrators can see the **Users** tab on the top navigation bar. If you click the **Users** tab, the **Users** category and links appear. These pages enable you to grant and edit permissions for those who administer your system groups. Click in the **User List** to modify users within your organization.

To add new users to your organization, click the **create new user** link on the to right corner of the page. The next page is the **Create User** page. Carefully fill in each of the required values for the new user.

Once all fields are complete, select the **Create Login** button. RHN now sends an email to the specified address and redirects you to the **Users** ▢ **User List** page. If you wish to select permissions and options for the newly created user, select their name from the list. Doing so displays the **User Details** page for that user, which provides several subtabs of options from which to choose. Refer to

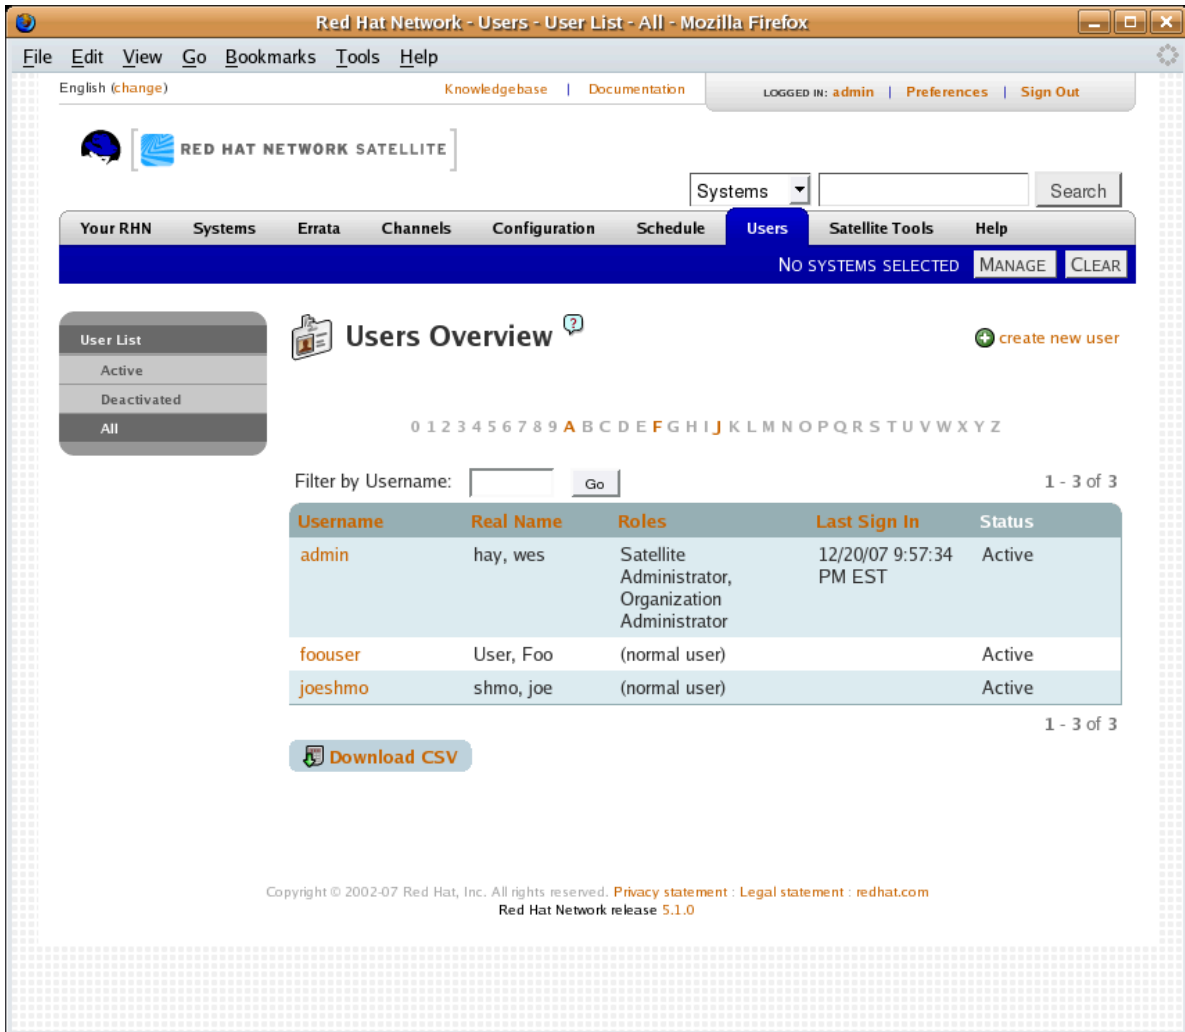
[Section 7.9.1.1, “User List ▢ Active ▢ User Details — !\[\]\(3dfb8d66e81160ad61421a3452093d1b_img.jpg\) ”](#) for detailed descriptions of each subtab.

7.9.1. User List ▢ Active —

This tab lists all active users of your RHN account. It displays the following basic information about each user: their username, real name, roles, and the date of their last sign in.

As shown in [Figure 7.20, “User List”](#), each row in the **User List** represents a user within your organization. There are four columns of information for each user:

- **Username** — The login name of the user. If you click on a username, the **User Details** page for the user is displayed. Refer to [Section 7.9.1.1, “User List ▢ Active ▢ User Details — !\[\]\(4e333a6106fc298d0ae6dff272a736ef_img.jpg\) ”](#) for more information.
- **Real Name** — The full name of the user (last name first).
- **Roles** — List of the user's privileges, such as Organization Administrator, Channel Administrator and normal user. Users can have multiple roles.
- **Last Sign In** — Shows when the user last logged into RHN.



The screenshot shows the Red Hat Network Users interface in a Mozilla Firefox browser window. The page title is "Red Hat Network - Users - User List - All". The interface includes a navigation menu with options like "Your RHN", "Systems", "Errata", "Channels", "Configuration", "Schedule", "Users", "Satellite Tools", and "Help". The "Users" tab is selected, and the "Users Overview" section is active. A sidebar on the left shows "User List" with options for "Active", "Deactivated", and "All". The main content area displays a table of users with columns for Username, Real Name, Roles, Last Sign In, and Status. The table lists three users: admin, fouser, and joeshmo. A "Download CSV" button is located below the table. The footer contains copyright information for Red Hat, Inc. and links to privacy and legal statements.

Username	Real Name	Roles	Last Sign In	Status
admin	hay, wes	Satellite Administrator, Organization Administrator	12/20/07 9:57:34 PM EST	Active
fouser	User, Foo	(normal user)		Active
joeshmo	shmo, joe	(normal user)		Active

Figure 7.20. User List

7.9.1.1. User List ▢ Active ▢ User Details —

The **User Details** page allows Satellite Administrators to manage the permissions and activity of all users. Included in the **User Details** page is the ability to delete or deactivate users.

Users may now be deactivated directly from the RHN web interface. RHN Satellite customers may deactivate or delete users from their systems, although non-Satellite customers must contact Customer Service to delete a user. Users may be deactivated or deleted by Satellite Administrators, or users may deactivate their own accounts.

Deactivated users cannot log in to the RHN web interface, nor may they schedule any actions. Satellite Administrators may not be deactivated until that role is removed from their account. Actions scheduled by a user prior to their deactivation remain in the action queue. For added flexibility, deactivated users may be reactivated by Satellite Administrators.

User deletion from the web interface is available exclusively to RHN Satellite customers. The Satellite Administrator role must be removed from a user before that individual may be deleted.



Warning

User deletion is irreversible; exercise it with caution. Consider disabling the user first in order to assess the effect deletion will have on your infrastructure.

To deactivate a user:

1. Navigate to the user's **User Details** tab.
2. Verify that the user is not an Satellite Administrator. If they are, uncheck the box to the left of that role and click the **Submit** button in the lower right of the screen.
3. Click the **deactivate user** link in the upper right of the screen.
4. Click the **Deactivate User** button in the lower right to confirm.

To delete a user:

1. Navigate to the user's **User Details** tab.
2. Verify that the user is not an Satellite Administrator and remove that role if necessary.
3. Click the **delete user** link in the upper right.
4. Click the **Delete User** button to permanently delete the user.

For instructions regarding deactivating your own account, refer to [Section 7.3.1.3, “Account Deactivation”](#).

7.9.1.1.1. User List ▾ Active ▾ User Details ▾ Details —

This is the default **User Details** tab, which displays the username, first name, last name, email address, and user roles for the user. All of this information is modifiable. To do so, make your changes and click the **Update** button. Remember, when changing a user's password, you will see only asterisks as you type the password.

To delegate responsibilities within your organization, Red Hat Network provides several roles with varying degrees of responsibility and access. This list describes the permissions of each and the differences between them:

- **User** — Also known as a *System Group User*, this is the standard role associated with any newly created user. This person may be granted access to manage system groups and software channels. The systems must be in system groups to which the user has permissions for them to be manageable or even visible. Remember, however, all globally subscribable channels may be used by anyone.
- **Activation Key Administrator** — This role is designed to manage your organization's collection of activation keys. This person can create, modify, and delete any key within your overarching account.
- **Channel Administrator** — This role has complete access to the software channels and related associations within your organization. It requires RHN Satellite or RHN Proxy Server. This person may change the base channels of systems, make channels globally subscribable, and create entirely new channels.

- **Configuration Administrator** — This role enables the user to manage the configuration of systems in the organization using either the RHN Satellite web-based interface or the **Red Hat Network Configuration Manager**.
- **Monitoring Administrator** — This role allows for the scheduling of probes and oversight of other Monitoring infrastructure. This role is available only on Monitoring-enabled RHN Satellite version 3.6 or later.
- **Satellite Administrator** — This role can perform any function available within Red Hat Network. As the master account for your organization, the person holding this role can alter the privileges of all other accounts, as well as conduct any of the tasks available to the other roles. Like the other roles, multiple Satellite Administrators may exist.
- **System Group Administrator** — This role is one step below Satellite Administrator in that it has complete authority over the systems and system groups to which it is granted access. This person can create new system groups, delete any assigned systems groups, add systems to groups, and manage user access to groups.

While it is possible for one Satellite Administrator to remove Satellite Administrator rights from another user, it is impossible to remove Satellite Administrator rights from the sole remaining Satellite Administrator. It is possible to remove your own Satellite Administrator privileges so long as you are not the last Satellite Administrator.

To assign a user a new role, select the appropriate checkbox. Remember that Satellite Administrators are automatically granted administration access to all other roles, signified by grayed-out checkboxes. To grant a user the ability to manage the configuration of systems, select the **Configuration Administrator** checkbox. When satisfied with the changes, click **Update**.

7.9.1.1.2. User List ▾ Active ▾ User Details ▾ System Groups —

This tab displays a list of system groups that the user may administer. Satellite Administrators may use the check boxes to set this user's access permissions to each system group. Check or uncheck the box to the left of the system group and click the **Update Permissions** button to save the changes.

Satellite Administrators may select one or more default system groups for this user. When the user registers a system, that system is assigned to the selected group or groups. This allows the user to have access to the newly-registered system immediately, if he or she has permissions to one or more of the groups to which the system is assigned. System Groups to which this user has access are preceded by an (*).

7.9.1.1.3. User List ▾ Active ▾ User Details ▾ Systems —

This tab lists all systems to which the user has access permission. These systems come from the system groups assigned to the user on the previous tab. You may choose a set of systems to work with by checking the boxes to the left of the systems and clicking the **Update List** button. Use the System Set Manager page to execute actions on those systems. Clicking the name of a system takes you to its **System Details** page. Refer to [Section 7.4.2.9, "System Details"](#) for more information.

7.9.1.1.4. User List ▢ Active ▢ User Details ▢ Channel Permissions —

This tab lists all channels available to your organization. You may grant explicit channel subscription permission to this user for each of the channels listed by checking the box to the left of the channel and clicking the **Update Permissions** button. Permissions granted through Satellite Administrator status, certificate authority status, or because the channel is globally subscribable have no checkbox, but display a check icon instead.

7.9.1.1.4.1. User List ▢ Active ▢ User Details ▢ Channel Permissions ▢ Subscription —



Identifies channels to which the user may subscribe systems. To change these, select or unselect the appropriate checkboxes and click the **Update Permissions** button. Note that channels subscribable through the user's admin status or the channel's global setting cannot be altered. They are identified with a check icon.

7.9.1.1.4.2. User List ▢ Active ▢ User Details ▢ Channel Permissions ▢ Management —



Identifies channels the user may manage. To change these, select or unselect the appropriate checkboxes and click the **Update Permissions** button. This status does not enable the user to create new channels. Note that channels automatically manageable through the user's admin status cannot be altered. They are identified with a check icon. Remember, Satellite Administrators and Channel Administrators can subscribe to or manage any channel.

7.9.1.1.5. User List ▢ Active ▢ User Details ▢ Preferences —

This page allows you to configure whether the user receives email notifications, the number of entries displayed per list page, and the timezone of the user. Make selections and click the **Save Preferences** button to update.

- **Email Notification** — Determine whether this user should receive email every time an Errata Alert is applicable to one or more systems in his or her RHN account, as well as daily summaries of system events.
- **RHN List Page Size** — Maximum number of items that appear in a list on a single page. If more items are in the list, clicking the **Next** button displays the next group of items. This preference applies to the user's view of system lists, Errata lists, package lists, and so on.
- **Time Zone** — Set this user's time zone so that scheduled actions are arranged according to the time in the relevant time zone.
- **Red Hat Contact Options** — Identify what ways (email, phone, fax, or mail) Red Hat may contact the user.

To modify any of these options, make your changes and click the **Save Preferences** button.

7.9.1.1.6. User List ▢ Active ▢ User Details ▢ Addresses —

This tab lists the addresses associated with the user's account. To update this information, click the appropriate **Edit this address** link, enter the relevant information, and click the **Update** button.

7.9.1.1.7. User List ▢ Active ▢ User Details ▢ Notification Methods —

This tab lists email and pager addresses designated to receive alerts from Monitoring probes. To create a method, click **create new method** and complete the fields. If you will receive these alerts via pager, select the associated checkbox to have the messages sent in a shorter format. When finished, click **Create Method**. The method shows up in the Methods list, from which it can be edited and deleted.

You may delete notification methods here, as well. If the notification method has probes attached to it, you are presented with a list of the probes. Note that if you are a Monitoring Administrator and cannot manage the system in question, the **System Details** and probe's **Current State** page are not accessible via links in their names. As always, Satellite Administrators have full access to all aspects of your RHN account.

7.9.2. User List ▢ Deactivated —

This page lists all users who have been deactivated. To reactivate any of the users listed here, click the check box to the left of their name and click the **Reactivate** button followed by the **Confirm** button. Reactivated users retain the permissions and system group associations they had when they were deactivated. Clicking on the User Name of any individual takes you to their User Details page.

7.9.3. User List ▢ All —

The **All** page lists all users that belong to your organization. In addition to the fields listed in the previous two screens, the table of users includes a **Status** field. This field indicates whether the user is **Active** or **Deactivated**. Deactivated users are also grayed out to indicate their status. Click on the username to move to the user's **User Details** page.

7.10. Monitoring —

If you click the **Monitoring** tab on the top navigation bar, the **Monitoring** category and links appear. These pages, which require Monitoring entitlements, enable you to view the results of probes you have set to run against Monitoring-entitled systems and manage the configuration of your monitoring infrastructure.






Initiate monitoring of a system through the **Probes** tab of the **System Details** page. Refer to [Section 7.4.2.9, "System Details"](#) for a description of the tab. See [Appendix D, Probes](#) for the complete list of available probes.

7.10.1. Probe Status —

The **Probe Status** page is shown by default when you click **Monitoring** in the top navigation bar.

The **Probe Status** page displays the summary count of probes in the various states and provides a simple interface to find problematic probes quickly. Please note that the probe totals in the tabs at the top of the page may not match the numbers of probes displayed in the tables below. The counts at the top include probes for all systems in your organization, while the tables display probes on only those systems to which you have access through the System Group Administrator role. Also, the probe counts displayed here may be out of sync by as much as one minute.

The following list describes each state and identifies the icons associated with them:

-  — *Critical* - The probe has crossed a CRITICAL threshold.
-  — *Warning* - The probe has crossed a WARNING threshold.
-  — *Unknown* - The probe is not able to accurately report metric or state data.
-  — *Pending* - The probe has been scheduled but has not yet run or is unable to run.
-  — *OK* - The probe is running successfully.

The **Probe Status** page contains tabs for each of the possible states, as well as one that lists all probes. Each table contains columns indicating probe state, the monitored system, the probes used, and the date and time the status was last updated.

In these tables, clicking the name of the system takes you to the **Probes** tab of the **System Details** page. Clicking the name of the probe takes you to its **Current State** page. From there, you may edit the probe, delete it, and generate reports based upon its results.

Monitoring data and probe status information that was previously available only through the web interface of the Satellite can now be exported as a CSV file. Click on the **Download CSV** links throughout the Monitoring pages to download CSV files of relevant information. The exported data may include, but is not limited to:

- Probe status
- All probes in a given state (OK, WARN, UNKNOWN, CRITICAL, PENDING)
- A Probe Event history

7.10.1.1. Probe Status Critical —

The probes that have crossed their CRITICAL thresholds or reached a critical status by some other means. For instance, some probes become critical (rather than unknown) when exceeding their timeout period.

7.10.1.2. Probe Status Warning —

The probes that have crossed their WARNING thresholds.

7.10.1.3. Probe Status Unknown —

The probes that cannot collect the metrics needed to determine probe state. Most but not all probes enter an unknown state when exceeding their timeout period. This may mean that the timeout period should be increased, or the connection cannot be established to the monitored system.

It is also possible the probes' configuration parameters are not correct and their data cannot be found. Finally, this state may indicate that a software error has occurred.

7.10.1.4. Probe Status Pending —

The probes whose data have not been received by RHN. This state is expected for a probe that has just been scheduled but has not yet run. If all probes go into a pending state, your monitoring infrastructure may be failing.

7.10.1.5. Probe Status OK —

The probes that have run successfully without exception. This is the state desired for all probes.

7.10.1.6. Probe Status All —

All probes scheduled on systems in your account, listed in alphabetical order by the name of system.

7.10.1.7. Current State —


Identifies the selected probe's status and when it last ran, while providing the ability to generate a report on the probe. Although this page is integral to monitoring, it is found under the **Probes** tab within the **System Details** page since its configuration is specific to the system being monitored.

To view a report of the probe's results, choose a relevant duration using the **date** fields and decide whether you would like to see metric data, the state change history or both. To obtain metric data, select the metric(s) on which you wish to see a report, and decide (using the checkboxes) whether the results should be shown in a graph, an event log, or both. Then click **Generate report** at the bottom of the page. If no data exist for the probe's metrics, you are presented with the following message: **NO DATA SELECTED TIME PERIOD AND METRIC.**

7.10.2. Notification —

Identifies the contact methods that have been established for your organization. These methods contain email or pager addresses designated to receive alerts from probes.

The various notification methods available to your organization are listed here on the default **Notification** screen. The methods are listed according to the user to which they apply.

To create a new notification method, click on the name of the user to whom the notification will apply. The user's User Details  Notification Methods page appears. Refer to [Section 7.9.1.1.7, "User List](#)

[Active](#)  [User Details](#)  [Notification Methods](#) —  " for further information. Click on the title of the notification method to edit the properties of the method.

7.10.2.1. Notification Filters

Notification filters allow you to create long-term rules that suspend, redirect, or automatically acknowledge standard notifications or send supplemental notifications. This can be helpful in managing verbose or frequent probe communication.

7.10.2.1.1. Notification Notification Filters Active Filters

This is the default screen for the Notification Filters tab. It lists all active filters available for your organization. Click the name of the filter to edit the properties of the filter.

To create a notification filter, click the **create new notification filter** link in the upper right of the screen. Configure each option listed below and click the **Save Filter** button to create the filter.

1. *Description*: Enter a value that allows you to distinguish this filter from others.
2. *Type*: Determine what action the filter should take: redirect, acknowledge, suspend, or supplement the incoming notification.
3. *Send to*: The **Redirect Notification** and **Supplemental Notification** options in step two require an email address to which to send the notifications. The remaining options require no email address.
4. *Scope*: Determine which monitoring components are subject to the filter.
5. *Organization/Scout/Probe*: This option allows you to select the organization, scout(s), or probe(s) to which this filter applies. To select multiple items from the list, hold the **Ctrl** key while clicking the names of the items. To select a range of items, hold the **Shift** key while clicking on the first and last items in the range.
6. *Probes in State*: Select which probe state(s) relate to the filter. For example, you may choose to create a supplemental notification for critical probes only. Un-check the box to the left of any state you want the filter to ignore.
7. *Notifications sent to*: This is the method to which the notification would be sent if no filter were in place. You may, for example, redirect notifications that would normally go to a user should that individual go on vacation, leaving all other notifications from the probe unchanged.
8. *Match Output*: Select precise notification results by entering a regular expression here. If the "Message:" portion of the notification does not match the regular expression, the filter is not applied.
9. *Recurring*: Select whether a filter runs continuously or on a recurring basis. A recurring filter runs multiple times for a period of time smaller than the duration of the filter. For example, a recurring filter could run for 10 minutes of every hour between the start and end times of the filter. A non-recurring filter runs continuously between the start and end times of the filter.
10. *Beginning*: Enter a date and time for the filter to begin operation.
11. *Ending*: Enter an end date and time for the filter.
12. *Recurring Duration*: How long a recurring filter instance is active. This field, applicable to recurring filters only, begins at the **Beginning** time specified above. Any notification generated outside of the specified duration is not filtered.
13. *Recurring Frequency*: How often the filter activates.

Notification filters cannot be deleted. However, a filter may be canceled by setting the end date to some time in the past. (Note that the end date must be equal to or later than the start date, or the change fails.) Another method is to select a set of filters from the **Active** page and click the **Expire Notification Filters** button in the lower right. These filters are then canceled and appears in the **Expired Filters** tab.

7.10.2.1.2. Notification ▯ Notification Filters ▯ Expired Filters

This tab lists all notification filters whose end date has passed. Expired filters are stored indefinitely; this allows an organization to recycle useful filters as needed and provides a historical record for troubleshooting.

7.10.3. Probe Suites

Probe Suites allow you to configure and apply one or more probes to a system or systems. Probe Suites may be configured once and then applied to any number of systems in a batch. This results in time savings and consistency for Monitoring customers.

To create and apply a Probe Suite, first create an empty Probe Suite, then configure member probes, and finally apply the Suite to selected systems.

1. From the Monitoring ▯ Probe Suites page, select the **create probe suite** link. Enter an easily distinguishable name for the Probe Suite. You may also choose to add a brief description of the Suite. Click the **Create Probe Suite** button to continue.
2. Add and configure the probes that comprise the Suite. Click the **create new probe** link in the upper right.
3. As described in [Section 7.4.2.9.5.2, “System Details ▯ Monitoring — !\[\]\(67ff022fd78f943b679992c2874bbfd1_img.jpg\)”](#), configure the probe and click the **Create Probe** button in the lower right. Repeat this process until all desired probes have been added.



Note

Sendmail must be configured correctly on your RHN Satellite and each client system to which the Probe Suite is applied must have the **rhnmmd** daemon installed and running. Refer to the *RHN Satellite 5.2.0 Installation Guide* for additional information.

4. Add the systems to which the Probe Suite applies. Click the **add systems to probe suite** link in the upper right of the screen to continue.
5. The next page displays a list of all systems with Monitoring entitlements. Check the box to the left of the system(s) to which you wish to apply the Probe Suite, select the monitoring scout you wish to use, and click the **Add systems to probe suite** button to complete the creation of the Probe Suite.

You can either delete or detach probes from the suite. Detaching a probe disassociates the probes from the suite and converts them to system-specific probes for the specified system. This means that changes to the detached probes only effect that system. Deleting a probe removes it from the Suite for all systems.

To remove probes from the Probe Suite:

1. From the Monitoring ▢ Probe Suites page, click on the title of the Probe Suite you wish to alter.
2. Select the **Probes** sub-tab.
3. Check the box next to the probe you wish to remove.
4. Click the **Delete probes from Probe Suites** button.

You may also remove a system from the Probe Suite. There are two ways to accomplish this. The first method is to detach the system from the Probe Suite. When you do so, the system still has the same probes assigned to it. However, you now have the ability to configure these probes individually without affecting any other systems. For more information about removing probes from an individual system,

refer to [Section 7.4.2.9.5.2, “System Details ▢ Monitoring — !\[\]\(dfbd6b3763a6d1d9afaa974f64e2e4b5_img.jpg\)”](#).

To detach a system from the suite:

1. From the **Monitoring ▢ Probe Suites** page, click on the title of the Probe Suite you wish to alter.
2. Select the **Systems** sub-tab.
3. Check the box next to the system(s) you wish to remove from the Probe Suite.
4. Click the **Detach System(s) from Probe Suite** button

The second method is to remove the system from the suite. This removes the system from the suite and deletes all running probes from the system.



Note

This action deletes all of the Probe Suites' probes from the system as well as all of the historical Time Series and Event Log data. This action is irreversible.

To remove a system from the Probe Suite and delete all associated probes from the system:

1. From the Monitoring ▢ Probe Suites page, click on the title of the Probe Suite you wish to alter.
2. Select the **Systems** sub-tab.
3. Check the box next to the system(s) you wish to remove from the Probe Suite.
4. Click the **Remove System(s) from Probe Suite** button.

Finally, as with single Probes, you may download a CSV file containing information about Probe Suites. Click the **Download CSV** link at the bottom of the **Monitoring ▢ Probe Suites** page to download the file.

7.10.4. Scout Config Push —

Displays the status of your monitoring infrastructure. Anytime you make a change to your monitoring configuration, such as adding a probe to a system or editing a probe's thresholds, you must reconfigure your monitoring infrastructure. Do this by selecting the RHN Server's checkbox and clicking **Push Scout Configs**. The table on this page identifies the date and time of requested and completed pushes.

Clicking the name of the server opens its Red Hat Network Monitoring Daemon SSH Public Key. This allows you to copy and paste the SSH key to the systems that are monitored by the scout. This is required in order for the Red Hat Network Monitoring Daemon to connect to the Satellite.

7.10.5. General Config —

Collects information that is universally applicable to your Monitoring infrastructure. Modifying anything on this page causes the Monitoring services on the RHN Satellite to reset. It also schedules restart events for the Monitoring services on all Monitoring-enabled RHN Proxy Servers that connect to this Satellite. This is done so that the Monitoring services on these servers immediately reload their configuration.

Typically, the defaults provided in other fields are acceptable, since they are derived from your Satellite installation. Nevertheless, you may use the fields on this page to alter your Monitoring configuration. For instance, you may change your mail exchange server here. This page also allows you to alter the destination of all administrative emails from the Satellite. When finished, click **Update Config**.

7.11. Admin

The **Admin** page allows RHN Satellite customers to manage the basic configuration of the Satellite, including creating and managing the Organizations feature of RHN Satellite. Only the Satellite Administrator can access the **Admin** page.

7.11.1. Admin ▾ Organizations

The *multiple organizations* feature allows administrators to create and manage multiple organizations across the Satellite. The Organizations feature allows administrators to appropriate software and system entitlements across various organizations, as well as control an organization's access to systems management tasks. For more information about using the multiple organizations feature, refer to [Chapter 9, Multiple Organizations](#).

7.11.2. Admin ▾ RHN Satellite Configuration

This tab is broken down into subtabs that allow you to configure most aspects of the RHN Satellite. Once changes have been made, it is important to restart the Satellite, which may be accomplished on the final tab.

7.11.2.1. Admin ▾ Satellite Configuration ▾ General

The **Satellite Configuration ▾ General Configuration** page allows you to alter the most basic Satellite settings, such as the admin email address and whether Monitoring is enabled.

7.11.2.2. Admin ▾ Satellite Configuration ▾ Monitoring

The **RHN Satellite Configuration ▾ Monitoring** page allows you to configure the monitoring aspects of this Satellite. The local mail exchanger and local main domain are used to mail monitoring notification messages to administration. This is required only if you intend to receive alert notifications from probes. If you do, provide the mail server (exchanger) and domain to be used. Note that **sendmail** must be configured to handle email redirects of notifications. When finished, click **Update Config**.

7.11.2.3. Admin ▢ Satellite Configuration ▢ Certificate

The **RHN Satellite Configuration ▢ Certificate** page allows you to either upload a new Satellite certificate. To identify the certificate's path, click **Browse**, navigate to the file, and select it. To input its contents, open your certificate in a text editor, copy all lines, and paste them directly into the large text field at the bottom. Red Hat recommends using the file browser as it is less error prone. Click **Update** to continue. If you receive errors related to DNS, ensure your Satellite is configured correctly.

7.11.2.4. Admin ▢ Satellite Configuration ▢ Bootstrap Script

The **RHN Satellite Configuration ▢ Bootstrap** page allows you to generate a bootstrap script for redirecting client systems from the central RHN Servers to the Satellite. This script, to be placed in the `/var/www/html/pub/bootstrap/` directory of the Satellite, significantly reduces the effort involved in reconfiguring all systems, which by default obtain packages from the central RHN Servers. The required fields are pre-populated with values derived from previous installation steps. Ensure this information is accurate.

Checkboxes offer options for including built-in security SSL and GNU Privacy Guard (GPG) features, both of which are advised. In addition, you may enable remote command acceptance and remote configuration management of the systems to be bootstrapped here. Both features are useful for completing client configuration. Finally, if you are using an HTTP proxy server, complete the related fields. When finished, click **Generate Bootstrap Script**.

7.11.2.5. Admin ▢ Satellite Configuration ▢ Organizations

The **RHN Satellite Configuration ▢ Organizations** page contains details about the Organizations feature of RHN Satellite, as well as links to quickly get started creating and configuring organizations. For more information about configuring Organizations, refer to [Section 7.11.1, "Admin ▢ Organizations"](#).

7.11.2.6. Admin ▢ Satellite Configuration ▢ Restart

The **RHN Satellite Configuration ▢ Restart** page contains the final step in configuring the Satellite. Click the **Restart** button to restart the Satellite in order to incorporate all of the configuration options added on the previous screens. Note that it will take between four and five minutes for the restart to finish.

7.12. Help

The **Help** pages provide access to the full suite of documentation and support available to RHN users. Click **Help** in the **Overview** category to see a list of options available to you.

7.12.1. Reference Guide

The **Reference Guide** page takes you to this same document, the most comprehensive set of instructions for using Red Hat Network. Note that links to other technical guides may also appear in the left navigation bar, depending on the entitlement level and product offering of the account with which you logged in.

7.12.2. Satellite Installation Guide

Implementing a fully functional RHN Satellite requires more than installing software and a database. Client systems must be configured to use the Satellite. Custom packages and channels should be

created for optimal use. Since these tasks extend beyond the basic installation, they are covered in detail in other guides, as well as this *RHN Satellite Installation Guide*.

Detailed information regarding RHN Satellite server and its installation and initial configuration.

7.12.3. Proxy Guide

RHN Proxy Server is a package-caching mechanism that reduces the bandwidth requirements for RHN and Satellite servers and enables custom package deployment. Proxy customers cache RPMs, such as Errata Updates from Red Hat or custom RPMs generated by their organization, on an internal, centrally-located server. Client systems then receive these updates from the Proxy rather than by accessing the Internet individually.

The *RHN Proxy Server Installation Guide* provides detailed information regarding RHN Proxy server installation and initial configuration.

7.12.4. Client Configuration Guide

By default, all Red Hat Network client applications are configured to communicate with central Red Hat Network Servers. When connecting clients to RHN Satellite or RHN Proxy Server instead, many of these settings must be altered. Altering client settings for a system or two may be relatively simple. A large enterprise environment, containing hundreds or thousands of systems, will likely benefit from the mass reconfiguration steps described here.

The *Client Configuration Guide* is a best practices manual intended to help customers of RHN Satellite and RHN Proxy Server configure their client systems efficiently.

7.12.5. Channel Management Guide

An Red Hat Network and RHN Satellite software channel is a collection of software packages. Channels help you segregate packages by sensible rules: That is, a channel may contain packages from a specific Red Hat Enterprise Linux version, for instance. A channel may contain packages for an application or family of applications. Users may also define channels for their own particular needs; a company may create a channel that contains packages for all of the organization's laptops, for example.

The *Channel Management Guide* documents the creation and maintenance of custom channels using RHN Satellite.

7.12.6. Release Notes

The **Release Notes** page lists the notes accompanying every recent release of Red Hat Network. These notes describe all significant changes occurring in a given release cycle, from major enhancements to the user interface to minor changes to the related documentation.

7.12.7. API

Documentation for using the Red Hat Network Application Programming Interface (API) for creating tools and programs to automate common tasks via Red Hat Network.

The **API** page contains an overview of the API, with links to detailed descriptions of various API calls available to administrators and developers. There is also an **FAQ** page for answers to common

questions about the Red Hat Network API. Finally, there is a **Sample Scripts** page that shows users example code using API calls.

7.12.8. Search

The **Documentation Search** page features a robust search engine that indexes and searches RHN Satellite and RHN Proxy Server documentation.

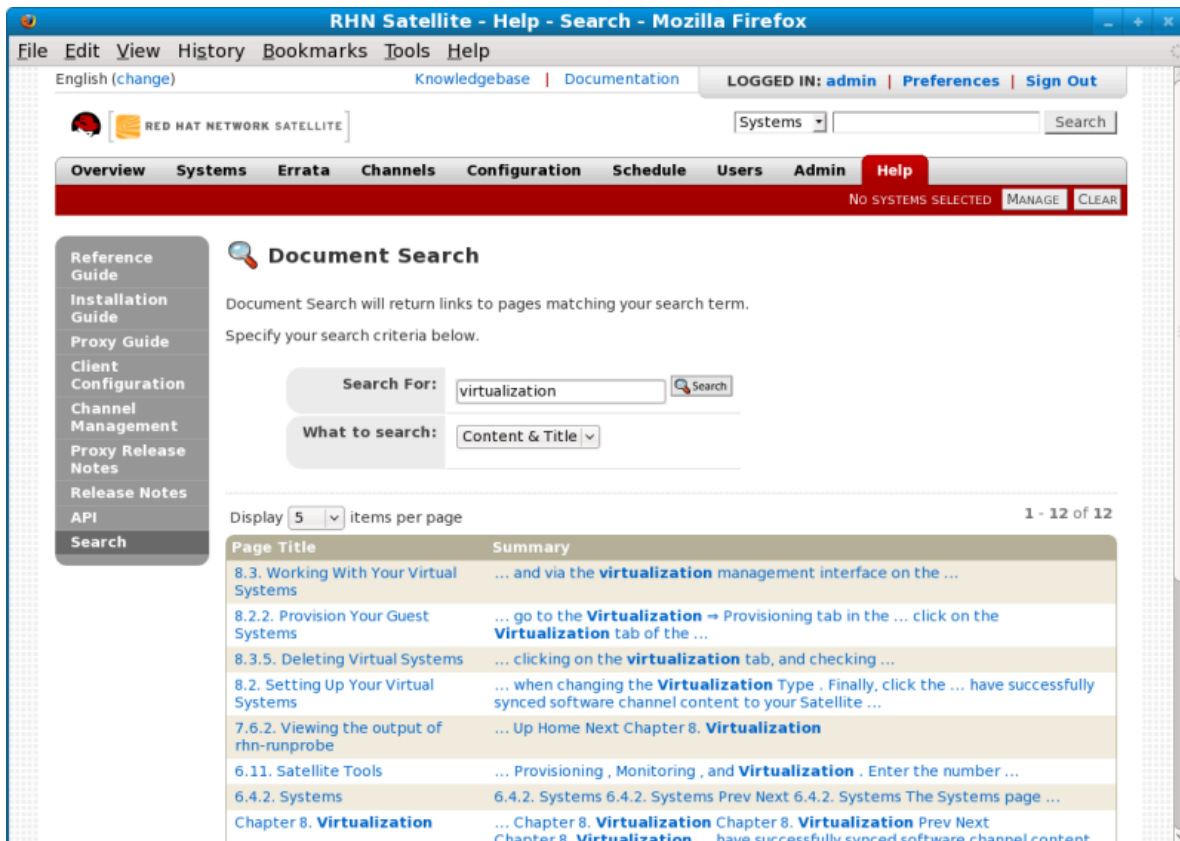



Figure 7.21. Documentation Search

Users can search the available online documentation and filter them according to the following choices in the **What to Search** drop-down menu:

- **Content & Title** — Search both the title heading or body content of all available documents
- **Free Form** — Search documents and indices for any keyword matches, which broadens search results.
- **Content** — Search only the body content of documentation for more specific matches
- **Title** — Search only the titles heading of the documentation for targeted, specific search results.

The **Free Form** field additionally allows you to search using *field names* that you prepend to search queries and filter results in that field.

For example, if you wanted to search all of the Satellite manuals for the word **Virtualization** in the title and **kickstart** in the content, type the following in the **Free Form** field:




```
title:Virtualization and content:kickstart
```

Other supported field names for Documentation search include:

- **url** — Search the URL for a particular keyword
- **title** — Search titles for a particular keyword
- **content** — Search the body of the documentation for a particular keyword

If there are several pages of search results, you can limit the amount of visible results shown on one page by clicking the **Display *quantity* items per page** drop-down menu, which offers between 10 and 500 results per page.

To move between pages, click the right or left angle brackets (> to go forward or < to go backward)

Monitoring

The Red Hat Network Monitoring entitlement allows you to perform a whole host of actions designed to keep your systems running properly and efficiently. With it, you can keep close watch on system resources, network services, databases, and both standard and custom applications.

Monitoring provides both real-time and historical state-change information, as well as specific metric data. You are not only notified of failures immediately and warned of performance degradation before it becomes critical, but you are also given the information necessary to conduct capacity planning and event correlation. For instance, the results of a probe recording CPU usage across systems would prove invaluable in balancing loads on those systems.

Monitoring entails establishing notification methods, installing probes on systems, regularly reviewing the status of all probes, and generating reports displaying historical data for a system or service. This chapter seeks to identify common tasks associated with the Monitoring entitlement. Remember, virtually all changes affecting your Monitoring infrastructure must be finalized by updating your configuration, through the **Scout Config Push** page.

8.1. Prerequisites

Before attempting to implement RHN Monitoring within your infrastructure, ensure you have all of the necessary tools in place. At a minimum, you need:

- Monitoring entitlements — These entitlements are required for all systems that are to be monitored. Monitoring is supported only on Red Hat Enterprise Linux systems.
- RHN Satellite with Monitoring — Monitoring systems must be connected to a Satellite with a base operating system of Red Hat Enterprise Linux AS 4, Red Hat Enterprise Linux 5 or later. Refer to the RHN Satellite Installation Guide within **Help** for installation instructions.
- Monitoring Administrator — This role must be granted to users installing probes, creating notification methods, or altering the monitoring infrastructure in any way. (Remember, the Satellite Administrator automatically inherits the abilities of all other roles within an organization and can therefore conduct these tasks.). Assign this role through the **User Details** page for the user.
- Red Hat Network Monitoring Daemon — This daemon, along with the SSH key for the scout, is required on systems that are monitored in order for the internal process monitors to be executed. You may, however, be able to run these probes using the systems' existing SSH daemon (**sshd**). Refer to [Section 8.2, “Red Hat Network Monitoring Daemon \(rhnmd\)”](#) for installation instructions and a quick list of probes requiring this secure connection. Refer to [Appendix D, Probes](#) for the complete list of available probes.

8.2. Red Hat Network Monitoring Daemon (rhnmd)

To get the most out of your Monitoring entitlement, Red Hat suggests installing the Red Hat Network Monitoring Daemon on your client systems. Based upon **OpenSSH**, **rhnmd** enables the RHN Satellite to communicate securely with the client system to access internal processes and retrieve probe status.

Please note that the Red Hat Network Monitoring Daemon requires that monitored systems allow connections on port 4545. You may avoid opening this port and installing the daemon altogether by using **sshd** instead. Refer to [Section 8.2.3, “Configuring SSH”](#) for details.

8.2.1. Probes requiring the daemon

An encrypted connection, either through the Red Hat Network Monitoring Daemon or **sshd**, is required on client systems for the following probes to run:

- Linux::CPU Usage
- Linux::Disk IO Throughput
- Linux::Disk Usage
- Linux::Inodes
- Linux::Interface Traffic
- Linux::Load
- Linux::Memory Usage
- Linux::Process Counts by State
- Linux::Process Count Total
- Linux::Process Health
- Linux::Process Running
- Linux::Swap Usage
- Linux::TCP Connections by State
- Linux::Users
- Linux::Virtual Memory
- LogAgent::Log Pattern Match
- LogAgent::Log Size
- Network Services::Remote Ping
- Oracle::Client Connectivity
- General::Remote Program
- General::Remote Program with Data

Note that all probes in the Linux group have this requirement.

8.2.2. Installing the Red Hat Network Monitoring Daemon

Install the Red Hat Network Monitoring Daemon to prepare systems for monitoring with the probes identified in [Section 8.2.1, “Probes requiring the daemon”](#). Note that the steps in this section are optional if you intend to use **sshd** to allow secure connections between the RHN monitoring infrastructure and the monitored systems. Refer to [Section 8.2.3, “Configuring SSH”](#) for instructions.

The **rhnmmd** package can be found in the RHN Tools channel for all Red Hat Enterprise Linux distributions. To install it:

1. Subscribe the systems to be monitored to the RHN Tools channel associated with the system. This can be done individually through the **System Details** ▢ **Channels** ▢ **Software** subtab or for multiple systems at once through the **Channel Details** ▢ **Target Systems** tab.
2. Once subscribed, open the **Channel Details** ▢ **Packages** tab and find the **rhnmmd** package (under 'R').
3. Click the package name to open the **Package Details** page. Go to the **Target Systems** tab, select the desired systems, and click **Install Packages**.
4. Install the SSH public key on all client systems to be monitored, as described in [Section 8.2.4, "Installing the SSH key"](#).
5. Start the Red Hat Network Monitoring Daemon on all client systems using the command:

```
service rhnmmd start
```

6. When adding probes requiring the daemon, accept the default values for **RHNMD User** and **RHNMD Port**: **nocpu1se** and **4545**, respectively.

8.2.3. Configuring SSH

If you wish to avoid installing the Red Hat Network Monitoring Daemon and opening port 4545 on client systems, you may configure **sshd** to provide the encrypted connection required between the systems and RHN. This may be especially desirable if you already have **sshd** running. To configure the daemon for monitoring use:

1. Ensure the SSH package is installed on the systems to be monitored:

```
rpm -qi openssh-server
```

2. Identify the user to be associated with the daemon. This can be any user available on the system, as long as the required SSH key can be put in the user's `~/.ssh/authorized_keys` file.
3. Identify the port used by the daemon, as identified in its `/etc/ssh/sshd_config` configuration file. The default is port 22.
4. Install the SSH public key on all client systems to be monitored, as described in [Section 8.2.4, "Installing the SSH key"](#).
5. Start the **sshd** on all client systems using the command:

```
service sshd start
```

6. When adding probes requiring the daemon, insert the values derived from steps 2 and 3 in the **RHNMD User** and **RHNMD Port** fields.

8.2.4. Installing the SSH key

Whether you use **rhnm** or **sshd**, you must install the Red Hat Network Monitoring Daemon public SSH key on the systems to be monitored to complete the secure connection. To install it:

1. Navigate to the **Monitoring** ▢ **Scout Config Push** page on the Satellite interface and click the name of the RHN Server that will monitor the client system. The SSH **id_dsa.pub** key is visible on the resulting page.
2. Copy the character string (beginning with **ssh-dss** and ending with the hostname of the RHN Server).
3. On the command line of the system to be monitored, switch to the user aligned with the daemon. This is accomplished for **rhnm** with the command:

```
su - nocpulse
```

4. Paste the key character string into the `~/.ssh/authorized_keys` file for the daemon's user. For **rhnm**, this is `/var/lib/nocpulse/.ssh/authorized_keys`.

If config management is enabled on the systems to be monitored, you may deploy this file across systems using a config channel. Refer to [Section 7.7.1, "Preparing Systems for Config Management"](#) for details.



Note

If valid entries already exist in **authorized_keys**, add the daemon key to the file rather than replacing the existing key. To do so, save the copied text to **id_dsa.pub** in the same **.ssh/** directory and then run the following command: `cat ~/.ssh/id_dsa.pub >> ~/.ssh/authorized_keys`.

5. Finally, ensure the **.ssh/** directory and **authorized_keys** file have the appropriate permissions set. This can be done as the daemon's user with the following commands:

```
chmod 700 ~/.ssh
chmod 600 ~/.ssh/authorized_keys
```

Once the key is in place and accessible, all probes that require it should allow **ssh** connections between the Monitoring infrastructure and the monitored system. You may then schedule probes requiring the monitoring daemon to run against the newly configured systems.

8.3. mysql package

If your RHN Satellite will serve Monitoring-entitled client systems against which you wish to run **MySQL** probes, you must configure the **mysql** package on the RHN Satellite. Refer to [Appendix D, Probes](#) for a listing of all available probes.

Subscribe the Satellite to the Red Hat Enterprise Linux Base channel and install the **mysql** package either through the **up2date**, **yum** or RHN Hosted.

Once finished, your Satellite may be used to schedule MySQL probes.

8.4. Notifications

In addition to viewing probe status within the RHN interface, you may be notified whenever a probe changes state. This is especially important when monitoring mission-critical production systems. For this reason, Red Hat recommends taking advantage of this feature.

To enable probe notifications within RHN, you must have identified a mail exchange server and mail domain during installation of your RHN Satellite and configured **sendmail** to properly handle incoming mail. Refer to the *Installation* chapter of the *RHN Satellite Installation Guide* for details.

8.4.1. Creating Notification Methods

Notifications are sent via a *notification method*, an email or pager address associated with a specific RHN user. Although the address is tied to a particular user account, it may serve multiple administrators through an alias or mailing list. Each user account can contain multiple notification methods. To create a notification method:

1. Log into the RHN website as either an Satellite Administrator or Monitoring Administrator.
2. Navigate to the **User Details** ▾ **Notification Methods** tab and click **create new method**.
3. Enter an intuitive, descriptive label for the method name, such as **DBA day email**, and provide the correct email or pager address. Remember, the labels for all notification methods are available in a single list during probe creation, so they should be unique to your organization.
4. Select the checkbox if you desire abbreviated messages to be sent to the pager. This shorter format contains only the probe state, system hostname, probe name, time of message, and Send ID. The standard, longer format displays additional message headers, system and probe details, and instructions for response.
5. When finished, click **Create Method**. The new method shows up in the **User Details** ▾ **Notification Methods** tab and the **Notification** page under the top **Monitoring** category. Click its name to edit or delete it.
6. While adding probes, select the **Probe Notifications** checkbox and select the new notification method from the resulting dropdown menu. Notification methods assigned to probes cannot be deleted until they are dis-associated from the probe.

8.4.2. Receiving Notifications

If you create notification methods and associate them with probes, you must be prepared to receive them. These notifications come in the form of brief text messages sent to either email or pager addresses. Here is an example of an email notification:

```
Subject: CRITICAL: [hostname]: Satellite: Users at 1
From: "Monitoring Satellite Notification" (rogerthat01@redhat.com)
Date: Mon, 6 Dec 2004 13:42:28 -0800
To: user@organization.com
```

```
This is RHN Monitoring Satellite notification 01dc8hqw.
```

```
Time: Mon Dec 06, 21:42:25 PST
State: CRITICAL
System: [hostname] ([IP address])
```

```
Probe: Satellite: Users
Message: Users 6 (above critical threshold of 2)
Notification #116 for Users

Run from: RHN Monitoring Satellite
```

As you can see, the longer email notifications contain virtually everything you would need to know about the associated probe. In addition to the probe command, run time, system monitored, and state, the message contains the *Send ID*, which is a unique character string representing the precise message and probe. In the above message, the Send ID is **01dc8hqw**.

Pager notifications, by necessity, contain only the most important details, namely the subject of the email message (containing state, system, probe, and time) and the Send ID. Here is an example pager notification:

```
CRITICAL: [hostname]: Satellite: Users at 21:42 PST, notification 01dc8hqw
```

8.4.3. Redirecting Notifications

Upon receiving a notification, you may redirect it by including advanced notification rules within an acknowledgment email. Just reply to the notification and include the desired option. These are the possible redirect options, or *filter types*:

- **ACK METOO** — Sends the notification to the redirect destination(s) *in addition to* the default destination.
- **ACK SUSPEND** — Suspends the notification method for a specified time period.
- **ACK AUTOACK** — Does not change the destination of the notification, but automatically acknowledges matching alerts as soon as they are sent.
- **ACK REDIR** — Sends the notification to the redirect destination(s) *instead of* the default destination.

The format of the rule should be *filter_type probe_type duration email_address* where *filter_type* indicates one of the previous advanced commands, *probe_type* indicates probe or system, *duration* indicates the length of time for the redirect, and *email_address* indicates the intended recipient. For example:

```
ACK METOO system 1h boss@domain.com
```

Capitalization is not required. Duration can be listed in minutes (m), hours (h), or days (d). Email addresses are needed only for redirects (REDIR) and supplemental (METOO) notifications.

The description of the action contained in the resulting email defaults to the command entered by the user. The reason listed is a summary of the action, such as **email ack redirect by user@domain.com** where user equals the sender of the email.



Note

You can halt or redirect almost all probe notifications by replying to a notification emails with a variation of the command **ack suspend host**. However, you cannot halt Satellite

probe notifications by responding to a probe with **ack**, **suspend**, **host** or other redirect responses. These probes require you to change the notifications within the web interface of the Satellite.

8.4.4. Filtering Notifications

Since notifications can be generated whenever a probe changes state, simple changes in your network can result in a flood of notifications. The creation, cancellation, and application of Notification filters is discussed in detail in [Section 7.10.2.1, “Notification ▯ Filters”](#).

8.4.5. Deleting Notification Methods

Theoretically, removing notification methods should be as easy as creating them. After all, you must populate no fields to conduct the deletion and a button exists for this explicit purpose. However, existing relationships between methods and probes can complicate this process. Follow these steps to remove a notification method:

1. Log into the RHN website as an Satellite Administrator or Monitoring Administrator.
2. Navigate to the **Monitoring ▯ Notifications** page and click the name of the method to be removed.
3. On the **User Details ▯ Notification Methods** tab, click **delete method**. If the method is not associated with any probes, you are presented with a confirmation page. Click **Confirm Deletion**. The notification method is removed.



Tip

Since both the notification method name and address can be edited, consider updating the method rather than deleting it. This redirects notifications from all probes using the method without having to edit each probe and create a new notification method.

4. If the method is associated with one or more probes, you are presented with a list of the probes using the method and the systems to which the probes are attached instead of a confirmation page. Click the probe name to go directly to the **System Details ▯ Probes** tab.
5. On the **System Details ▯ Probes** tab, select another notification method and click **Update Probe**.
6. You may now return to the **Monitoring ▯ Notifications** page and delete the notification method.

8.5. Probes

Now that the Red Hat Network Monitoring Daemon has been installed and notification methods have been created, you may begin installing probes on your Monitoring-entitled systems. If a system is entitled to Monitoring, a **Probes** tab appears within its **System Details** page. This is where you will conduct most probe-related work.

8.5.1. Managing Probes

To add a probe to a system, the system must be entitled to Monitoring. Further, you must have access to the system itself, either as the system's root user, through the System Group Administrator role, or as the Satellite Administrator. Then:

1. Log into the RHN website as either an Satellite Administrator or the System Group Administrator for the system.
2. Navigate to the **System Details** ▢ **Probes** tab and click **create new probe**.
3. On the **System Probe Creation** page, complete all required fields. First, select the Probe Command Group. This alters the list of available probes and other fields and requirements. Refer to [Appendix D, Probes](#) for the complete list of probes by command group. Remember that some probes require the Red Hat Network Monitoring Daemon to be installed on the client system.
4. Select the desired Probe Command and the Monitoring Scout, typically **RHN Monitoring Satellite** but possibly an RHN Proxy Server. Enter a brief but unique description for the probe.
5. Select the **Probe Notifications** checkbox to receive notifications when the probe changes state. Use the **Probe Check Interval** dropdown menu to determine how often notifications should be sent. Selecting **1 minute** (and the **Probe Notification** checkbox) means you will receive notifications every minute the probe surpasses its CRITICAL or WARNING thresholds. Refer to [Section 8.4, "Notifications"](#) to find out how to create notification methods and acknowledge their messages.
6. Use the **RHNMD User** and **RHNMD Port** fields, if they appear, to force the probe to communicate via **sshd**, rather than the Red Hat Network Monitoring Daemon. Refer to [Section 8.2.3, "Configuring SSH"](#) for details. Otherwise, accept the default values of **nocpulse** and **4545**, respectively.
7. If the **Timeout** field appears, review the default value and adjust to meet your needs. Most but not all timeouts result in an UNKNOWN state. If the probe's metrics are time-based, ensure the timeout is not less than the time allotted to thresholds. Otherwise, the metrics serve no purpose, as the probe will time out before any thresholds are crossed.
8. Use the remaining fields to establish the probe's alert thresholds, if applicable. These CRITICAL and WARNING values determine at what point the probe has changed state. Refer to [Section 8.5.2, "Establishing Thresholds"](#) for best practices regarding these thresholds.
9. When finished, click **Create Probe**. Remember, you must commit your Monitoring configuration change on the **Scout Config Push** page for this to take effect.

To delete a probe, navigate to its **Current State** page (by clicking the name of the probe from the **System Details** ▢ **Probes** tab), and click **delete probe**. Finally, confirm the deletion.

8.5.2. Establishing Thresholds

Many of the probes offered by RHN contain alert thresholds that, when crossed, indicate a change in state for the probe. For instance, the Linux::CPU Usage probe allows you to set CRITICAL and WARNING thresholds for the percent of CPU used. If the monitored system reports 75 percent of its CPU used, and the WARNING threshold is set to 70 percent, the probe will go into a WARNING state. Some probes offer a multitude of such thresholds.

In order to get the most out of your Monitoring entitlement and avoid false notifications, Red Hat recommends running your probes without notifications for a time to establish baseline performance for each of your systems. Although the default values provided for probes may suit you, every organization has a different environment that may require altering thresholds.

8.5.3. Monitoring the RHN Server

In addition to monitoring all of your client systems, you may also use RHN to monitor your RHN Server, whether that be an RHN Satellite or RHN Proxy Server. To monitor your RHN Server, find a system monitored by the server, and go to that system's **System Details** ▢ **Probes** tab.

Click **create new probe** and select the **Satellite** Probe Command Group. Next, complete the remaining fields as you would for any other probe. Refer to [Section 8.5.1, “Managing Probes”](#) for instructions.

Although the RHN Server appears to be monitored by the client system, the probe is actually run from the server on itself. Thresholds and notifications work normally.



Note

Any probes that require Red Hat Network Monitoring Daemon connections cannot be used against a RHN Satellite or RHN Proxy Server on which Monitoring software is running. This includes most probes in the Linux command group as well as the Log Agent probes and the Remote Program probes. Use the Satellite command group probes to monitor RHN Satellites and RHN Proxy Servers. In the case of Proxy scouts, the probes are listed under the system for which they are reporting data.

8.6. Troubleshooting

Though all Monitoring-related activities are conducted through the RHN website, Red Hat provides access to some command line diagnostic tools that may help you determine the cause of errors. To use these tools, you must be able to become the **nocpulse** user on the RHN Server conducting the monitoring.

First log into the RHN Server as root. Then switch to the **nocpulse** user with the following command:

```
su - nocpulse
```

You may now use the diagnostic tools described within the rest of this section.

8.6.1. Examining Probes with `rhncatalog`

To thoroughly troubleshoot a probe, you must first obtain its probe ID. You may obtain this information by running **rhncatalog** on the RHN Server as the **nocpulse** user. The output will resemble:

```
2 ServiceProbe on example1.redhat.com (199.168.36.245): test 2
3 ServiceProbe on example2.redhat.com (199.168.36.173): rhe12.1 test
4 ServiceProbe on example3.redhat.com (199.168.36.174): SSH
5 ServiceProbe on example4.redhat.com (199.168.36.175): HTTP
```

The probe ID is the first number, while the probe name (as entered in the RHN website) is the final entry on the line. In the above example, the **5** probe ID corresponds to the probe named **HTTP**.

Further, you may pass the **--commandline (-c)** and **--dump (-d)** options along with a probe ID to **rhncatalog** to obtain additional details about the probe, like so:

```
rhncatalog --commandline --dump 5
```

The **--commandline** option yields the command parameters set for the probe, while **--dump** retrieves everything else, including alert thresholds and notification intervals and methods.

The command above will result in output similar to:

```
5 ServiceProbe on example4.redhat.com (199.168.36.175 ):
linux:cpu usage
    Run as: Unix::CPU.pm --critical=90 --sshhost=199.168.36.175
--warn=70 --timeout=15 --sshuser=nocpulse
--shell=SSHRemoteCommandShell --sshport=4545
```

Now that you have the ID, you use it with **rhncatalog** to examine the probe's output. Refer to [Section 8.6.2, "Viewing the output of rhncatalog"](#) for instructions.

8.6.2. Viewing the output of rhncatalog

Now that you have obtained the probe ID with **rhncatalog**, use it in conjunction with **rhncatalog** to examine the complete output of the probe. Note that by default, **rhncatalog** works in test mode, meaning no results are entered in the database. Here are its options:

Option	Description
--help	List the available options and exit.
--probe=PROBE_ID	Run the probe with this ID.
--prob_arg=PARAMETER	Override any probe parameters from the database.
--module=PERL_MODULE	Package name of alternate code to run.
--log=all=LEVEL	Set log level for a package or package prefix.
--debug=LEVEL	Set numeric debugging level.
--live	Execute the probe, enqueue data and send out notifications (if needed).

Table 8.1. **rhncatalog** Options

At a minimum, you should include the **--probe** option, the **--log** option, and values for each. The **--probe** option takes the probeID as its value and the **--log** option takes the value "all" (for all run levels) and a numeric verbosity level as its values. Here is an example:

```
rhncatalog --probe=5 --log=all=4
```

The above command requests the probe output for probeID 5, for all run levels, with a high level of verbosity.

More specifically, you may provide the command parameters derived from **rhn-catalog**, like so:

```
rhn-runprobe 5 --log=all=4 --sshuser=nocpulse --sshport=4545
```

This yields verbose output depicting the probe's attempted execution. Errors are clearly identified.

Multiple Organizations

RHN Satellite supports the creation and management of *multiple organizations* within one Satellite installation, allowing for the division of systems, content, and subscriptions across different organizations or specific groups. This chapter guides the user through basic setup tasks and explains the concepts of multiple organization creation and management within RHN Satellite.

9.1. Recommended Models for Using Multiple Organizations

The following examples detail two possible scenarios using the multiple organizations (or multi-org) feature. Installing or upgrading to RHN Satellite 5.1 or later does not require that you make use of the multi-org feature. You may create additional organizations on your Satellite and start using those organizations at whatever pace makes the most sense for you. It is a good idea to create an additional organization and use it on a trial basis for a limited set of systems/users to fully understand the impact of a multi-org Satellite on your organization's processes and policies.

9.1.1. Centrally-Managed Satellite for A Multi-Department Organization

In this first scenario, the RHN Satellite is maintained by a central group within a business or other organization (refer to [Figure 9.1, "Centralized Satellite Management for Multi-Department Organization"](#)). The Satellite administrator of Organization 1 (the administrative organization created during Satellite configuration) treats Organization 1 (the 'Administrative Organization') as a staging area for software and system subscriptions and entitlements.

The Satellite administrator's responsibilities include the configuration of the Satellite (any tasks available under the **Admin** area of the web interface), the creation and deletion of additional Satellite organizations, and the allocation and removal of software and system subscriptions and entitlements.

Additional organizations in this example are mapped to departments within a company. One way to decide what level to divide the various departments in an organization is to think about the lines along which departments purchase subscriptions and entitlements for use with RHN Satellite. To maintain centralized control over organizations in the Satellite, create an Organization Administrator account in each subsequently created organization so that you may access that organization for any reason.

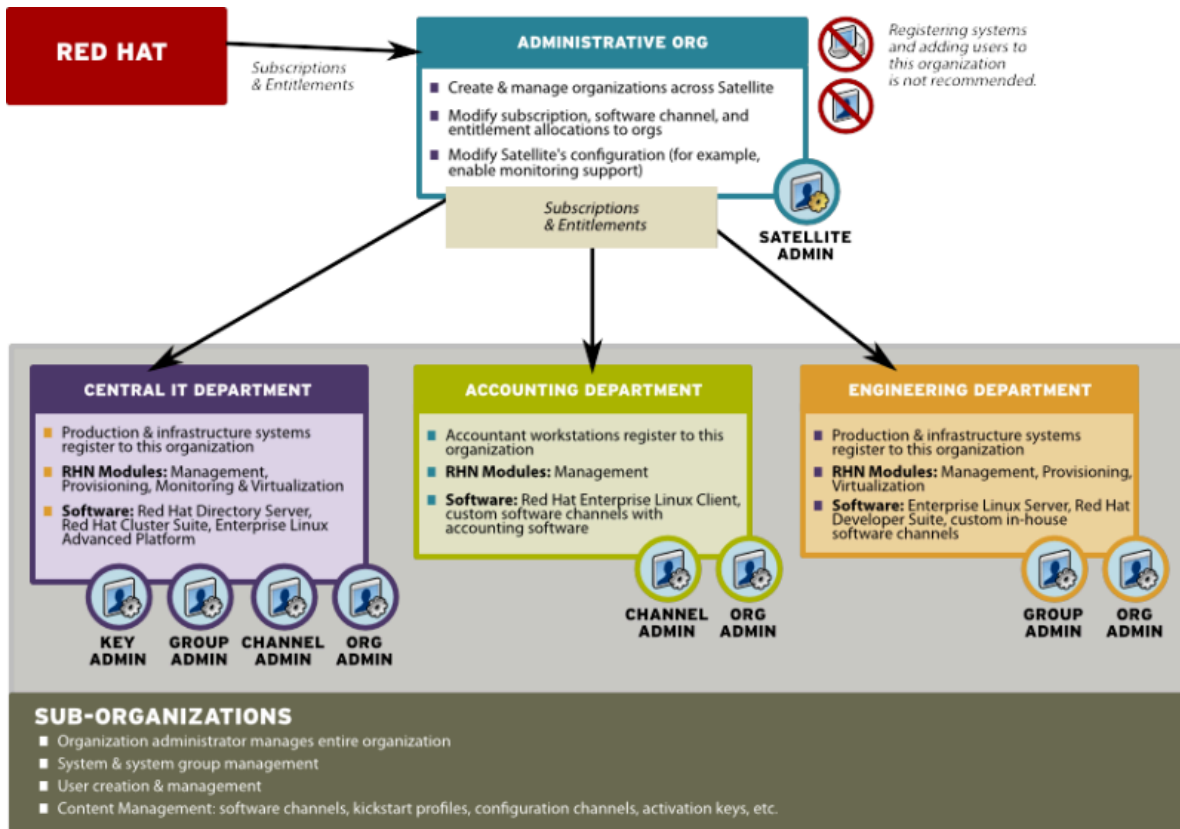


Figure 9.1. Centralized Satellite Management for Multi-Department Organization

9.1.2. Decentralized Management of Multiple Third Party Organizations

In this example, the Satellite is maintained by a central group, but each organization is treated separately without relations or ties to the other organizations on the Satellite. Each organization may be a customer of the group that manages the Satellite application itself.

While a Satellite consisting of sub-organizations that are all part of the same company may be an environment more tolerant of sharing systems and content between organizations, in this decentralized example sharing is less tolerable. Administrators can allocate entitlements in specific amounts to each organization. Each organization will have access to all Red Hat content synced to the Satellite if the organization has software channel entitlements for the content.

However, if one organization pushes custom content to their organization, it will not be available to other organizations. You cannot provide custom content that is available to all or select organizations without re-pushing that content into each organization.

In this scenario, Satellite Administrators may want to reserve an account in each organization to have login access. For example, if you are using Satellite to provide managed hosting services to external parties, you could reserve an account for yourself so to access systems in that organization and push content.

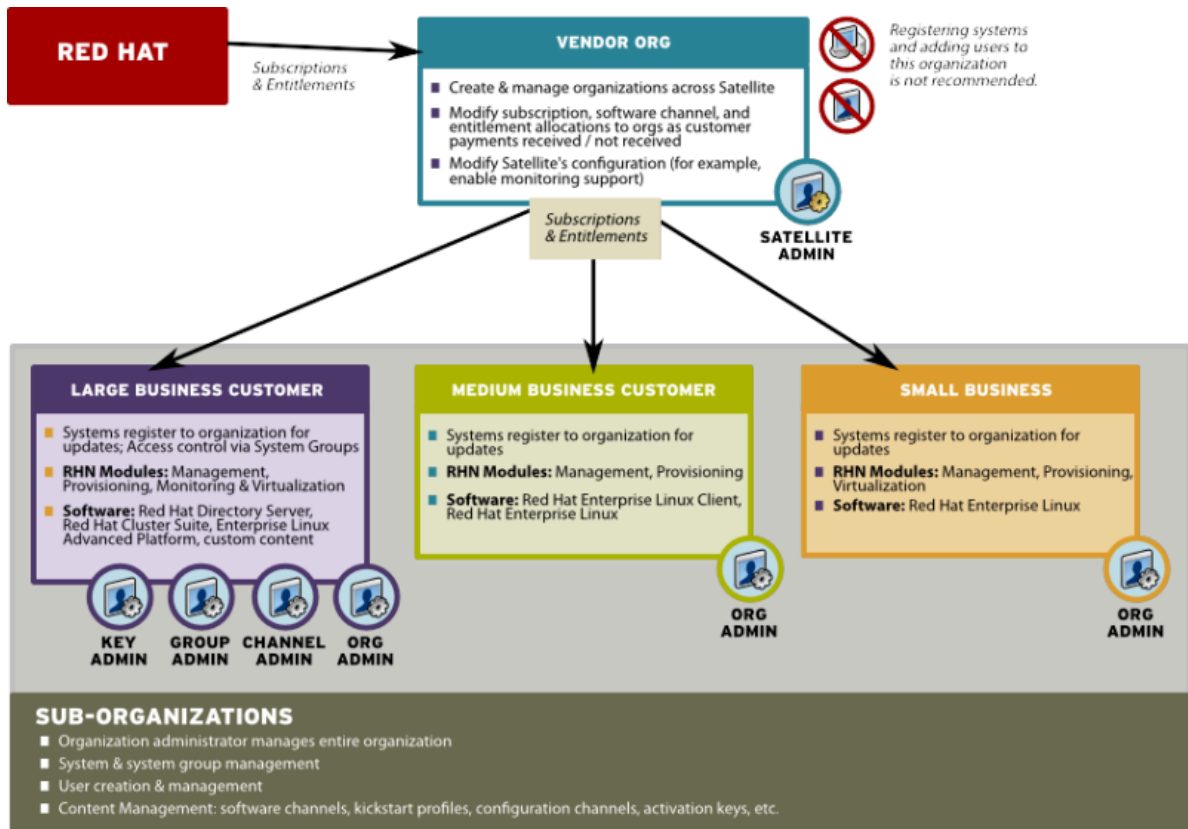


Figure 9.2. Decentralized Satellite Management for Multi-Department Organization

9.1.3. General Tips for Multi-Org Usage

Regardless of the specific model above you choose in the management of your multi-org Satellite, the following best practices tips can help.

It is not recommended to use the administrative organization (organization #1) for registering systems and creating users in any situation unless you intend to use Satellite as a single organization Satellite or are in the process of migrating from a single organization Satellite to a multiple organization Satellite. This is due to the following reasons:

1. The administrative organization is treated as a special case with respect to entitlements. You can only add or remove entitlements to this organization implicitly by removing them or adding them from the other organizations on the Satellite.
2. The administrative organization is intended to be a staging area for subscriptions and entitlements. When you associate the Satellite with a new certificate, any new entitlements will be granted to this organization by default. In order to make those new entitlements available to other organizations on the Satellite, you will need to explicitly allocate those entitlements to the other organizations from the administrative organization.

9.1.3.1. Certificate Has Less Entitlements Than I Am Using

If you are issued a new Satellite certificate and it contains less entitlements than the systems in the organizations on your Satellite are consuming, you will be unable to activate this new certificate when uploading it through the Satellite's web interface under **Admin** → **Satellite Configuration** → **Certificate**, uploading it through the <http://rhn.redhat.com> profile of the Satellite system under the **Satellite** tab,

or by running the **rhn-satellite-activate** command. You will get an error stating that there are insufficient entitlements in the certificate.

There are a few ways you can reduce Satellite entitlement usage in order to activate your new certificate. Red Hat recommends evaluating each organization's entitlement usage on the Satellite and decide which organizations should relinquish some entitlements and still function properly. You can then contact each organization administrator directly and request that they unentitle or delete the system profiles of any extraneous systems in their organizations. If you have login access to these organizations, you can do this yourself. Logged in under a Satellite administrator, you cannot decrement the allocated entitlements to an organization below the number of entitlements that organization has actively associated with system profiles.

There are some situations in which you need to free entitlements and do not have a lot of time to do so, and may not have access to each organization in order to do this yourself. There is an option in Multi-Org Satellites that allows the Satellite administrator to decrement an organization's entitlement count below their usage. This method must be done logged into the administrative organization.

For example, logged into the administrative organization, if your certificate is 5 system management entitlements shy of being able to cover all registered systems on your Satellite, the 5 systems that were most recently registered to that organization will be unentitled. This process is described below:

1. In the `/etc/rhn/rhn.conf` file, set `web.force_unentitlement=1`
2. Restart the Satellite
3. Reduce the allocated entitlements to the desired organizations either via each organization's **Subscriptions** tab or via individual entitlement's **Organizations** tabs.
4. A number of systems in the organization should now be in an **unentitled** state. The number of systems unentitled in the organization will be equal to the difference between the total number of entitlements you removed from the organization and the number of entitlements the organization did not have applied to the systems.

For example, if you removed 10 entitlements from the organization in step 3, and the organization has 4 entitlements that were not in use by systems, then 6 systems in the organization will be unentitled.

After you have the sufficient number of entitlements required, you should then be able to activate your new Satellite certificate. Note that modifying the `web.force_unentitlement` variable is only necessary to decrement an organization's allocated entitlements below what they are using. If an organization has more entitlements than are being actively used, you do not need to set this variable to remove them.

9.1.3.2. Certificate Has More Entitlements Than I Am Using

If you are issued a new Satellite certificate and it has more entitlements than are being consumed on your Satellite, any extra entitlements will be assigned to the administrative organization. If you log into the web interface as the Satellite administrator, you will then be able to allocate these entitlements to other organizations. The previously-allocated entitlements to other organizations will be unaffected.

9.2. Admin ¶ Organizations

The **Organizations** Web interface allows administrators to view, create, and manage multiple organizations across the Satellite. Administrators can allocate software and system entitlements

across various organizations, as well as control an organization's access to systems management tasks.



Figure 9.3. Admin

The **Organizations** page contains a listing of organizations across the Satellite, with both User and System counts assigned to each organization. The **Organizations** page also features a **Trusts** page for any organizational trusts established. Refer to [Section 9.6, “Organizational Trusts”](#) for more information about establishing organizational trusts.

9.2.1. Admin ▢ Organizations ▢ Details

Clicking on an organization displays the **Details** page, where administrators are provided a summary of various aspects of the organization.

- **Active Users** — The number of users in the organization
- **Systems** — The number of systems subscribed to the organization.
- **System Groups** — The number of groups subscribed to the organization.
- **Activation Keys** — The number of activation keys available to the organization.
- **Kickstart Profiles** — The number of kickstart profiles available to the organization.
- **Configuration Channels** — The number of Configuration Channels available to the organization.

From this page, you can delete the organization by clicking the **Delete Organization** link.

The **Details** page also contains three subtabs: **Users**, **Subscriptions**, and **Trusts**.

9.3. Creating an Organization

The **Create New Organization** page in the RHN Satellite web interface can be accessed by proceeding to **Admin** ▾ **Organizations** ▾ **Create New Organization**.

Administrators can create new organizations and assign entitlements, groups, systems, and users to the group so that organizations can perform administrative tasks on their own without affecting other organizations.

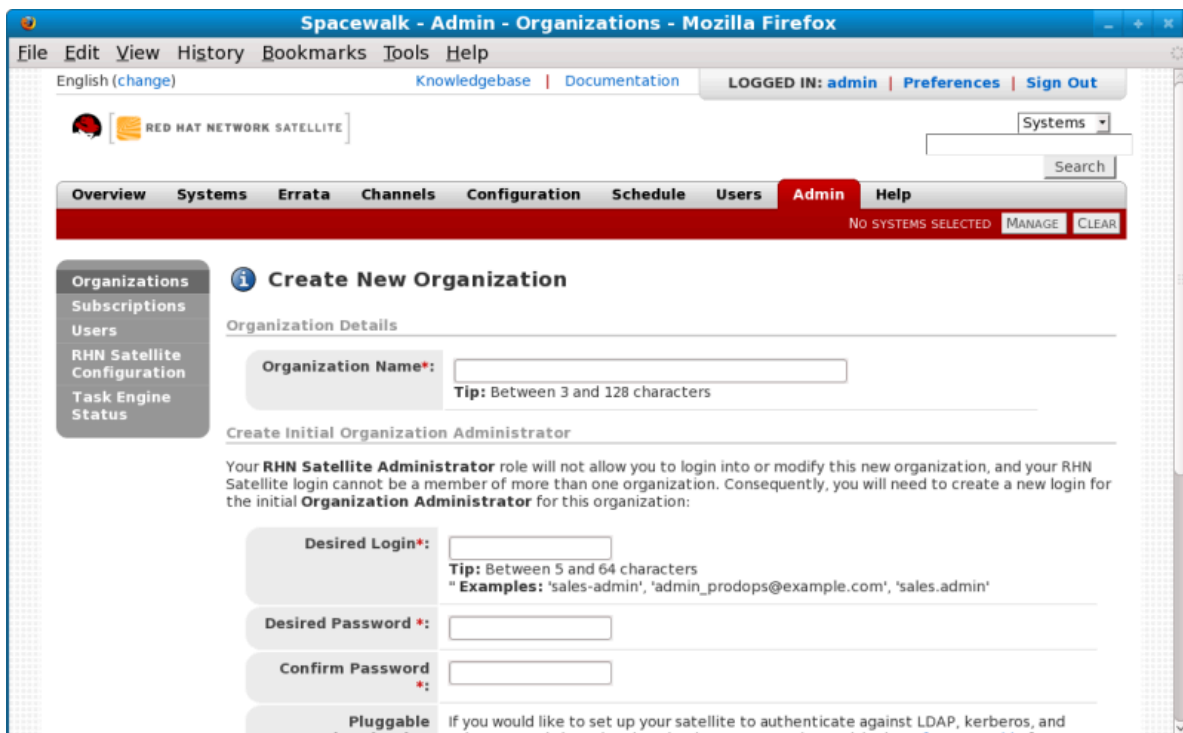


Figure 9.4. Create New Organization

1. Input the **Organization Name** in the provided text box. The name should be between 3 and 128 characters.
2. Create an administrator for the organization:
 - a. Enter a **Desired Login** for the organization administrator, which should be between 3 and 128 characters long.
 - b. Create a **Desired Password** and **Confirm** the password.
 - c. Type in the **Email** for the organization administrator.
 - d. Enter the **First Name** and **Last Name** of the organization administrator.
3. Click the **Create Organization** button to complete the process.

Once the new organization is created, the **Organizations** page will display with the new organization listed.



Tip

Satellite Administrators should consider reserving the administrative Organization Administrator account for themselves to have the option of logging into this organization for various reasons. If your Satellite is configured for PAM authentication, avoid using PAM accounts for the administrative organization administrator account in new organizations. Instead, create a Satellite-local account for organization administrators and reserve PAM-authenticated accounts for Satellite logins with less elevated privileges in order to discourage users to frequently log into the Satellite with elevated privileges, as the potential for making mistakes is higher using these accounts.

Additionally, consider creating a login name for the administrative Organization Administrator account that describes (for example, **orgadmin-mktg** or **eng-dept-admin**), to match admin login names with the organization.

9.4. Managing Entitlements

One important task after creating a new organization is to assign management entitlements to the new organization. Management system entitlements are a base requirement for an organization to function on the Satellite. The number of management entitlements allocated to an organization is equivalent to the maximum number of systems that may register to that organization on the Satellite, regardless of the number of software entitlements available. For example, if there are 100 Red Hat Enterprise Linux Client entitlements but only 50 management system entitlements to an organization, only 50 systems are able to register to that organization.

You must also grant RHN Tools software channel entitlements to each organization. The RHN Tools channel contains various client software required for extended Satellite functionality, such as clients necessary for configuration management and kickstart support as well as the **rhn-virtualization** package, which is necessary for the entitlements of Xen virtual guests to be counted correctly corresponding to the number of Red Hat Enterprise Linux subscriptions to which they are associated.

Access the **Subscriptions** tab by clicking **Admin** ▾ **Organizations** ▾ **Details** ▾ **Subscriptions**.

The **Subscriptions** tab has two subtabs for managing the software channel and system entitlements for the organization.

9.4.1. Admin ▾ Subscriptions ▾ Software Channel Entitlements

The **Software Channel Entitlements Across Satellite** page lists of all entitlements on the Satellite, throughout all organizations, as well their usage. Click on a **Entitlement Name** for a more detailed view.

The **Details** subtab for the software channel entitlement contains information about the software channel access granted when subscribed to the entitlement.

The **Organizations** subtab allows Satellite administrators to adjust the number of software channels available to each organization. Type in the number (within the range listed in **Possible Values**) and click the **Update** button for that organization.



Note

Organization Administrators that create a custom channel can only use that channel within their organization unless an Organizational Trust is established between the organizations that want to share the channel. For more information about organizational trusts, refer to [Section 9.6, "Organizational Trusts"](#).

The **Organizations** subtab also contains broad usage information in the **System-Wide Entitlement Usage** section, including:

- **Total** — The total number of channel entitlements for the Satellite.
- **Available** — The number of entitlements currently available for allocation.
- **Usage** — The number of entitlements currently in use by all organizations (aside from the base organization), compared to the total number of entitlements allocated.

For example, if the **Total** column is 100 and the **Available** column is 70, that means 30 entitlements are allocated for organizations. The **Usage** column shows how many of those 30 allocated entitlements are in use by organizations besides the base organization. So if the **Usage** column reads *24 of 30 (80%)*, that means 24 channel entitlements are distributed to Satellite organizations (other than the base organization) out of 30 total allocated.

9.4.2. Admin ▯ Subscriptions ▯ System Entitlements

The **System Entitlements Across Satellite** page lists all system entitlements on this Satellite, across all organizations, as well as their usage. Click on the entitlement's name for more details about it.

System entitlements include **Management**, **Provisioning**, **Monitoring**, and **Virtualization**. Enter the number of allocations of each system entitlement in the text box, not to exceed the limit indicated in the **Possible Values**.

The **Details** subtab for the system entitlement contains information about the entitlement and what access it grants.

The **Organizations** subtab allows Satellite administrators to adjust the number of system entitlement allocations available to each organization. Type in the number (within the range listed in **Possible Values**) and click the **Confirm Changes** button for that organization.

The **Organizations** subtab for the system entitlement also contains broad usage information in the "Satellite-Wide Entitlement Usage" section, including:

- **Total Allocated** — The number of total entitlements available for the entire Satellite.
- **Entitlement Usage** — The number of entitlements currently being used.
- **Organization Usage** shows the number of organizations that have access to the entitlement.

9.5. Configuring Systems in an Organization

Now that an organization has been created and requisite entitlements assigned to it, you can then assign systems to each organization.

There are two basic ways to register a system against a particular organization:

1. Registering Using Login and Password — If you provide a login and password created for a specified organization, the system will be registered to that organization. For example, if **user -123** is a member of the **Central IT** organization on the Satellite, the following command on any system would register that system to the **Central IT** organization on your Satellite:

```
rhncattool --username=user-123 --password=foobaz
```



Note

The `--orgid` (for Red Hat Enterprise Linux 4 and 5) and `--orgpassword` (in RHEL 4) parameters in `rhncattool` are *not related* to Satellite registration or RHN Satellite's multiple organizations support.

2. Registering Using An Activation Key — You can also register a system to an organization using an activation key from the organization. Activation keys will register systems to the organization in which the activation key was created. Activation keys are a good registration method to use if you want to allow users to register systems into an organization without providing them login access to that organization. If you want to move systems between organizations, you may also automate the move with scripts using the activation keys.



Note

Activation keys have a new format since RHN Satellite 5.1.0, so the first few characters of the activation key are used to indicate which organization (by ID number) owns the activation.

9.6. Organizational Trusts

Organizations can share their resources with each other by establishing an *organizational trust* in the Satellite. An organizational trust is bi-directional, meaning that once a Satellite Administrator establishes a trust between two or more organizations, the Organization Administrator from each organization is free to share as much or as little of their resources as they need to. It is up to each Organization Administrator to determine what resources to share, and what shared resources from other organizations in the trust to use.



Note

Only Organization Administrators are able to share their custom content; Satellite Administrators only allocate system and software entitlements to each organization.

9.6.1. Establishing an Organizational Trust

A Satellite Administrator can create a trust between two or more organizations. To do this, click the **Organizations** link on the side menu on the **Admin** main page.

Click the name of one of the organizations and within the **Details** page, click the **Trusts** subtab.

On the **Trusts** subtab, there is a listing of all the other trusts on the RHN Satellite. Here you may use the **Filter by Organization** text box to narrow down a long list of organizations to a specific subset.

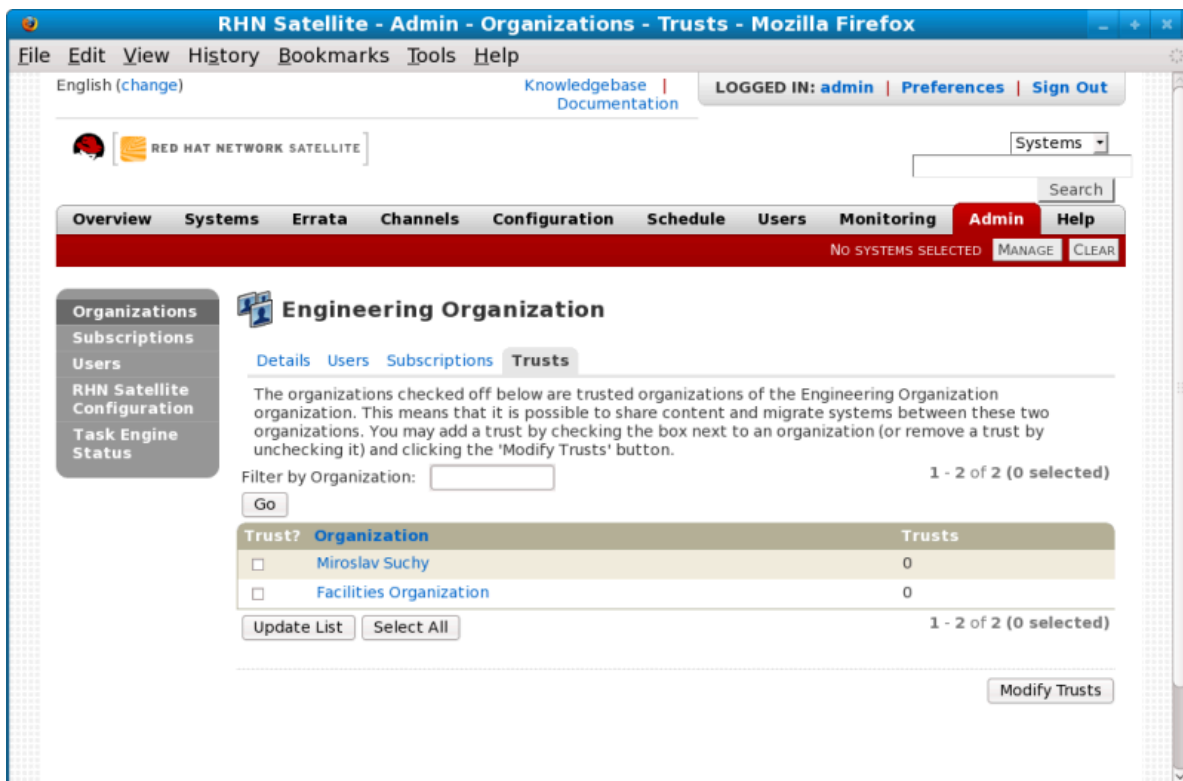


Figure 9.5. Organizational Trusts

Click the checkbox next to the names of the organizations you want to be in the organizational trust with the current organization and click the **Modify Trusts** button.

9.6.2. Sharing Content Channels between Organizations in a Trust

Once an organizational trust has been established, organizations can now share content such as custom software channels with the other organizations in the trust. There are also three levels of channel sharing that can be applied to each channel for finer-grained channel access control.



Note

Organizations cannot share Red Hat Channels because they are available to all organizations that have entitlements to those channels.

To share a custom channel with another organization, perform the following steps:

1. Login to the Satellite with the username of the Organization Administrator.
2. Click on the **Channels** tab.
3. On the side menu, click **Manage Software Channels**.
4. Click the custom channel that you want to share with the other organizations.

5. From the **Channel Access Control** section of the **Details** page, there are three choices for sharing in **Organizational Sharing**.
 - **Private** — Make the channel private so that it cannot be accessed by any organizations except the channel's owner.
 - **Protected** — Allow the channel to be accessed by specific trusted organizations of your choice.



Note

Choosing **Protected** sharing displays a separate page that prompts you to confirm that you are granting channel access to the organizations by clicking **Grant Access and Confirm**.

- **Public** — Allow all organizations within the trust to access the custom channel.

Click the radio button next to your selection and click **Update Channel**.

Now, any other Organization Administrators within the trust for which you have granted access to your custom channel can allow their client systems to install and update packages from the shared channel.



Note

If you have a system subscribed to a shared channel, and the organizational administrator of the shared channel changes access rights to the channel, then the system loses that channel. If he changes a base channel right, then the system will have no base channel on the **Systems** page and will not receive updates.

9.6.3. Migrating Systems from One Trusted Organization to Another

In addition to sharing software channels, organizations in a trust can migrate systems to other trusted organizations by using a utility called **migrate-system-profile**.

migrate-system-profile usage is based on the command-line, and uses systemIDs and orgIDs as arguments to specify what is being moved and its destination organization.

To use the **migrate-system-profile** command, you must have the **spacewalk-backend-tools** package installed. You do not need to be logged into the Satellite server to use **migrate-system-profile**; however, if you do not you will need specify the hostname or IP address of the server as a command-line switch.



Note

When an organization migrates a system with the **migrate-system-profile** command, the system does not carry any of the previous entitlements or channel subscriptions from the source organization. However, the system's history is preserved, and can be accessed by the new Organization Administrator in order to simplify the rest of the migration process, which includes subscribing to a base channels and granting entitlements.

9.6.3.1. Using `migrate-system-profile`

Using `migrate-system-profile` is straightforward. You need to ascertain the ID of the system to be migrated, the ID of the organization the system will migrate to, and the hostname or IP address of the Satellite server if you are running the command from another machine.

The usage from the command line is the following:

```
migrate-system-profile --satellite {SATELLITE HOSTNAME OR IP} --systemId={SYSTEM ID} --to-org-id={DESTINATION ORGANIZATION ID}
```

For example, if the Finance department (created as an organization in RHN Satellite with OrgID 2) wants to migrate a workstation (with SystemID 10001020) from the Engineering department, but the Finance Organization Administrator does not have shell access to the RHN Satellite server. The RHN Satellite hostname is **satserver.example.com**.

The Finance Organization Administrator would type the following from a shell prompt:

```
migrate-system-profile --satellite satserver.example.com --systemId=10001020 --to-org-id=2
```

The Finance Organization Administrator is then prompted for their username and password (unless they specified it using `--username=` and `--password=` at the command-line).

The Finance Organization Administrator would then be able to see the system from the **Systems** page when logged into the RHN Satellite web interface. The Finance Organization Administrator can then finish the migration process by assigning a base channel and granting entitlements to the client as he would any other system registered to his organization, which is available from the system's **History** page in the **Events** subtab.

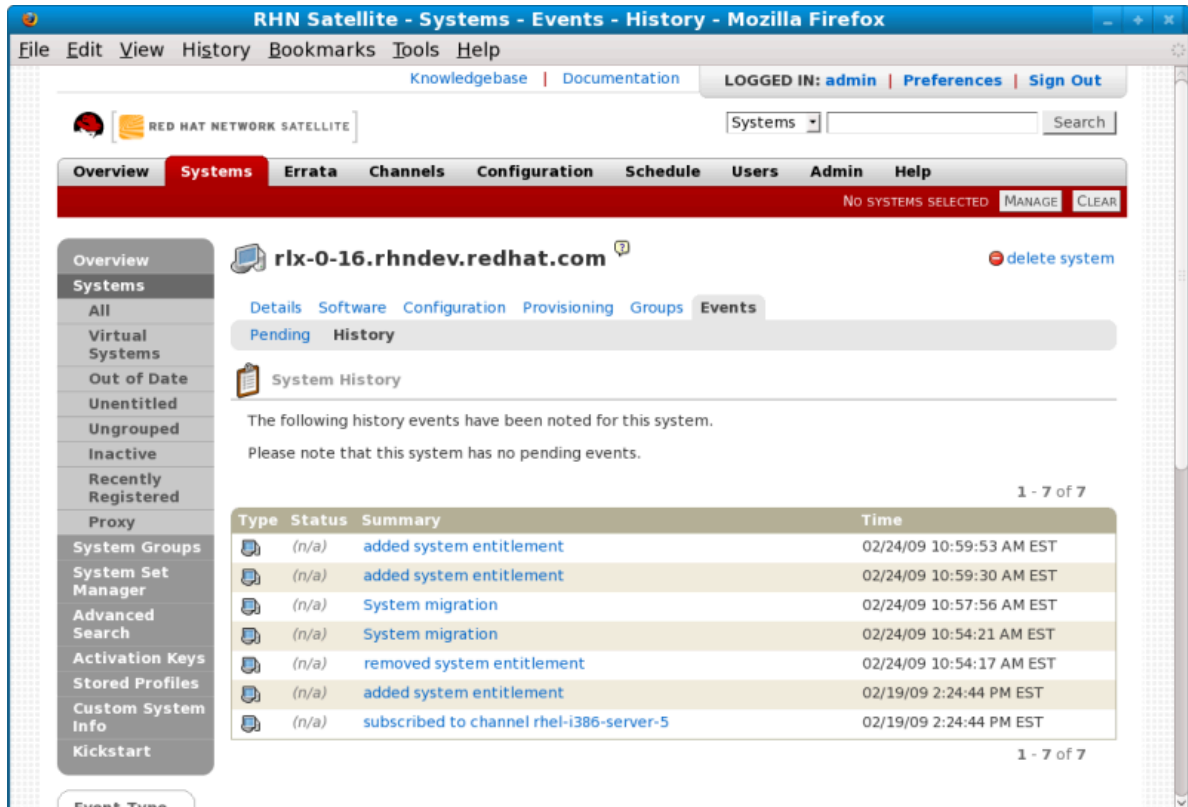


Figure 9.6. System History

Note

The Satellite Administrator can migrate a system from one trusted organization to any other in the trust. However, Organization Administrators can only migrate a system from their own organization to another in the trust.

Satellite Administrators that need to migrate several systems at once can use the `--csv` option of `migrate-system-profile` to automate the process using a simple comma-separated list of systems to migrate.

A line in the CSV file should contain the ID of the system to be migrated as well as destination organization's ID in the following format:

```
systemId,to-org-id
```

the `systemId`, for example could be `1000010000`, while the `to-org-id` could be `4`. So, a compatible CSV could look like the following:

```
1000010000,3
1000010020,1
1000010010,4
```

For more information about using `migrate-system-profile` refer to the manual page by typing `man migrate-system-profile` or for a basic help screen type `migrate-system-profile -h`.

9.7. Admin ▾ Users

The **Users Across Satellite** page contains a list of all users on the Satellite, throughout all organizations.



Note

You are only able to modify the details of organization users if you are logged in as that Organization Administrator.

Clicking the **Username** displays the **User Details** page. Refer to [Section 7.9, “Users — !\[\]\(c694a3ff3b077d76910920a6a1593ab4_img.jpg\)](#)” for more information on user configuration.

9.7.1. Admin ▾ Organizations ▾ Details ▾ Users

The **Users** subtab lists the users assigned to the organization, including their real names, email address, and a check mark indicating that the user is an administrator of the organization.

If you are the Organization Administrator, you can click the username to display the **User Details** page for the user. For instructions regarding user management, refer to [Section 7.9.1.1, “User List ▾ Active ▾](#)

[User Details — !\[\]\(dd161862f9164df98f62b726e9846241_img.jpg\)”](#).



Note

You must be logged in as the Organization Administrator to edit the User details for an organization. The Satellite Administrator cannot edit user details for organization users.

Virtualization

In order to manage and provision your client systems, you must first synchronize content from RHN's central servers to your Satellite.

RHN recommends that you sync at least the following channels:

- Red Hat Network Tools for RHEL Server (v. 5 for 32-bit x86) — `rhn-tools-rhel-i386-server-5`
- RHN Tools — `rhn-tools-rhel-4-as-i386`
- Red Hat Enterprise Linux Server (v. 5 for 32-bit x86) — `rhel-i386-server-5` (and all child channels)
- Red Hat Enterprise Linux Server Virtualization (v. 5 for 32-bit x86) — `rhel-i386-server-vt-5` (and all child channels)

10.1. Setting Up the Host System for Your Virtual Systems

Before creating guest systems, you must first prepare your host system. To do this, create a Red Hat Enterprise Linux 5 Server kickstart profile, then use that kickstart profile to install the operating system on your host. Once these steps are complete, you can proceed to provision virtual guests.

10.1.1. Create a Kickstart Profile for the Guest Systems

1. Login to the Satellite's web interface. Navigate to the **Kickstart Overview** screen by clicking the **Manage Kickstarts** link in the **Tasks** widget in **Your RHN**, or by clicking on the **Systems** tab, followed by the **Kickstart** subtab in the left navigation bar.
2. On the **Kickstart Overview** page, click the **Create a New Kickstart Profile** link in the **Kickstart Actions** widget in the upper right corner.
3. You should now find yourself on Step 1 of the kickstart profile creation process:
 - a. Enter a label for your profile that will enable you to distinguish it from your other profiles. For the remaining instructions, we'll assume the label is **host-system-for-virtual-guests**.
 - b. For the **Base Channel** field, select Red Hat Enterprise Linux (v.5 for \$ARCH) (where \$ARCH is the architecture of your host system).



Note

You may install 32-bit Red Hat Enterprise Linux 5 on a 64-bit host system. If you choose to do this, however, please be aware that your guest systems must also run the 32-bit version of Red Hat Enterprise Linux.

- c. In the **Kickstartable Tree** field, select **ks-rhel-\$ARCH-server-5** where \$ARCH is the architecture of your host system.
- d. Please select **Para-Virtualized Host** for the **Virtualization Type** field.



Note

If you are changing the **Virtualization Type** of an existing kickstart profile, it may also modify the bootloader and partition options, potentially overwriting any user customizations. Be sure to review the **Partitioning** tab to verify these settings when changing the **Virtualization Type**.

- e. Finally, click the **Next** button in the lower right of the screen to continue on to the next step.



Note

If any of the fields are missing the options indicated above, you may not have successfully synced software channel content to your Satellite from Red Hat's servers.

4. For Step 2 of the kickstart profile creation process, select the location of the distribution files for the installation of your host system. There should already be a **Default Download Location** filled out and selected for you on this screen. Click the **Next** button on this screen to continue to Step 3.



Note

As in the previous step, if the default download location is missing, you may not have successfully synced software channel content to your Satellite from Red Hat's server.

5. For Step 3 of the kickstart profile creation process, please choose a root password to set on the host system you will be provisioning, and click **Finish** to finish creation of the profile.
6. This completes kickstart profile creation. After completing Step 3, you are taken to the newly-created kickstart profile. You may browse through the various tabs of the profile and modify the settings as you see fit, but this is not necessary as the default settings should work well for the majority of cases.

10.1.2. Kickstart Your Host System

Next, kickstart your host system using your newly-created kickstart profile. There are three different scenarios for kickstarting your host system. Please read through these three scenarios below, and follow the instructions for the scenario that applies best to you:

10.1.2.1. Your Host System Has Red Hat Enterprise Linux 4 or Earlier Installed

In this case, register your host system to your Satellite and schedule the kickstart process via the Satellite's web interface.

1. First, register your host system to your Satellite. Use **ssh** to connect to your host system. Register your host system to your satellite issuing the following command as root:

```
rhncat /usr/bin/rhnreg_ks
```

```
--serverUrl=http://your-satellite.example.com/XMLRPC \  
--username=username --password=password
```



Note

If your host system is already registered to a different Red Hat Network server, add the **--force** option to the command above.

2. Next, open up the host system's profile in the Satellite web interface. Log into the web interface of your Satellite at <https://your-satellite.example.com/>. Click on the **Systems** tab in the top navigation bar. You should see the host system you just registered — click on its profile name to access its system profile page.
3. Add a provisioning entitlement to your host system. From your host system profile page, click on **Details** ▾ **Properties** tab. Check the **Provisioning** checkbox in the **Add-On Entitlements** field, and click the **Update Properties** button in the lower right hand corner of the screen.
4. Next, schedule the kickstart. You are brought back to the host system's profile page. You should now see a **Provisioning** tab in the system profile. Click on this tab. This should bring up the **Schedule Kickstart** page for the system.
5. Select the kickstart profile we created for this host earlier. Then, select the **Schedule Kickstart and Finish** button in the lower right-hand corner of the screen.



Note

If you do not see the kickstart profile you created earlier on the host system's **Schedule Kickstart** page, you may have created a kickstart profile for an architecture that does not match the architecture of the host system you have registered. If this is the case, open the kickstart profile by navigating to **Systems** ▾ **Kickstart** ▾ **Profiles** within the Satellite web interface, and clicking on the label for the host system's kickstart profile. Click on the **Kickstart Detail** ▾ **Operating System** tab, and select items under the **Base Channel** and **Available Trees** selections that match the architecture of your host system. Click on the **Update Kickstart** button in the lower right hand corner of the screen, and navigate back to the host system's **Schedule Kickstart** page, following the steps above this note.

6. After scheduling the kickstart, you will be taken to a **Kickstart Status** screen in the Satellite's web interface. Keep your web browser open to that page to follow along with the host system's progress.
7. Use **ssh** to connect to the host system, and run the command **rhncpck**. This should cause the kickstart process to run immediately rather than the next time the **rhncpck** process runs on the system. You should immediately see output indicating the start of a kickstart process on the host system, and it will eventually warn you that the system is going down for reboot in three minutes.
8. After three minutes have passed, the system will reboot. Follow the progress of the kickstart via the Satellite web interface.

- Depending on various factors, the kickstart process may take between ten and thirty minutes. At the end of this time period, the Satellite kickstart status page should indicate if the kickstart finished successfully.



Tip

If the kickstart fails, the Satellite kickstart status page should indicate that there was a failure. For more details on why the kickstart failed, click on the **Events** ▢ **History** tab in the host system's profile, and click on the name of the kickstart event that failed to get more details on the failure. It may also be useful to consult `/var/log/up2date` on the host system for troubleshooting purposes.

10.1.2.2. Your Host System Does Not have Red Hat Enterprise Linux Installed

First, create a boot CD to initiate the kickstart on your host system. You will be able to use the kickstart profile we created in earlier steps to provision the host. Note you must have physical access to the machine you intend to use in order to follow these steps:

- You will find an ISO to create a boot CD for you host by using **ssh** to log into your Satellite. It is at the following location on your satellite:

```
/var/satellite/rhn/kickstart/ks-rhel-i386-server-5/images/boot.iso
```

For details on how to use this ISO image to burn a CD using Linux, please refer to the following Red Hat Knowledgebase Article:

http://kbase.redhat.com/faq/FAQ_80_446.shtml

If you must burn this ISO image to CD using another operating system, please refer to the following Knowledgebase Article:

http://kbase.redhat.com/faq/FAQ_35_1897.shtml



Tip

It is possible to use a flash-memory USB key to boot your system in order to kickstart it. Refer to the *Red Hat Enterprise Linux System Administration Guide* (available at <http://www.redhat.com/docs/manuals/enterprise/>) for tips on how to do this. Note that your host system's hardware must support boot via these devices.

- Insert the boot CD in the drive and reboot the system, making sure the CD-ROM drive is set as the primary boot device in the system's BIOS.
- After reboot, you should find yourself at a boot prompt. Type the following command at this prompt to start your kickstart:

```
linux \  
ks=http://your-satellite.example.com/ks/label/the_profile_label_you_created_earlier
```




Note

For some systems you may either need to add `ksdevice=eth0` to the command above or disable one of two or more NICs in the system's BIOS to avoid confusion during the kickstart process.

4. The kickstart for your host system should begin. It should take around fifteen minutes to complete. Upon successful completion of this kickstart, you will have provisioned a host system for your virtual guest and registered it to your Satellite.

10.1.2.3. Your Host System Has Red Hat Enterprise Linux 5 Installed

You should register your host system to your Satellite and check to see if the required `xen` packages are installed on the system. If they are not, install them using the Satellite

1. First, register your host system to your Satellite. Use `ssh` to connect to your host system. Register your host system to your satellite issuing the following command as root:

```
rhnreg_ks --serverUrl=http://your-satellite.example.com/XMLRPC \  
--username=username --password=password
```



Note

If your host system is already registered to a different Red Hat Network server, add the `--force` to the command above.

2. Next, open up the host system's profile in the Satellite web interface. Log into the web interface of your Satellite at <https://your-satellite.example.com/>. Click on the **Systems** tab in the top navigational bar. You should see the host system you just registered - click on its profile name to access its system profile page.
3. Make sure your system has access to the software channels it needs to access the software required for hosting virtual guests. From your host system's profile page, click on the **Alter Channel Subscriptions** link on the profile page under the **Subscribed Channels** header. Check the **RHEL Virtualization** and **Red Hat Network Tools for RHEL Server** checkboxes and click the **Change Subscriptions** button underneath the list of channels.
4. Next, check to see if you have the necessary software installed for hosting virtual guest on the system. On the host system, issue the following command as root:

```
rpm -q xen kernel-xen rhn-virtualization-host
```

If `rpm` indicates these packages are not installed, you must install them by running the following command as root on the system:

```
yum install xen kernel-xen rhn-virtualization-host
```

You will then need to edit the `/etc/grub.conf` configuration file to boot the new xen kernel by default. To do this, select the lines in `grub.conf` that pertain to the xen kernel from the beginning of the `title` line to the end of the `initrd` line, copy the lines, delete them, and paste them so they are the first kernel entry in `grub.conf`. Also ensure that the value of the default variable at the top of `grub.conf` is set to a value of '0'.



Note

If you ever update the kernel on the host system, the standard kernel is the default choice upon reboot. To ensure that the Xen kernel is chosen by default, change the following value in the `/etc/sysconfig/kernel` file:

```
DEFAULTKERNEL=kernel
```

Change the value to `kernel-xen`:

```
DEFAULTKERNEL=kernel-xen
```

5. Reboot the system, boot it into the xen kernel. The system should not automatically boot into the xen kernel on reboot but if you would like to make sure it has for troubleshooting purposes, use the command `uname -r` to see if the running kernel is a xen kernel. If you do not see the `xen` string in the name of the kernel, you have not booted into the correct kernel.



Note

If the system already has `xen` and `kernel-xen` installed you do not need to reboot after installing `rhn-virtualization-host`.

6. You will also need to install and run the `osad` package in order for your host system to be responsive to commands sent from the Satellite, such as start, pause, resume, and shutdown. To install:

```
yum install -y osad
```

after installation, you should then start the `osad` process:


```
/sbin/service osad restart
```


7. Your host system should now be ready for RHN virtual guest provisioning.

10.2. Setting Up Your Virtual Systems


In order to work with virtual guest systems, you must first create a kickstart profile that will allow you to easily provision virtual guests, then you must provision the guests.

10.2.1. Create a Kickstart Profile for the Guest Systems

1. Log on to the Satellite's web interface. Navigate to the **Kickstart Overview** screen by clicking on the **Manage Kickstarts** link in the **Tasks** widget in **Overview**, or by clicking on **Systems** in the top navigation bar  **Kickstart** from the left navigation bar.
2. On the **Kickstart Overview** page, click the **Create a new Kickstart Profile** link in the **Kickstart Actions** widget in the upper right corner.
3. The next page displayed is Step 1 of the kickstart profile creation process:
 - a. Enter a label for the profile that will allow you to distinguish it from the other profiles. A good choice would be **guest-system**.
 - b. For the **Base Channel** field, select **Red Hat Enterprise Linux \$PRODUCT (v.5 for \$ARCH)** where \$ARCH is the architecture of your host system's operating system and \$PRODUCT is either Server or Client.




Note
Red Hat Enterprise Linux Client 5 may not be available for selection if you did not sync the Client software channels to your Satellite.



Tip
Please note that the channel labels for Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 5 Desktop refer to 'server' and 'client' respectively.

- c. For the **Kickstartable Tree** field, you should select **ks-rhel-\$ARCH-\$PRODUCT-5** where \$ARCH is the architecture of your host system and \$PRODUCT is either 'server' or 'client', depending on which product with which you would like to provision your guest.
- d. Select **Para-Virtualized Guest** for the **Virtualization Type** field.



Note
If you are changing the **Virtualization Type** of an existing kickstart profile, it may also modify the bootloader and partition options, potentially overwriting any user customizations. Be sure to review the **Partitioning** tab to verify these settings when changing the **Virtualization Type**.

- e. Finally, click the **Next** button in the lower right of the screen to continue on to the next step.
4. For Step 2 of the kickstart profile creation process, select the location of the distribution files for the installation of your guest system. There should already be a **Default Download Location** filled out and selected for you on this screen. Click the **Next** button on this screen to continue to Step 3.



Note

As in the previous step, if the default download location is missing, you may not have successfully synced software channel content to your Satellite from Red Hat's servers.

5. For Step 3 of the kickstart profile creation process, choose a root password for the guest system you are provisioning, and click **Next** to finish creation of the profile.
6. This completes kickstart profile creation. After completing Step 3 you should be taken to the profile details. You may browse through the various tabs of the profile and modify the settings as you see fit, but this is not necessary as the default settings should work well for the majority of cases. While the interface allows you to allocate less, we strongly recommend allocating at least 2 GB of storage for your guest system with this kickstart profile.

10.2.2. Provision Your Guest Systems

1. Log into the Satellite's web interface. Browse to your host system's profile by clicking on the **Systems** tab in the top navigation bar, and click on the system's name.
2. To schedule a kickstart for a guest system, go to the **Virtualization** ▾ **Provisioning** tab in the host system's profile. For the **Guest Name** field choose **guest1**. For the **Memory Allocation**, **Virtual CPUs**, and **Storage** fields, the default values should be fine. Feel free to change these as desired, taking note of the advice provided for each field in the interface. For the **Kickstart Profile** field, select the guest system profile we created in the last step.
3. Finally, click on the **Schedule Kickstart and Finish** button in the lower-right corner of the screen. You will be taken to the **Kickstart Status** page where you can follow along with the guest's kickstart progress. After ten to fifteen minutes the status screen should indicate the kickstart successfully completed. To view your new guest, click on the **Virtualization** tab of the host system's profile on the Satellite. To view a list of virtual systems, navigate to **Systems** ▾ **Systems** ▾ **Virtual Systems**.



Note

If you do not see the **Initiate a kickstart for a Xen** guest message on the **Kickstart Status** page shortly after scheduling the kickstart of the guest, you may be missing **osad** on your host.

Host systems require the **osad** package in order to be responsive to commands sent from the Satellite, such as start, pause, resume, and shutdown. If **osad** is not installed and running, the host system will not receive these commands from the web interface for 2.5 hours, or the next time that the RHN daemon runs.

You can check whether or not **osad** is installing and running by checking the **OSA Status** field in the host system's profile on the Satellite. If the field does not exist or indicates a failure of that the system has not contact Satellite in several minutes, then you will need to install **yum** (using the command **yum install -y osad**) before you can successfully provision a guest on the host.



Tip

You may receive the following message from the **Kickstart Status** page during the guest's kickstart:

```
The install process on the guest system has not communicated to RHN in the past n minutes. This may be due to a hung install process, or it may just be due to a slow install because of hardware constraints. A log of the installation process is available, you may wish to review it to troubleshoot this issue.
```

Be patient and do not worry if you see this message unless more than twenty minutes have passed. To check if the kickstart is continuing, check the installation log to make sure there are no errors, and as you reload the Kickstart Status page check to see that the Last File Request field continues to be updated.

4. If you would like to register additional guests to your host, repeat the steps above. It is important to remember that you can only provision one guest at a time. If you attempt to schedule a guest kickstart while another is currently taking place, the current guest kickstart process will be canceled and the new guest kickstart process will begin.
5. View your newly-created virtual guest's system in the Satellite's web interface by clicking on the **Virtualization** tab in the host system's profile. Then, click on the profile name of your virtual system. You will be brought to its Satellite system profile.

10.3. Working With Your Virtual Systems

Once you have set up your virtual system, you can then manage and customize them via various methods, including connecting via SSH and via the virtualization management interface on the host system.

10.3.1. Logging into Virtual Systems Directly via SSH

1. You will need to locate the virtual system's IP address. Locate it by navigating to the **Systems** ▾ **Virtual Systems** tab and clicking on the virtual system's profile name.
2. On the virtual system's profile page, you'll find the IP address in the left-hand informational column in the **IP Address** field.
3. Connect to the IP address by using **ssh** as root, using the password you set for the virtual system in the kickstart profile you created for it earlier.

10.3.2. Gaining Console Access Via the Host

1. First you will need to connect to the host system and determine the ID number of the guest you would like to work with. Connect to the host system via **ssh** and run the following command:

```
xm list
```

This should provide you with a list all of the guests you created on your Satellite, including their ID number. Look for the guest, **guest1**, that we created earlier in this list. If, for example, this guest has been assigned an ID of 2, then:

2. Run the following command to access the console of this virtual system:

```
xm console 2
```

You should immediately be able to view a login prompt on **guest1**.

3. Login to **guest1** as root using the same password you set in the kickstart profile you used to provision the system.

(There may be some messages on the screen. In this case, hit the **Enter** key on your keyboard to receive a fresh login prompt.)

4. To exit the guest console and return to the host system's command prompt, you may hit the **Ctrl** and **]** keys on your keyboard simultaneously.

10.3.3. Installing Software Via the Satellite Web Interface

1. Browse to the virtual system's profile in your Satellite's web interface by logging in and navigating to **Systems** ▾ **Systems** ▾ **Virtual Systems** and clicking on the name of your virtual system's profile.
2. In the virtual system's profile, click on the **Software** ▾ **Packages** tab.
3. Click on **Install New Packages** in the **Packages** tab menu.
4. Select the packages you wish to install and click the **Install Selected Packages** button in the lower right-hand corner of the screen.
5. Review the package install details and click on the **Confirm** button in the lower right-hand corner of the screen.
6. The package install will take place the next time the guest system checks in with the Satellite. To force the install to take place immediately, you may run the **rhn_check** command on the guest system.

10.3.4. Installing Software Via Yum From the Virtual System

Your virtual system registered to your Satellite as part of the guest provisioning process, so you may simply use the **yum** command to install and update software. For example, to install the text editor vim, issue the following command:

```
yum install -y vim-enhanced
```

10.3.5. Restarting Guests when Host Reboots

By default, when a host system reboots, the guests are not restarted and must be manually started by the administrator.

However, the **rhn-virtualization-host** service can restart guests automatically in the event of a host system reboot.

To use this service, follow these steps:

1. Locate the guest's config file on the host in `/etc/sysconfig/rhn/virt/`. It will be named by UUID, but the correct file can be found by using the **grep** command to search for the guest name within the UUID files.
2. When you have found the UUID file corresponding to your guest system, create a symbolic link from the UUID file to the `/etc/sysconfig/rhn/virt/auto/` directory.

```
ln -s /etc/sysconfig/rhn/virt/GUEST_UUID.xml /etc/sysconfig/rhn/virt/auto/
```

10.3.6. Deleting Virtual Systems

Deleting a virtual system is a multi-step process.

1. First you must shut down the virtual system that you wish to delete. You may do this by browsing to the host system's profile in the Satellite web interface, clicking on the virtualization tab, and checking off the virtual systems that you would like to delete. Finish shutting down by clicking the **Shutdown Systems** button at the bottom of the screen.
2. Next, delete the virtual system from Satellite. This is accomplished by checking off the virtual system's checkbox and clicking the **Delete System** button at the bottom of the screen>



Tip

Please allow for at least two minutes between shutting down a virtual system and deleting it. Otherwise, the virtual system may not shut down properly and you will delete it while it is running. If you delete a virtual system from Satellite while it is running, it will reappear on the Satellite the next time it checks in. If this happens, simply shutdown the system, wait two minutes, and delete it again.

3. Delete the disk image for the virtual system you would like to delete. You will find the disk image for **guest1**, for example, at the following location on the host system:

```
/var/lib/xen/disk-images/guest1.disk
```

Delete it with the following command:

```
rm /var/lib/xen/disk-images/guest1.disk
```

4. Finally, you must delete the RHN configuration files from the host system. To locate the RHN configuration file for **guest1**, run the following command:

```
grep guest1 /etc/sysconfig/rhn/virt/*.xml
```

Then delete the file indicated. For example:

```
rm /etc/sysconfig/rhn/virt/14e5cfbf72342515236ad74b260c2f6b.xml
```

5. You have successfully deleted a guest system from your host system and from Satellite.

Cobbler

RHN Satellite features the *Cobbler* server that allows administrators to centralize their system installation and provisioning infrastructure. Cobbler is an installation server that collects the various methods of performing unattended system installations, whether it be server, workstation, or guest systems in a full or para-virtualized setup.

Cobbler has several tools to assist in pre-installation guidance, kickstart file management, content channel management, and more. Features of Cobbler include:

- Installation environment analysis using the **cobbler check** command
- Multi-site installation server configuration with **cobbler replicate**
- Automation of kickstart file creation using **kickstart import**
- Kickstart template creation and management using the Cheetah template engine and Kickstart Snippets
- Virtual machine guest installation automation with the **koan** client-side tool.

11.1. Cobbler Requirements

To use Cobbler as a PXE boot server, you should check the following guidelines:

- If you plan to use Cobbler to install systems using PXE, you must have **tftp-server** installed and configured.
- If you plan to use Cobbler to PXE boot systems for installation, you must have either the ability to act as a DHCP server for Cobbler PXE booting or access to your network DHCP server **/etc/dhcp.conf** to change **next-server** to the hostname or IP address of your Cobbler server.

11.1.1. Using **cobbler check**

Once your Satellite server is installed, you can check that your system is correctly configured to run as a proper Cobbler boot server.

The **cobbler check** command can be run at a shell prompt as root. It will automatically run checks on several services' runtime status and configuration files for proper configuration. The following lists the services and configuration files that **cobbler check** tests.

- **/etc/cobbler/settings** — The central cobbler configuration file; **cobbler check** will check to see if you have set the **server** setting to the addressable IP or hostname of the Cobbler server, as well as the **next_server** field set to the hostname or IP of the Cobbler server if it will be also be managing DHCP services for PXE boots.
- SELinux — Will verify that you have the correct SELinux setting for HTTPD services and the proper content rules for **tftp** and files in **/var/www/cobbler/images/**.
- **cobblerd** — Will check to see if you have the Cobbler daemon running.
- **xinetd** — Will inform you that Xinetd services running and that you have changed the parameter *disabled* to **no** in **/etc/xinetd.d/tftp**.
- **httpd** — Will check to see if the HTTPD service is running.

- **iptables** — Will remind you that if you are running an IPTables firewall, that you have rules set to allow ports 69 (TFTP), 80 (HTTPD), 25150 and 25151 (Cobbler).

Run the **cobbler check** command as root on your system to see what settings and services need to be enabled to properly run Cobbler on your boot server.



Note

If you make any changes to the `/etc/cobbler/settings` file, you must run **service cobblerd restart** and **cobbler sync** for the changes to take effect.

11.1.2. Configuring Cobbler with `/etc/cobbler/settings`

Cobbler configuration is mainly done within the `/etc/cobbler/settings` file. The file contains several configurable settings, and offers detailed explanations for each setting regarding how it affects the functionality of Cobbler and whether it is recommended for users to change the setting for their environment.

Most of the settings can be left default and Cobbler will run as intended. For more information about configuring Cobbler settings, consult the `/etc/cobbler/settings` file, which documents each setting in detail.

11.1.3. Cobbler and DHCP

Cobbler supports bare-metal kickstart installation of systems configured to perform network boots using a PXE boot server. To properly implement a Cobbler installation server, administrators need to either have administrative access to the network's DHCP server or implement DHCP on the Cobbler server itself.

11.1.3.1. Configuring an Existing DHCP Server

If you have a DHCP server deployed on another system on the network, you will need administrative access to the DHCP server in order to edit the DHCP configuration file so that it points to the Cobbler server and PXE boot image.

As root on the DHCP server, edit the `/etc/dhcpd.conf` file and append a new class with options for performing PXE boot installation. For example:

```
allow booting;
allow bootp;
class "PXE" {
match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
next-server 192.168.2.1;
filename "pxelinux.0";
}
```

Following each action step-by-step in the above example:

1. The administrator enables network booting with the **bootp** protocol.
2. Then, the administrator creates a class called **PXE**, which, if a system that is configured to have PXE first in its boot priority, identifies itself as **PXEClient**.

3. Then DHCP server then directs the system to the Cobbler server at 192.168.2.1.
4. Finally, the DHCP server refers to the boot image file (in this case, at `/var/lib/tftpboot/pxelinux.0`).

11.1.4. Xinetd and TFTP

Xinetd is a daemon that manages a suite of services, including TFTP, the FTP server used for transferring the boot image to a PXE client.

To configure TFTP, you must first enable the service via Xinetd. To do this, edit the `/etc/xinetd.d/tftp` as root and change the `disable = yes` line to `disable = no`.

Before TFTP can start serving the `pxelinux.0` boot image, you must start the Xinetd service.

```
chkconfig --level 345 xinetd on
/sbin/service xinetd start
```

The `chkconfig` command turns on the `xinetd` service for all user runlevels, while the `/sbin/service` command turns on `xinetd` immediately.

11.1.5. Configuring SELinux and IPTables for Cobbler Support

Red Hat Enterprise Linux is installed with SELinux support in addition to secure firewall enabled by default. To properly configure a Red Hat Enterprise Linux server to use Cobbler, you must first configure these system and network safeguards to allow connections to and from the Cobbler Server.

11.1.5.1. SELinux Configuration

To enable SELinux for Cobbler support, you must set the SELinux boolean to allow HTTPD web service components. Run the following command as root on the Cobbler server:

```
setsebool -P httpd_can_network_connect true
```

The `-P` switch is essential, as it enables HTTPD connection persistently across all system reboots.

You must also set SELinux file context rules to ensure Cobbler properly functions in an SELinux system.

Run the following as root on the Cobbler server:

```
semanage fcontext -a -t public_content_t "var/lib/tftpboot/*"
```

The command sets file context for TFTP to serve the boot image file.

11.1.5.2. IPTables Configuration

Once you have configured SELinux, you must then configure IPTables to allow incoming and outgoing network traffic on the Cobbler server.

If you have an existing firewall ruleset using IPTables, you need to add the following rules to open the requisite Cobbler-related ports. The following lists each of the requisite rules with their associated service.

- For TFTP:

```
/sbin/iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 69 -j ACCEPT
/sbin/iptables -A INPUT -m state --state NEW -m udp -p udp --dport 69 -j ACCEPT
```

- For HTTPD:

```
/sbin/iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 80 -j ACCEPT
/sbin/iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 443 -j ACCEPT
```

- For Cobbler:

```
/sbin/iptables -A INPUT -m state --state NEW -m tcp -p udp --dport 25150 -j ACCEPT
```

- For Koan:

```
/sbin/iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 25151 -j ACCEPT
```

Once those firewall rules are entered, be sure to save the firewall configuration:

```
/sbin/iptables-save
```

11.1.6. Syncing and Starting the Cobbler Service

Once all the prerequisites specified in **cobbler check** are met, you can now start the Cobbler service.

First, ensure that the configuration files are all synchronized by running the following command:

```
cobbler sync
```

Then, start the Satellite server:

```
/usr/sbin/rhn-satellite start
```



Warning

Do not start or stop the **cobblerd** service independent of the Satellite service, as doing so may cause errors and other issues.

Always use **/usr/sbin/rhn-satellite** to start or stop RHN Satellite.

11.2. Adding a Distribution to Cobbler

If all Cobbler prerequisites have been met and Cobbler is now running, you can now begin adding a distribution to the Cobbler if you have the content on the Cobbler server.

For information about creating and configuring kickstart distributions from the RHN Satellite interface, refer to [Section 7.4.9.6, “Kickstart Distributions — !\[\]\(bd1a142de767a21e5362c595f844a4ff_img.jpg\) ”](#).

Using **cobbler** to create a distribution from the command line is as follows:

```
cobbler distro add --name=string --kernel=path --initrd=path
```

The **--name=*string*** switch is a label used to differentiate one distro choice from another (for example, **rhel5server**)

The **--kernel=*path*** switch specifies the path to the kernel image file

The **--initrd=*path*** switch specifies the path to the initial ramdisk (initrd) image file.

11.3. Adding a Profile to Cobbler

Once you have configured a distribution to Cobbler, you can then add profiles to Cobbler.

Cobbler profiles associate a distribution to additional options, like kickstart files. Profiles are the core unit of provisioning and there must be at least one Cobbler profile for every distribution added. For example, two profiles might be created for a web server and a desktop configuration. While both profiles use the same distro, the profiles are for different installations types.

For information about creating and configuring kickstart profiles from the RHN Satellite interface, refer to [Section 7.4.9.2, “Kickstart Profiles”](#).

The usage of **cobbler** to create profiles from the command line is as follows:

```
cobbler profile add --name=string --distro=string [--kickstart=url] [--virt-file-size=gigabytes] [--virt-ram=megabytes]
```

The **--name=*string*** is the unique label for the profile, such as **rhel5webserver** or **rhel4workstation**.

The **--distro=*string*** switch specifies the distribution that will be used for this particular profile. Distributions were added in [Section 11.2, “Adding a Distribution to Cobbler”](#).

The **--kickstart=*url*** option specifies the location of the kickstart file (if available).

The **--virt-file-size=*gigabytes*** option allows you to set the size of the virtual guest file image. The default is 5 gigabytes if left unspecified.

The **--virt-ram=*megabytes*** option specifies how many megabytes of physical RAM that a virtual guest system can consume. The default is 512 megabytes if left unspecified.

11.4. Adding a System to Cobbler

Once the distributions and profiles for Cobbler have been created, you can next add systems to Cobbler. System records map a piece of hardware on a client with the cobbler profile assigned to run on it.



Note

If you are provisioning via **koan** and PXE menus alone, it is not required to create system records, though they are useful when system-specific kickstart templating is required or to establish that a specific system should always receive a specific content installed. If there is a specific role intended for a specified client, system records should be created for it.

For information about creating and configuring kickstarts from the RHN Satellite interface, refer to

[Section 7.4.2.9.4, “System Details ▯ Provisioning — !\[\]\(10f8862fc183b400327470ea85afe9ae_img.jpg\) ”](#).

The following command adds a system to the Cobbler configuration:

```
cobbler system add --name=string --profile=string --mac=AA:BB:CC:DD:EE:FF
```

The **--name=string** is the unique label for the system, such as **engineeringserver** or **frontofficeworkstation**.

The **--profile=string** specifies one of the profile names added in [Section 11.3, “Adding a Profile to Cobbler”](#).

The **--mac=AA:BB:CC:DD:EE:FF** option allows systems with the specified MAC address to automatically be provisioned to the profile associated with the system record if they are being kickstarted.

For more options, such as setting hostname or IP addresses, refer to the Cobbler manpage by typing **man cobbler** at a shell prompt.

11.5. Cobbler Templates

Within the RHN Satellite web interface, there are facilities to create variables for use with kickstart distributions and profiles. For example, to create a kickstart profile variable, refer to [Section 7.4.9.3.3, “Kickstart Details ▯ Variables”](#).

Kickstart variables are a part of an infrastructural change in Satellite to support *templating* in kickstart files. In the context of kickstart files, templates are files that hold descriptions used to build actual kickstart files, rather than creating specific kickstarts.

These templates are then shared by various profiles and systems that have their own variables and corresponding values. These variables modify the templates and software called a *template engine* parses the template and variable data into a usable kickstart file. Cobbler uses an advanced template engine called *Cheetah* that provides support for templates, variables, and snippets.

Advantages of using templates include:

- Robust features that allow administrators to create and manage large amounts of profiles or systems without duplication of effort or manually creating kickstarts for every unique situation

- While templates can become complex and involve loops, conditionals and other enhanced features and syntax, it can also be used simply to make kickstart files without such complexity.

11.5.1. Using Templates

Kickstart templates can have static values for certain common items such as PXE image filenames, subnet addresses, and common paths such as `/etc/sysconfig/network-scripts/`. However, where templates differ from standard kickstart files are in their use of variables.

For example, a standard kickstart file may have a networking passage that looks similar to the following

```
network --device=eth0 --bootproto=static --ip=192.168.100.24 --netmask=255.255.255.0 --
gateway=192.168.100.1 --nameserver=192.168.100.2
```

However, in a kickstart template file, the networking passage may look similar to the following:

```
network --device=$net_dev --bootproto=static --ip=$ip_addr --netmask=255.255.255.0 --gateway=
$my_gateway --nameserver=$my_nameserver
```

These variables will be substituted with the values set in your kickstart profile variables or in your system detail variables. If there are the same variables defined in both the profile and the system detail, then the system variable takes precedence.

For more information about kickstart templates, refer to the Cobbler project page at the following URL:

<https://fedorahosted.org/cobbler/wiki/KickstartTemplating>

11.5.2. Kickstart Snippets

If you have common configurations that are the same across all kickstart templates and profiles, you can utilize the *Snippets* feature of Cobbler to take advantage of code reuse.

Kickstart snippets are sections of kickstart code that can be called by a `$SNIPPET()` function that will be parsed by Cobbler and substitute that function call with the contents of the snippet.

For example, if you had a common hard drive partition configuration for all servers, such as:

```
clearpart --all
part /boot --fstype ext3 --size=150 --asprimary
part / --fstype ext3 --size=40000 --asprimary
part swap --recommended

part pv.00 --size=1 --grow

volgroup vg00 pv.00
logvol /var --name=var vgroupname=vg00 --fstype ext3 --size=5000
```

You could take that snippet, save it to a file (such as `my_partition`), and place the file in `/var/lib/cobbler/snippets/` so that Cobbler can access them.

You can then use the snippet by using the `$SNIPPET()` function in your kickstart templates. For example:

```
$$SNIPPET('my_partition')
```

Wherever you invoke that function, the Cheetah parser will substitute the function with the snippet of code contained in the **my_partition** file.

For more information about kickstart snippets, refer to the Cobbler project page at the following URL:

<https://fedorahosted.org/cobbler/wiki/KickstartSnippets>

11.6. Using Koan

Whether you are provisioning guests on a virtual machine or reinstalling a new distribution on a running system, koan works in conjunction with Cobbler to provision systems on the fly.

11.6.1. Using Koan to Provision Virtual Systems

If you have created a virtual machine profile as documented in [Section 11.3, “Adding a Profile to Cobbler”](#), you can use **koan** to initiate the installation of a virtual guest on a system.

For example, say you've created a Cobbler profile such as the following:

```
cobbler add profile --name=virtualfileserv er --distro=rhel-i386-server-5 --virt-file-size=20
--virt-ram=1000
```

This profile is for a fileserver running Red Hat Enterprise Linux 5.3 with a 20GB guest image size and allotted 1GB of system RAM.

To find the name of the virtual guest system profile, run the following with **koan**:

```
koan --server=hostname --list=profiles
```

This command lists all of the available profiles created with **cobbler profile add**.

Then, begin the process of creating the image file and starting the installation of the virtual guest system.

```
koan --virt --server=cobbler-server.example.com --profile=virtualfileserv er --
virtname=marketingfileserv er
```

The command specifies that a virtual guest system be created from the Cobbler server (hostname `cobbler-server.example.com`) using the **virtualfileserv er** profile. The **virtname** option specifies a label for the virtual guest, which by default is labeled with the system's MAC address.

Once installation of the virtual guest is complete, it can be used as any other virtual guest system.

11.6.2. Using Koan to Re-install Running Systems

There may be instances where you need to re-install a machine with another operating system while it is still running. **koan** can help you by destructively replacing a running system with a new installation from the available Cobbler profiles.

To replace a running system and install a new one, run the following command *on the system itself*:

```
koan --replace-self --server=hostname --profile=name
```

This command, when executed on the running system to be replaced, will start the provisioning process and replace its own system using the profile in **--profile=name** on the Cobbler server specified in **--server=hostname**.

UNIX Support Guide

12.1. Introduction

This chapter documents the installation procedure for, and identifies differences in, Red Hat Network functionality when used to manage UNIX-based client systems. RHN offers UNIX support to help customers migrate from UNIX to Linux. Because of the limited scope of this task, the features offered for UNIX client management are not as comprehensive as those available for managing Red Hat Enterprise Linux systems.

Subsequent sections specify supported UNIX variants, RHN features supported by the UNIX management system, the prerequisites for managing a UNIX system with RHN, as well as the installation procedure for UNIX clients.

12.1.1. Supported UNIX Variants

The following UNIX variants, versions, and architectures are supported by RHN Satellite:

Solaris Version	sun4m	sun4d	sun4u	sun4v	sun4us	x86
Solaris 8	yes	no	yes	n/a	no	no
Solaris 9	yes	n/a	yes	n/a	no	yes
Solaris 10	n/a	n/a	yes	yes	no	yes

Table 12.1. Supported Solaris Architectures and Versions

12.1.2. Prerequisites

These items are needed to obtain UNIX support:

- RHN Satellite 5.0.0 or later
- A Satellite certificate with Management entitlements
- Management entitlements for each UNIX client
- RHN packages for UNIX including python, pyOpenSSL, and the Red Hat Network Client packages.
- Sunfreeware packages that provide supporting libraries. Some of these packages are available via the RHN Satellite. Refer to [Section 12.3.1, “Download and Install Additional Packages”](#) for the complete list.

12.1.3. Included Features

The following features are included in the UNIX support service level as they exist within RHN:

- The **Red Hat Network Service Daemon (rhnsd)**, which triggers `rhn_check` according to a configurable interval
- The **Red Hat Network Configuration Client (rhncfg-client)**, which executes all configuration actions scheduled from the Satellite
- The **Red Hat Network Configuration Manager (rhncfg-manager)**, which allows command line administration of RHN configuration channels

- The **rhn_check** program, which checks in with the Satellite and performs any actions scheduled from the server
- All Management-level functionality, such as system grouping, package profile comparison, and use of the System Set Manager to administer multiple systems at once
- A Provisioning feature called *Remote Command* that enables users to schedule root-level commands on any managed client through the Satellite's website, if the client allows this action

12.1.4. Differences in Functionality

The following RHN features work differently in a UNIX environment:

- The **Red Hat Update Agent for UNIX** offers a much smaller set of options than its Linux counterpart and relies upon the operating system's native toolset for package installation, rather than **rpm** - Refer to [Section 12.4.2.4, "Updating From the Command Line"](#) for the precise list of options.
- The **RHN Push** application has been similarly modified to upload native UNIX file types, including packages, patches, and patch clusters.

Since Solaris package, patch and patch cluster files are different from RPM files, the channel upload mechanism is somewhat different. There are two applications in the **rhnpush** package for Solaris:

- The first, **solaris2mpm**, is an RHN utility that create an MPM file for each Solaris package or patch. The neutral format of the MPM file allows the Satellite to understand and manage the uploaded files.
- The second, **rhnpush**, has been extended so that it can handle MPM as well as RPM files. Otherwise, it operates identically to the Linux version of **rhnpush**.
- The **Channels** tab of the RHN website has been augmented to accommodate the storage and installation of native UNIX file types.

12.1.5. Excluded Features

The following RHN features are not available with the UNIX support system:

- All Provisioning-level functionality, such as kickstarting and package rollback, with the exception of configuration file management
- All Errata-related options, since the concept of Errata Updates is not understood in UNIX
- Source files for packages

Answer files are not yet supported. Support for such files is planned for a future release.

Additionally, relocating **RHAT* .pkg** files during installation is not yet supported.

12.2. Satellite Server Preparation/Configuration

You must configure the Satellite to support UNIX clients before the required files are available for deployment to the client systems. This can be accomplished in one of two ways, depending on whether you have yet installed your Satellite server:

1. During the Satellite installation:

Enable UNIX support on the Satellite by checking the "Enable Solaris Support" box during the installation process, as pictured:

The screenshot shows a web browser window with the URL `http://your-satellite.example.com/install/configure.pxt`. The page is titled "Satellite Installation" and features a blue "Install" button at the top. Below the title, there is a paragraph of text explaining the configuration options. The main section is "Red Hat Network Configuration", which contains several form fields and checkboxes:

- Satellite Hostname*:
- HTTP proxy:
- HTTP proxy username:
- HTTP proxy password:
- RPM repository mount point*:
- Enable SSL:
- Enable Solaris Support:
- Disconnected Satellite:
- Enable monitoring backend:
- Enable monitoring scout:

A "Continue" button is located at the bottom right of the form.

Figure 12.1. Enabling UNIX Support During Satellite Installation

2. After the Satellite has been installed:

Enable UNIX support by configuring the Satellite after it has been installed. To do so, select **Satellite Tools** in the top menubar, then select **Satellite Configuration** in the left navigation bar. In the screen that follows, check the **Enable Solaris Support** box, as pictured:

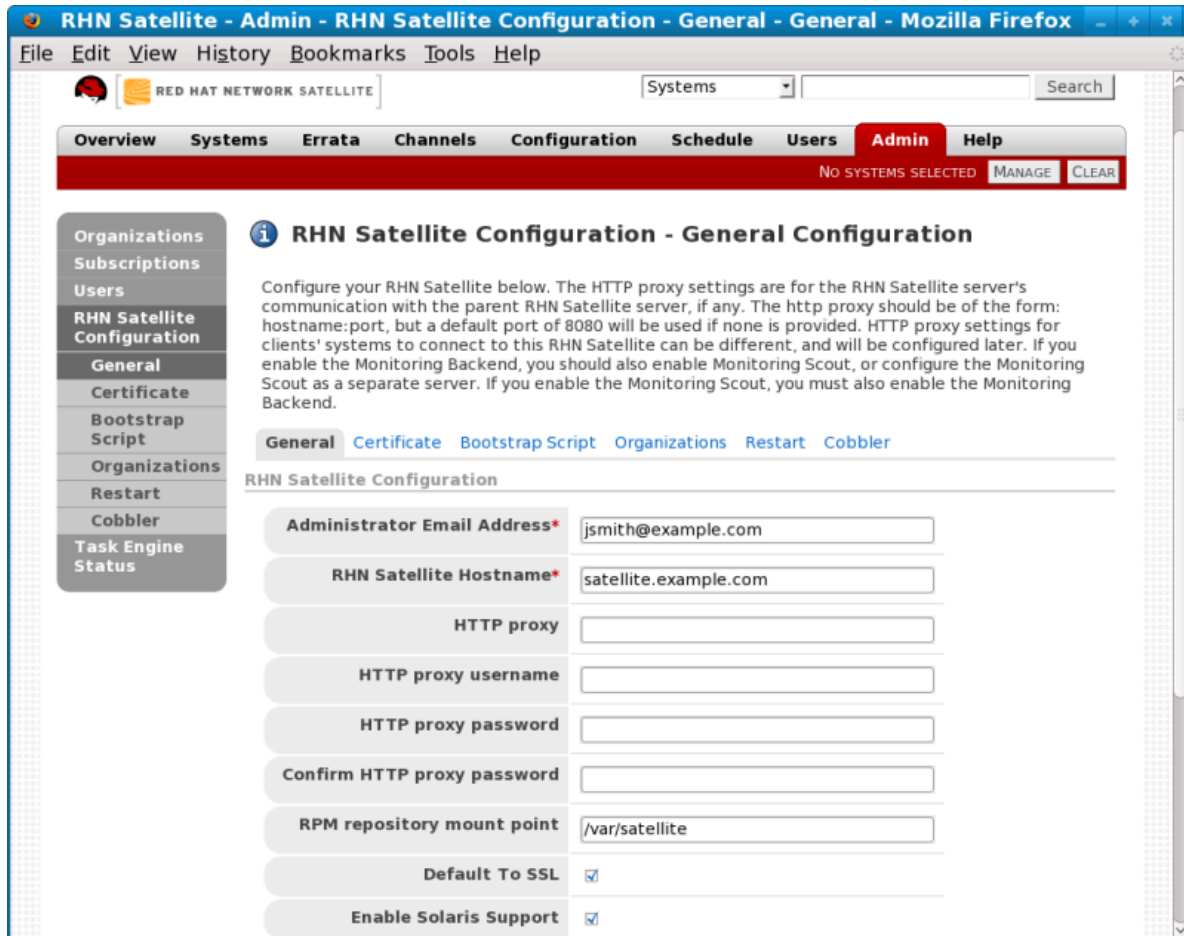


Figure 12.2. Enabling UNIX Support After Satellite Installation

Click the **Update Configuration** button to confirm the change.

3. Finally, you must create a base channel to which your client systems may subscribe. This is because RHN does not provide UNIX content; as a result, you cannot use `satellite-sync` to create the channel.

To create a Solaris channel, login to the web interface of the Satellite as either an Satellite Administrator or a certificate authority. Navigate to the **Channel** tab, followed by the **Manage Software Channels** from the left navigation bar. Click the **create new channel** link in the upper right of the resulting screen. Provide a name and label for your new channel, and select either **Sparc Solaris** or **i386 Solaris** as the architecture, depending on the architecture of your client.

12.3. Client System Preparation

Before your UNIX-based client systems benefit from Red Hat Network, they must be prepared for connection:

1. Download and install `gzip` and required third-party libraries.
2. Download the RHN application tarball from the Satellite to the client and install the contents.
3. Next, deploy the SSL certificates required for a secure connection.

4. Configure the client applications to connect to the RHN Satellite.

Once finished, your systems will be ready to begin receiving RHN updates. The following three sections explain these steps in detail.

12.3.1. Download and Install Additional Packages

This section steps you through the process of downloading and installing third-party applications and the RHN applications from the Satellite onto the UNIX client.

Of primary importance is the **Red Hat Update Agent for UNIX (up2date)**, which provides the link between your client systems and Red Hat Network. The UNIX-specific version of the **Red Hat Update Agent** is limited in functionality compared to its Linux counterpart but still enables system registration and facilitates package installs and patches. Refer to [Section 12.4, "Registration and Updates"](#) for a full description of the tool's options.



Note

It may be useful to enter the command `bash` when first logging into the Solaris client. If the BASH shell is available, it will make the system's behavior as Linux-like as possible.

12.3.1.1. Install Third-Party Packages

Installation of the RHN applications cannot proceed unless the following utility and libraries are present:

- `gzip`
- `libgcc`
- `openssl`
- `zlib`

The `gzip` utility is provided by the `SUNWgzip` package and may be downloaded from <http://www.sunfreeware.com>.

On recent versions of Solaris, the necessary libraries are provided by the following natively installed packages:

- `SUNWgccruntime`
- `SUNWopenssl*`
- `SUNWzlib`

For older Solaris versions, the following required packages may be downloaded from <http://www.sunfreeware.com>:

- `SMClibgcc` or `SMCgcc`
- `SMCoss1`
- `SMCzlib`

To verify if a package is installed on the client, use the **pkginfo** command. For example, to check for a package that contains "zlib" in the name, run the following command:

```
# pkginfo | grep zlib
```



Note

Solaris package archive names differ from the name of the installed package. For example, the package archive **libgcc<version>-sol<solaris-version>-sparc-*local*.gz** becomes **SMClibgcc** after installation

12.3.1.2. Configure the Library Search Path

In order to allow the Solaris client to use the libraries installed in the previous step, you must add their location to the library search path. To do so, first check the current library search path":

```
# crle -c /var/ld/ld.config
```

Make a note of the current Default Library Path. Next, modify the path to also include the components shown below. Note that the **-l** option resets the value, rather than appending it, so if there already were values set on your system, prepend them to the **-l** parameter.

On sparc:

```
# crle -c /var/ld/ld.config -l /other/existing/path:/lib:/usr/lib:/usr/local/lib
```

On x86:

```
# crle -c /var/ld/ld.config -l /other/existing/path:/lib:/usr/lib:/usr/local/lib:/usr/sfw/lib
```

12.3.1.3. Download RHN Client Packages

Download the appropriate tarball of packages from the **/var/www/html/pub/** directory of your Satellite. If you are able to use a GUI web browser like Mozilla, navigate to the **/pub** directory of the Satellite and save the appropriate tarball to your client:

```
http://your-satellite.example.com/pub/rhn-solaris-bootstrap-<version>-<solaris-arch>-<solaris-version>.tar.gz
```

If you must download the tarball from the command line, it should be possible to use **ftp** to transfer the file from the Satellite to the client.

Using **gzip**, decompress the tarball. You should have the following packages:

- **RHATposs1**
- **RHATrcfg**

- RHATrcfga
- RHATrcfgc
- THATrcfgm
- RHATrhnc
- RHATrhnl
- RHATrpush
- RHATsmart

SMClibgcc and **SMCosslg** may also be included in the tarball.

12.3.1.4. Install the RHN Packages

Change to the uncompressed directory and use the UNIX variant's native installation tool to install each package. For example, on Solaris, use the **pkgadd** command. Answer "yes" to any prompts during package install.

Here is how a typical installation might proceed:

```
# pkgadd -d RHATposs1-0.6-1.p24.6.pkg all
# pkgadd -d RHATpythn-2.4.1-2.rhn.4.sol9.pkg all
# pkgadd -d RHATrhnl-1.8-7.p23.pkg all
...
```



Note

You may choose to use the **-n** of **pkgadd**, which runs the command in non-interactive mode. However, this may cause the installation of some packages to fail silently on Solaris 10.

Continue until each package is installed in the RHN-specific path: **/opt/redhat/rhn/solaris/**.

12.3.1.5. Include RHN Packages in the PATH

In order to make the RHN packages available at each login, you may wish to add them to your PATH. To do so, add these commands to your login script:

```
# PATH=$PATH:/opt/redhat/rhn/solaris/bin
# PATH=$PATH:/opt/redhat/rhn/solaris/usr/bin
# PATH=$PATH:/opt/redhat/rhn/solaris/usr/sbin
# export PATH
```

To enable access to the RHN client command man pages, add them to your MANPATH. To do so, add the following commands to your login script:

```
# MANPATH=$MANPATH:/opt/redhat/rhn/solaris/man
# export MANPATH
```

Alternatively, you can also access the man pages from the command line, with the following command:

```
# man -M /opt/redhat/rhn/solaris/man <man page>
```

Finally, add the Red Hat Libraries to your PATH as you did with **libgcc**, **openssl** and **zlib**.

```
crle -c /var/ld/ld.config -l <current library paths>:/opt/redhat/rhn/solaris/lib
```

12.3.2. Deploying Client SSL Certificates

To ensure secure data transfer, Red Hat strongly recommends the use of SSL. The RHN Satellite eases implementation of SSL by generating the necessary certificates during its installation. The server-side certificate is automatically installed on the Satellite itself, while the client certificate is placed in the **/pub/** directory of the Satellite's Web server.

To install the certificate, follow these steps for each client:

1. Download the SSL certificate from the **/var/www/html/pub/** directory of the RHN Satellite onto the client system. The certificate will be named something similar to **RHN-ORG-TRUSTED-SSL-CERT**. It is accessible via the web at the following URL: **https://your-satellite.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT**.
2. Move the client SSL certificate to the RHN-specific directory for your UNIX variant. For Solaris, this can be accomplished with a command similar to:

```
mv /path/to/RHN-ORG-TRUSTED-SSL-CERT /opt/redhat/rhn/solaris/usr/share/rhn/
```

When finished, the new client certificate will be installed in the appropriate directory for your UNIX system. If you have a large number of systems to prepare for RHN management, you may script this entire process.

Now you must reconfigure the RHN client applications to refer to the newly installed SSL certificate. Refer to [Section 12.3.3, "Configuring the clients"](#) for instructions.

12.3.3. Configuring the clients

The final step before registering your client systems with Red Hat Network is to reconfigure their RHN applications to use the new SSL certificate and obtain updates from the RHN Satellite. Both of these changes can be made by editing the configuration file of the **Red Hat Update Agent**, which provides registration and update functionality.

Follow these steps on each client system:

1. As root, change to the RHN configuration directory for the system. For Solaris, the full path is **/opt/redhat/rhn/solaris/etc/sysconfig/rhn/**.
2. Open the **up2date** configuration file in a text editor.
3. Find the *serverURL* entry and set its value to the fully qualified domain name (FQDN) of your RHN Satellite:

```
serverURL[comment]=Remote server URL
serverURL=https://your-satellite.example.com/XMLRPC
```

4. Ensure the application refers to the RHN Satellite even when SSL is turned off by also setting the `noSSLServerURL` value to the Satellite:

```
noSSLServerURL[comment]=Remote server URL without SSL
noSSLServerURL=http://your-satellite.example.com/XMLRPC
```

5. With the `up2date` configuration file still open, find the `sslCACert` entry and set its value to the name and location of the SSL certificate described in [Section 12.3.2, “Deploying Client SSL Certificates”](#), for example:

```
sslCACert[comment]=The CA cert used to verify the ssl server
sslCACert=/opt/redhat/rhn/solaris/usr/share/rhn/RHN-ORG-TRUSTED-SSL-CERT
```

Your client systems are now ready for registration with Red Hat Network and management by your Satellite.

12.4. Registration and Updates

Now that you have installed RHN-specific packages, implemented SSL, and reconfigured your client systems to connect to the RHN Satellite, you are ready to begin registering systems and obtaining updates.

12.4.1. Registering Systems

This section describes the RHN registration process for UNIX systems. You must use the `rhnreg_ks` command to accomplish this; the use of activation keys for registering your systems is optional. These keys allow you to predetermine settings within RHN, such as base channels and system groups, and to apply those automatically to systems during their registration.

Since activation key generation and use is covered extensively in other chapters, this section focuses on differences when applying them to UNIX variants. Refer to [Section 7.4.6.1, “Managing Activation Keys”](#) for full descriptions of this process.

To register UNIX systems with your RHN Satellite, accomplish the following tasks in this order:

1. Log into the Satellite's web interface and click the **Systems** tab in the top navigation bar followed by **Activation Keys** in the left navigation bar. Then click the **create new key** link at the top-right corner of the page.
2. On the following page, select the base channel you created at the end of [Section 12.2, “Satellite Server Preparation/Configuration”](#).
3. After creating the key, click its name in the **Activation Keys** list to enhance its RHN settings by associating software and configuration channels and system groups.
4. Open a terminal on the client system to be registered and switch user to root.

5. Use `rhncg_ks` along with the `--activationkey` option to register the client with the Satellite. The string of characters that make up the key may be copied directly from the **Activation Keys** list on the website. The resulting command will look something like the following:

```
rhncg_ks --activationkey=b25fef0966659314ef9156786bd9f3af
```

6. Go back to the website, click the name of the activation key, and ensure the new system appears within the **Activated Systems** tab.

12.4.2. Obtaining Updates

Package updates in UNIX are handled much differently than in Linux. For instance, Solaris relies on Patch Clusters to update multiple packages at once, while Red Hat operating systems use Errata Updates to associate upgrades with specific packages. In addition, Solaris uses answer files to automate interactive package installations, something Linux doesn't understand, while Red Hat offers the concept of source packages. For this reason, this section seeks to highlight differences in using RHN tools on UNIX systems. (Note: RHN does not support Solaris answer files in the current release; such support is planned for future releases.)

Despite inherent differences, such as the lack of Errata, the channel and package management interfaces within the RHN website on the Satellite work largely the same for UNIX systems. All software channels designed to serve UNIX variants can be constructed almost exactly as the custom channels described in the *RHN Channel Management Guide*. The most significant difference is the architecture. When creating a UNIX software channel, ensure you select the base channel architecture appropriate for the systems to be served.

Furthermore, Red Hat recommends you break down your packages into base and child channels depending on their nature. For example, on Solaris, installation packages should go in the Solaris base channel, while patches and Patch Clusters should go in a child channel of the Solaris base channel. Extra installation packages can go in a separate Extras child channel.

RHN treats patches similarly to packages; they are listed and installed in the same way and with the same interface as normal packages. Patches are 'numbered' by Solaris, and will have names like "patch-solaris-108434". The version of a Solaris patch is extracted from the original Solaris metadata, and the release is always 1.

Patch Clusters are bundles of patches that are installed as a unit. RHN keeps track of the last time that a Patch Cluster was installed successfully on a system. However, Patch Clusters are not tracked on the client as installed entities so they do not appear in the installed packages or patches list. Patch Cluster names look like "patch-cluster-solaris-7_Recommended". The version is a datestring, such as "20040206", the release is always 1 and the epoch is always 0.

12.4.2.1. Uploading Packages to the Satellite

RHN does not provide UNIX content; any Solaris packages, patches or Patch Clusters must be uploaded to the Satellite in a format that it understands from a client system. That package can then be managed and distributed to other systems. RHN created `solaris2rpm` to translate Solaris packages, patches, and patch clusters to a format that the Satellite can understand.

12.4.2.1.1. solaris2mpm

As mentioned briefly in [Section 12.1.4, "Differences in Functionality"](#), **solaris2mpm** is part of RHN Push for Solaris. The content that is pushed to a Solaris channel on the Satellite must first be in .mpm format.

A .mpm file is an archive containing a description of the package data and the package or patch itself. The solaris2mpm command must be run on the client, never the Satellite.



Note

solaris2mpm requires free space equal to three times the size of any package, patch, or patch cluster it is converting. Normally, space in `/tmp/` will be used for this purpose. However, the `--tmpdir` option allows you to specify another directory if necessary.

Multiple files may be specified on the command line of solaris2mpm. Below is a usage example:

```
# solaris2mpm RHATrpush-3.1.5-21.pkg RHATrpush-3.1.5-23.pkg
Opening archive, this may take a while
Writing out RHATrpush-3.1.5-21.sparc-solaris.mpm
Opening archive, this may take a while
Writing out RHATrpush-3.1.5-23.sparc-solaris.mpm
```

Because no other directory was specified, the resulting .mpm files are written to the `/tmp/` directory. Note that the name of the resulting .mpm files includes the architecture of the client on which it was created. In this case, this was Sparc Solaris. The general format of mpm file names is:

```
name-version-release.arch.mpm
```

Patch clusters are "exploded" — .mpm files are generated for each patch in the cluster, as well as a top-level "meta" .mpm file containing information about the cluster as a whole.

Below are the options of solaris2mpm:

Option	Description
<code>--version</code>	Displays the program's version number and exits
<code>-h, --help</code>	Displays this information and exits
<code>-?, --usage</code>	Prints program usage information and exits
<code>--tmpdir=<tmpdir></code>	Temporary directory to work from
<code>--select-arch=<arch></code>	Selects the architecture (i386 or Sparc) for multi-arch packages.

Table 12.2. solaris2mpm options

12.4.2.1.2. rhnpush with .mpm Files

The Solaris version of **rhnpush** works like the standard utility, but with the added ability to handle .mpm files. Below is a usage example:

```
% rhnpush -v --server testbox.example.com --username myuser -c solaris-8 \
RHATrpush-3.1.5-*.mpm
Red Hat Network password:
```

```
Connecting to http://testbox.example.com/APP
Uploading package RHATrpush-3.1.5-21.sparc-solaris.mpm
Uploading package RHATrpush-3.1.5-23.sparc-solaris.mpm
```



Note

Patch cluster .mpm files must be pushed either concurrently with or after — never before — the .mpm files for the patches contained in that cluster.

Use `solaris2mpm` on each of the packages, patches, or patch clusters you wish to manage via the Satellite, then use RHN Push to upload them to the channel you created for them.

12.4.2.2. Updating Through the Website

To install packages or patches on an individual system, click the name of the system in the **Systems** category, select the packages from the Upgrade or Install lists of the **Packages** or **Patches** tab, and click **Install/Upgrade Selected Packages**.

To run a remote command while installing the package, click **Run Remote Command** rather than **Confirm**. Refer to [Section 12.5, “Remote Commands”](#) for instructions.

To install packages or patches on multiple systems at once, select the systems and click **System Set Manager** in the left navigation bar. Then, in the **Packages** tab, select the packages from the Upgrade or Install lists and click **Install/Upgrade Packages**. To complete the action, schedule the updates.

12.4.2.3. rhnsd

On Red Hat Enterprise Linux systems, the **rhnsd** daemon, which instructs the client system to check in with RHN, automatically starts at boot time. On Solaris systems, **rhnsd** does not start at boot time by default. It can be started from the command line in this way:

```
rhnsd --foreground --interval=240
```

The default location for **rhnsd** is `/opt/redhat/rhn/solaris/usr/sbin/rhnsd`. Below are the available options for **rhnsd** on Solaris:

Option	Description
-f, --foreground	Run in foreground
-i, --interval=MINS	Connect to Red Hat Network every MINS minutes
-v, --verbose	Log all actions to syslog
-h, --help	Give this help list
-u, --usage	Give this help list
-V, --version	Print program version

Table 12.3. **rhnsd** Options

12.4.2.4. Updating From the Command Line

Like the website, command line use of the **Red Hat Update Agent** is affected by the limitations of UNIX package management. That said, most core functions can still be accomplished through the

up2date command. The most significant difference is the absence of all options regarding source files. Refer to [Table 12.4, “Update Agent Command Line Arguments”](#) for the precise list of options available for UNIX systems.

The command line version of the **Red Hat Update Agent** accepts the following arguments on UNIX systems:

Argument	Description
--version	Show program version information.
-h, --help	Show this help message and exit.
-v, --verbose	Show additional output.
-l, --list	List the latest versions of all packages installed.
-p, --packages	Update packages associated with this System Profile.
--hardware	Update this system's hardware profile on RHN.
--showall	List all packages available for download.
--show-available	List all the packages available that are not currently installed.
--show-orphans	List all the packages currently installed that are not in channels the system is subscribed to.
--show-channels	Show the channel names along with the package names where appropriate.
--installall	Install all available packages. Use with --channel1 .
--channel=CHANNEL	Specify which channels to update from using channel labels.
--get	Fetch the package specified without resolving dependencies.

Table 12.4. Update Agent Command Line Arguments

12.5. Remote Commands

With UNIX support, RHN offers the flexibility of issuing remote commands on client systems through the Satellite's RHN website. This feature allows you to run virtually any (compatible) application or script on any system in your domain without ever having to open a terminal.

12.5.1. Enabling Commands

With the flexibility this tool offers comes great risk and the responsibility to mitigate that risk. For all practical purposes, this feature grants a root BASH prompt to anyone with administrative access to the system on the website.

This can be controlled, however, through the same config-enable mechanism used to determine which systems can have their configuration files managed by Red Hat Network. Refer to [Section 7.4.2.9.3,](#)

[“System Details ▯ Configuration — !\[\]\(d0262bbe9d2356661a2e89321dfcc781_img.jpg\) ”](#) for details.

In short, you must create a directory and file on the UNIX system that tell RHN it is acceptable to run remote commands on the machine. The directory must be named **script**, the file must be named

run, and both must be located in the `/etc/sysconfig/rhn/allowed-actions/` directory specific to your UNIX variant.

For instance, in Solaris, issue this command to create the directory:

```
mkdir -p /opt/redhat/rhn/solaris/etc/sysconfig/rhn/allowed-actions/script
```

To create the requisite file in Solaris, issue this command:

```
touch /opt/redhat/rhn/solaris/etc/sysconfig/rhn/allowed-actions/script/run
```

12.5.2. Issuing Commands

You may schedule a remote command in a variety of ways: on an individual system, on multiple systems at once, and to accompany a package action.

To run a remote command on an individual system by itself, open the **System Details** page and click the **Remote Command** subtab. (Note that this subtab only appears if the system has a Provisioning entitlement.) On this page, establish the settings for the command. You may identify a specific user, group, and timeout period, as well as the script itself. Select a date and time to begin attempting the command, and click the **Schedule Remote Command** link.

Similarly, you may issue a remote command on multiple systems at once through the **System Set Manager**. Select the systems, go to the **System Set Manager**, click the **Misc** tab, and scroll down to the **Remote Command** section. From there you may run a remote command on the selected systems at once.

To run a remote command with a package action, schedule the action through the **Packages** tab of the **System Details** page and click **Run Remote Command** while confirming the action. Use the radio buttons at the top to determine whether the command should run before or after the package action, establish the settings for the command, and click **Schedule Package Install/Upgrade**.

Note that installing multiple packages that have different remote commands requires scheduling the installs separately or combining the commands into a single script.

Appendix A. Red Hat Network Registration Client

Before you begin using Red Hat Network, you must create a username, password, and System Profile. The **Red Hat Network Registration Client** walks you through this process.



Warning

Only systems running Red Hat Enterprise Linux 2.1 need to use the **Red Hat Network Registration Client** before starting the **Red Hat Update Agent**. Systems running Red Hat Enterprise Linux 3 and later have this registration functionality built into the **Red Hat Update Agent**. After registering your system, refer to *Chapter 4, Red Hat Update Agent* for instructions on starting the **Red Hat Update Agent**.

A.1. Configuring the Red Hat Network Registration Client

To start the graphical interface for configuring the application to connect through an HTTP proxy server, type the following command at a shell prompt:

```
rhnc_register --configure
```

The window shown in [Figure A.1, “Red Hat Network Registration Client Configuration”](#) appears.




Figure A.1. Red Hat Network Registration Client Configuration

To start the command line version, use the command:

```
rhn_register --nox --configure
```

It has more configuration options than the graphical version.

You will be presented with a list of options and their current values:

```
0. enableProxyAuth No 1. noSSLServerURL http://xmlrpc.rhn.redhat.com/XMLRPC 2. oemInfoFile /
etc/sysconfig/rhn/oeminfo 3. enableProxy No 4. networkSetup Yes 5. httpProxy 6. proxyUser 7.
serverURL https://xmlrpc.rhn.redhat.com/XMLRPC 8. proxyPassword 9. debug No Enter number of
item to edit <return to exit, q to quit without saving>:
```

Enter the number of the item to modify and enter a new value for the option. When finished changing your configuration, press **Enter** to save your changes and exit. Press **q** and then **Enter** to quit without saving your changes.

The most common options configured are **enableProxy** and **httpProxy** to enable a proxy server. To enable a proxy server, change the value for **enableProxy** to **Yes** and the value of **httpProxy** to the name of the proxy server and port number in the format HOST:PORT. For example, to use the proxy server squid.mysite.org on port 3128, you would change the value to **squid.mysite.org:3128**.

If you require a proxy username and password, set **enableProxyAuth** to **Yes** to enable username/password authentication for the proxy, and set **proxyUser** and **proxyPassword** to the appropriate username and password for the proxy.

To bypass SSL, change the protocol for **serverURL** from **https** to **http** in the `/etc/sysconfig/rhn/rhn_register` file.

A.2. Starting the Red Hat Network Registration Client

You must be root to register a system with RHN. If started by a standard users, the **Red Hat Network Registration Client** prompts you to enter the root password before proceeding.



Important

If your username is part of a larger organizational account, be cautious when registering your systems. By default, all systems registered with the **Red Hat Network Registration Client** end up in the Ungrouped section of systems visible only to Satellite Administrators. To ensure that you retain management of these systems, Red Hat recommends that your organization create an activation key associated with a specific system group and grant you permissions to that group. You may then register your systems using that activation key and find those System Profiles within RHN immediately. Refer to [Section 4.5, “Registering with Activation Keys”](#) for instructions.

To start the **Red Hat Network Registration Client**, use one of the following methods:

1. On the GNOME desktop, go to Applications (the main menu on the panel) => **Programs** => **System** => **Red Hat Network**
2. On the KDE desktop, go to Applications (the main menu on the panel) => **System** => **Red Hat Network**
3. Type the command **rhn_register** at a shell prompt (for example an **XTerm** or **GNOME terminal**)
4. If you are not running the X Window System, type the command **rhn_register** at a shell prompt. Refer to [Section A.7, “Text Mode RHN Registration Client”](#) for further details.



Caution

You must use **Python 1.5.2-24** or later with Secure Sockets Layer (SSL) support. If not, the information transferred is not encrypted. If you have an earlier version of Python, you will see the message shown in [Figure A.2, “Use Python 1.5.2-24 or later”](#). To determine the version of Python on your system, use the command **rpm -q python**. It is strongly recommended that you use **Python 1.5.2-24** or later.



Figure A.2. Use Python 1.5.2-24 or later

If you have already registered your system and try to register it again, the dialog box shown in [Figure A.3, “Warning: This System Already Registered”](#) appears. If you continue, it overwrites your existing Digital Certificate file (`/etc/sysconfig/rhn/systemid`), and creates a different System Profile. You will no longer be able to use your previous System Profile — be sure this is what you want to do before you choose **Yes**.

If you overwrite an existing system registration, you can delete the unused profile via the website at <https://rhn.redhat.com>.

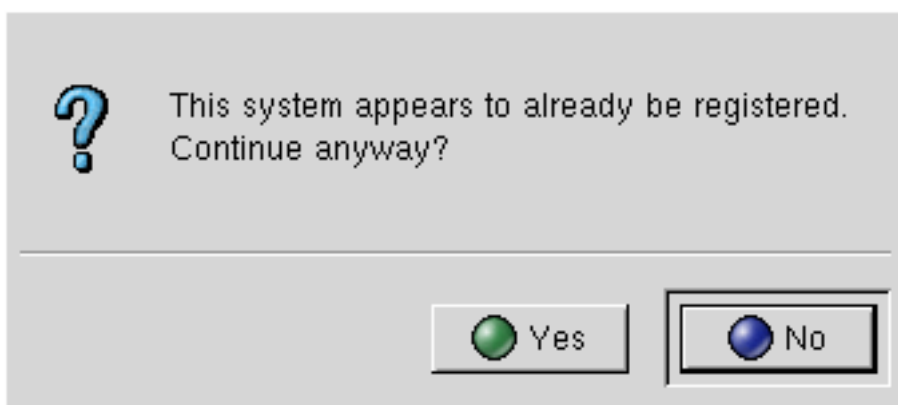


Figure A.3. Warning: This System Already Registered

The opening screen for the **Red Hat Network Registration Client** provides a brief overview of the services available and the steps required to register (see [Figure A.4, “Welcome Screen”](#)). Click **Next** to continue with the registration process. If you click **Cancel**, the registration process ends and no information is sent.

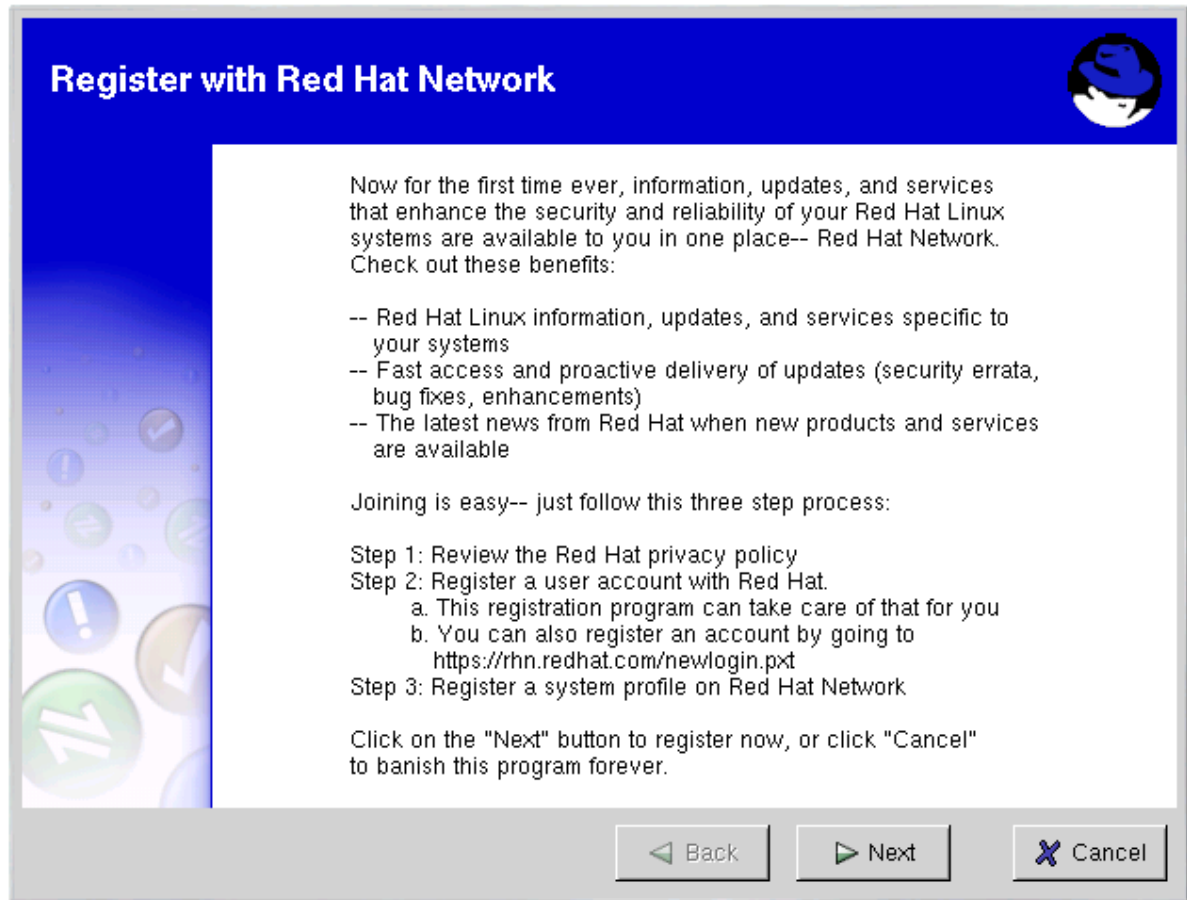


Figure A.4. Welcome Screen

Red Hat is committed to protecting your privacy (see [Figure A.5, "Red Hat Privacy Statement"](#)). The information gathered during the Red Hat Network registration process is used to create a System Profile. The System Profile is essential if you wish to receive update notifications about your system.



Figure A.5. Red Hat Privacy Statement

A.3. Registering a User Account

Before you can create a System Profile, you must create a user account. The only required information in this section is a unique username, password, and a valid email address.

In the screen shown in [Figure A.7, "Create a Unique Username and Password"](#), you must choose a username and password. Once logged in to Red Hat Network, you can modify your preferences, view your existing System Profile, or obtain the latest Red Hat software packages. You must choose a unique username. If you enter one already in use, you will see an error message (see [Figure A.6, "Error: Username Already Exists"](#)). Try different usernames until you find one that has not been used.




Figure A.6. Error: Username Already Exists



Note

If you are already a member of redhat.com, you can use the same user name and password. However, you must continue with the registration process to create your System Profile.

Your username has the following restrictions:

- Cannot contain any spaces
- Cannot contain the characters & +, %, or '
- Is not case-sensitive, thereby eliminating the possibility of duplicate usernames differing only by capitalization

In addition, the following restrictions apply to both your username and password:

- Must be at least four characters long
- Cannot contain any tabs
- Cannot contain any line feeds

Passwords are case-sensitive for obvious reasons.

If you have already registered a machine and created a System Profile, you can add a new machine to your account. Run the **Red Hat Network Registration Client** on the new machine you wish to add, and enter your existing Red Hat Network username and password. The new machine is added to your existing account, and you can log into Red Hat Network with your username and password to view all your systems simultaneously.

Step 2: Register or Update a User Account

Required Information

Are you already registered with redhat.com?
Yes: Enter your current user name and password below.
No: Choose a new user name and password and enter it below.

User name: myname

Password: *****

Password again, for verification: *****

E-mail address: user@example.com

Org Info

If you want this server to be registered as part of an existing organization, enter the information for that here.

organization ID:

organization password:

Back Next Cancel

Figure A.7. Create a Unique Username and Password

Most users can leave the **Org Info** section blank. If you have an existing organization account, work with your Satellite Administrator to ensure that your system is added to that account. This requires entering your organization's ID and password in the provided text fields. If the values are valid, the system is added to the organization's Red Hat Network account. Your Satellite Administrator can then create your user account through the **Users** category of the RHN website. Refer to [Section 7.9, "Users](#)

—  " for instructions.

Click **Next** to continue.

A.4. Registering a System Profile

Now that you have a user account, you can create a System Profile that consists of hardware and software information about your Red Hat Enterprise Linux system. The software System Profile information is used by Red Hat Network to determine what software update notifications you receive.

A.4.1. Hardware System Profile

After creating a username and password for your Red Hat Network account, the **Red Hat Network Registration Client** probes your system for the following information:

- Red Hat Enterprise Linux version
- Hostname

- IP address
- CPU model
- CPU speed
- Amount of RAM
- PCI devices
- Disk sizes
- Mount points

The next step is choosing a profile name for your system as shown in [Figure A.8, “System Profile - Hardware”](#). The default value is the hostname for the system. You may modify this to be a more descriptive string, such as **Email Server for Support Team**. Optionally, you can enter a computer serial or identification number for the system.

If you do not wish to include information about your hardware or network in your System Profile, deselect **Include information about hardware and network** (see [Figure A.8, “System Profile - Hardware”](#)).

Click **Next** to continue with the registration process.

Step 3: Register a System Profile – Hardware

A Profile Name is a descriptive name that you choose to identify this System Profile on Red Hat Network web pages. Optionally, include a computer serial or identification number.

Profile name: Service ID number:

Hardware information is important to determine what updated software and drivers are relevant to this system. The minimum set of information you can include will contain your system's architecture and Red Hat Linux version.

Include information about hardware and network

Included information

Red Hat Linux version: 7.0	CPU model: Pentium III (Coppermine)
Hostname: falcon.meridian.redhat.com	CPU speed: 730 MHz
IP address: 207.175.43.185	Memory: 256 megabytes

Additional hardware information including PCI devices, disk sizes and mount points will be included in the profile.

You will be able to update your hardware profile or create new hardware profiles when you login to Red Hat Network at <http://www.redhat.com/network>.

Figure A.8. System Profile - Hardware

A.4.2. Software System Profile

The software System Profile consists of a list of RPM packages for which you wish to receive notifications. The **Red Hat Network Registration Client** displays a list of all RPM packages listed in the RPM database on your system and then allows you to customize the list by deselecting packages.

A.4.2.1. Gathering RPM Database Information

Only those packages you choose during this part of the registration are included in your System Profile, and you will only receive notifications about the packages in your System Profile. Thus, if you use an older version of a package and deselect it from the list, it will not be replaced with a newer version. This RPM list can be modified through the Red Hat Network website or by using the **Red Hat Update Agent**. [Figure A.9, "Registration Wizard"](#) shows the progress bar displayed while the **Red Hat Network Registration Client** gathers a list of the RPM packages installed on your system. This operation may take some time depending on your system.




Figure A.9. Registration Wizard

Once the RPM package list is built, the list is displayed as shown in [Figure A.10, "RPM Package Information"](#). Deselecting **Include RPM Packages installed on this system in my System Profile** omits this information from your System Profile.

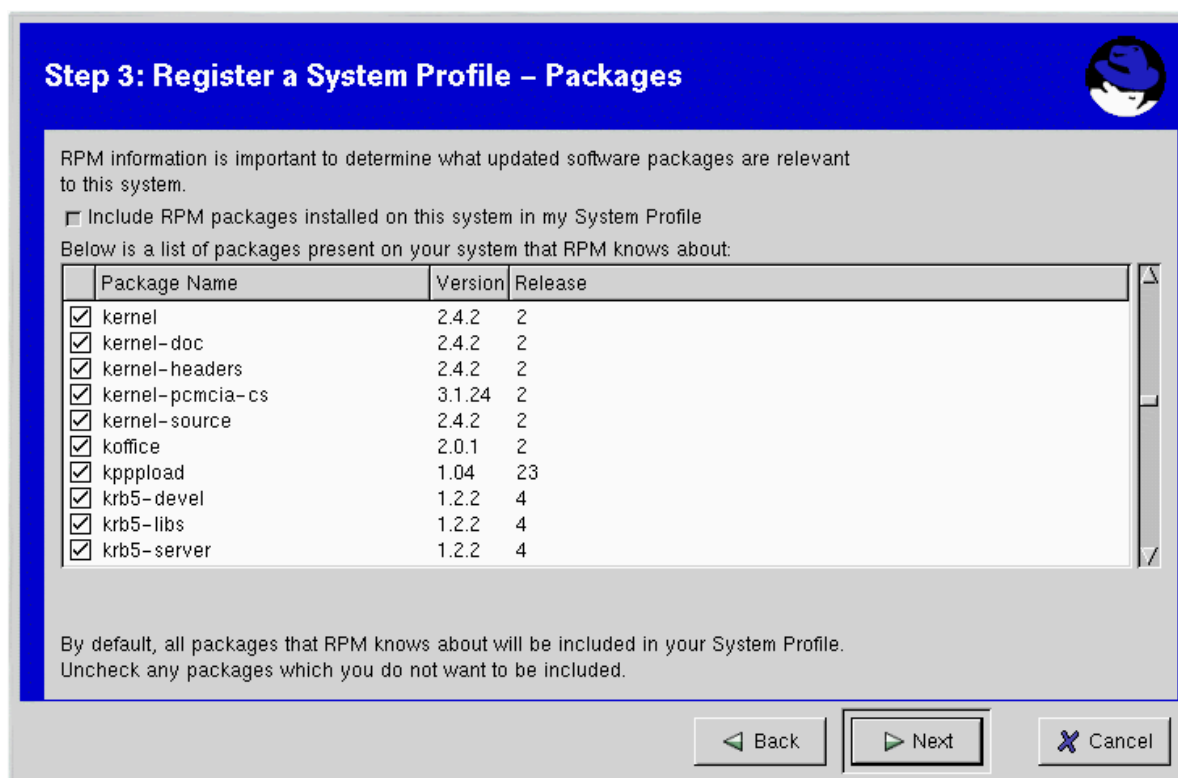


Figure A.10. RPM Package Information

A.4.2.2. Choosing RPM Packages to Exclude from the System Profile

By default, all RPM packages in your RPM database are included in your System Profile to be updated by Red Hat Network. To exclude a package, uncheck the package from the list by clicking the checkbox beside the package name. For example, [Figure A.11, “Choose which RPM Packages to Exclude from System Profile”](#) shows that the **procmail**, **procps**, and **psgml** packages have been omitted from the package list.

Choose which packages to exclude, if any, from the System Profile, and click **Next** to continue with the registration process.

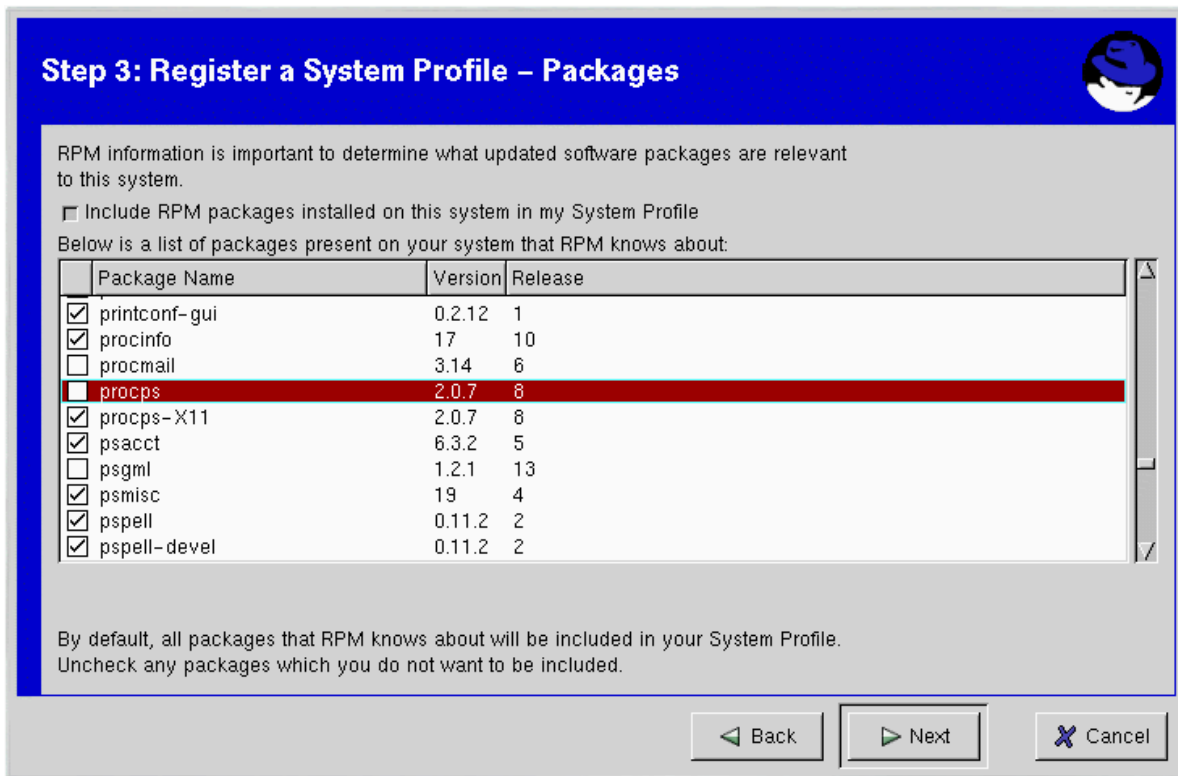


Figure A.11. Choose which RPM Packages to Exclude from System Profile

A.5. Finishing Registration

As seen in [Figure A.12, "Finished Collecting Information for System Profile"](#), the last step of registration is to confirm that you want to send your System Profile to the Red Hat Network. If you choose **Cancel** at this point, no information is sent. Clicking **Next** submits your RHN System Profile.

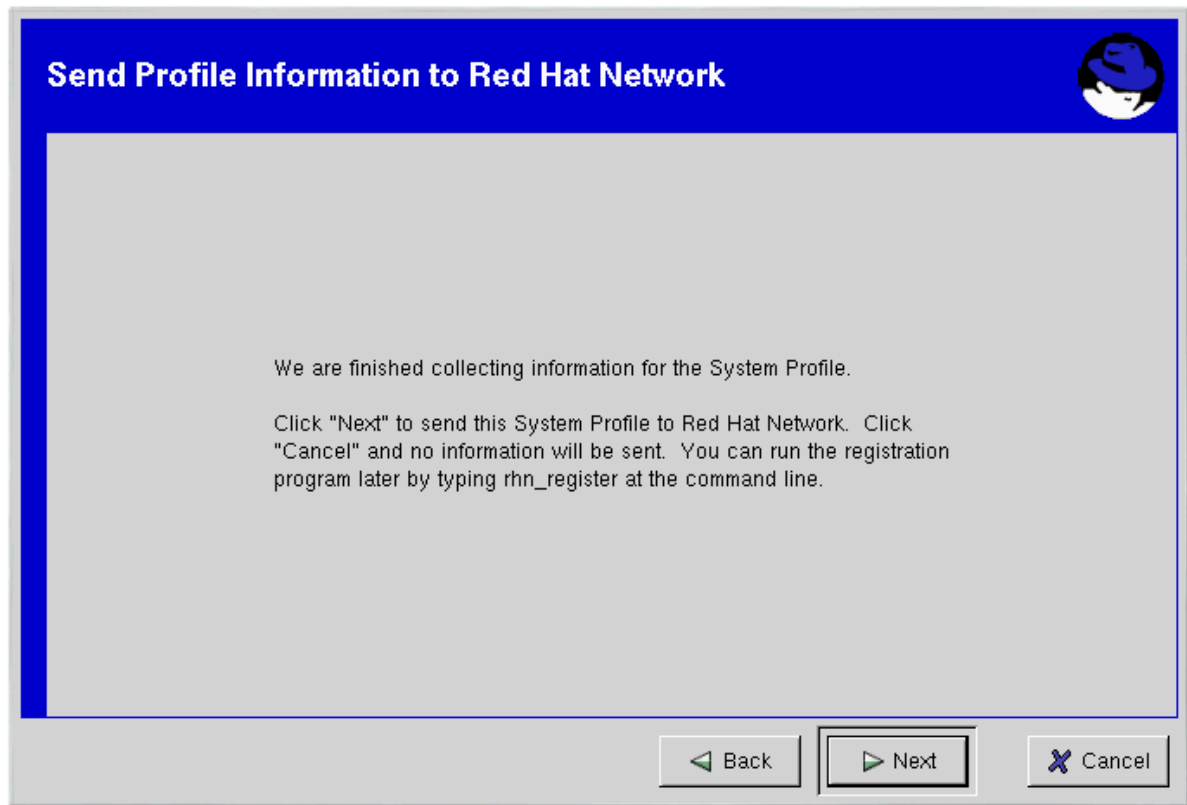


Figure A.12. Finished Collecting Information for System Profile

Figure A.13, "Send System Profile to Red Hat Network" shows the progress bar displayed while your profile is sent. This process may take some time depending on your connection speed.




Figure A.13. Send System Profile to Red Hat Network

The Red Hat Network Registration Client displays the **Registration Finished** screen (Figure A.14, "Registration Finished") once your System Profile has been sent successfully. Click **Finish** to exit the **Red Hat Network Registration Client**.

After completing the registration, you must entitle your system to an RHN service level. Refer to Section A.6, "Entitling Your System" for details.

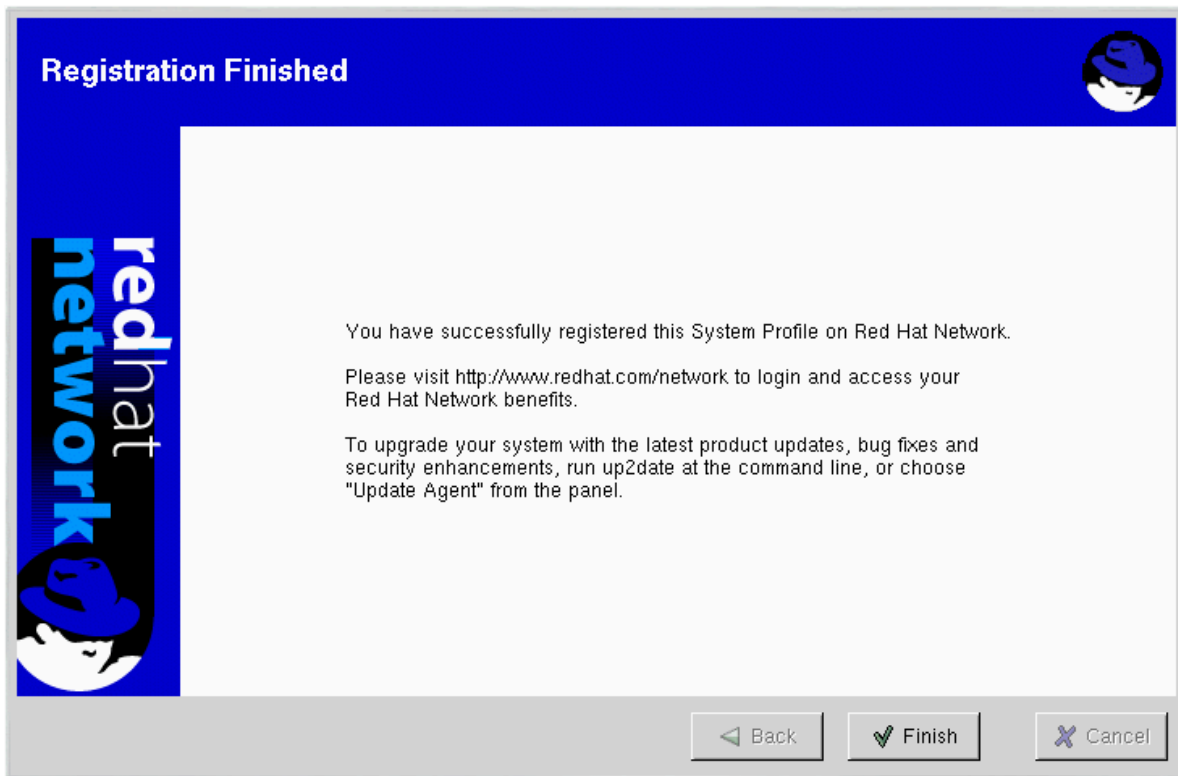


Figure A.14. Registration Finished

A.6. Entitling Your System

Now that you have registered your system, it must be entitled before you can receive updated packages. In other words, you must subscribe it to a service level offering.

To entitle a system, go to <http://rhn.redhat.com> and log in using the same username and password you just used in the **Red Hat Network Registration Client**. Click **Systems** on the top navigation bar and then **Systems Entitlements** in the left navigation bar.

The **System Entitlements** page displays the following items:

- a list of the system for which the user can choose an entitlement level
- the current entitlements applied to each of these systems
- buttons that allow the user to change entitlement level
- an overview of the number and types of purchased entitlements that remain available to the organization

To change the entitlement level of a system or systems, check the box to the left of the systems and click the appropriate button for the desired entitlement level. Note that you must apply a Management entitlement to a system before you can add a Provisioning entitlement. You can change entitlements to any available level at any time.



Note

Removing a required entitlement (such as Provisioning) will not cancel a previously scheduled action (such as a kickstart).

As you change the selected entitlements for your systems, the number of available entitlements is updated at the bottom of the screen.

A.7. Text Mode RHN Registration Client

If you are not running the X Window System, the **Red Hat Network Registration Client** starts in text mode.

You can force the **Red Hat Network Registration Client** to run in text mode with the command:

```
rhn_register --nox
```

The screens for the text mode **Red Hat Network Registration Client** are almost identical to the screens for the graphical **Red Hat Network Registration Client**. Some of the text in the text mode version is more concise due to lack of space in the interface. However, there are equal numbers of screens and fields in both versions. Thus, if you are using the text mode version, you can still follow the instructions that begin in [Section A.2, "Starting the Red Hat Network Registration Client"](#).

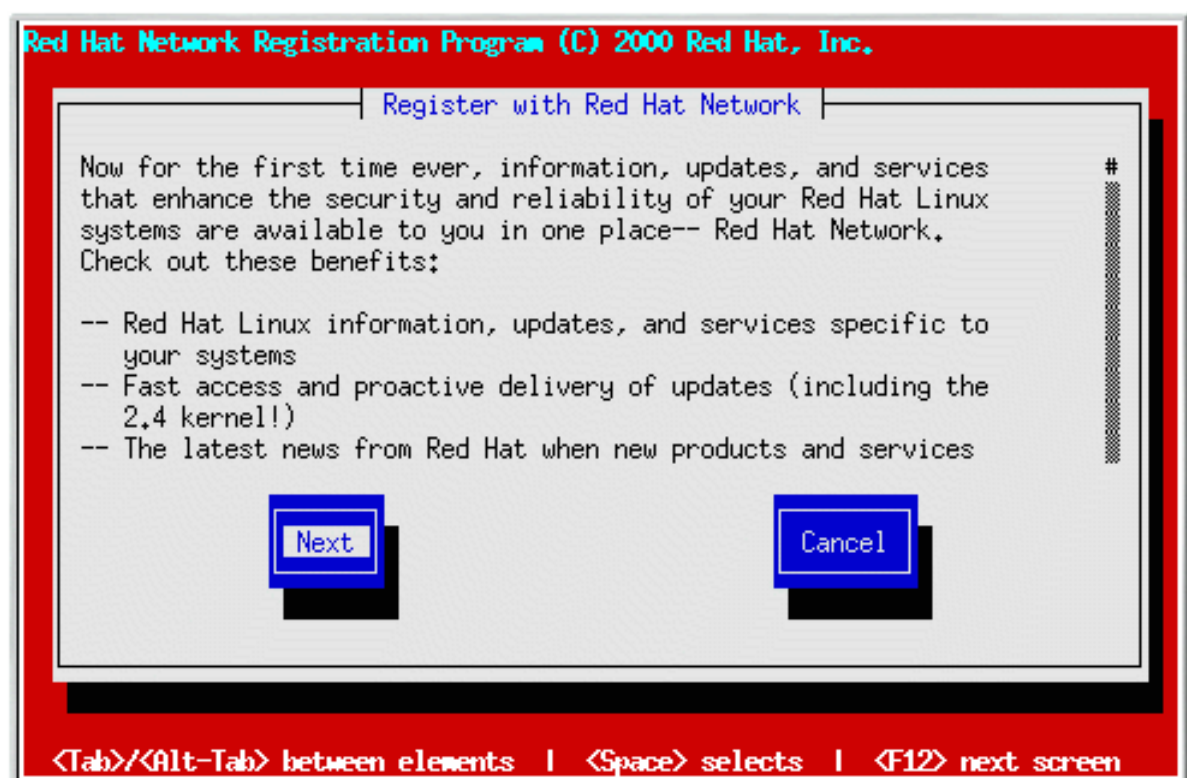


Figure A.15. Text Mode Welcome Screen

Appendix B. Command Line Config Management Tools

In addition to the options provided in the RHN website, Red Hat Network offers two command line tools for managing a system's configuration files: the **Red Hat Network Configuration Client** and the **Red Hat Network Configuration Manager**. There is a complementary **Red Hat Network Actions Control** tool that is used to enable and disable configuration management on client systems. If you do not yet have these tools installed, they can be found within the **RHN Tools** child channel for your operating system.



Tip

Keep in mind, whenever a configuration file is deployed via RHN, a backup of the previous file including its full path is made in the `/var/lib/rhncfg/backups/` directory on the affected system. The backup retains its filename but has a `.rhn-cfg-backup` extension appended.

B.1. Red Hat Network Actions Control

The **Red Hat Network Actions Control** (`rhncfg-actions-control`) application is used to enable and disable configuration management of a system. Client systems cannot be managed in this fashion by default. This tool allows Satellite Administrators to enable or disable specific modes of allowable actions such as: *deploying* a configuration file onto the system, *uploading* a file from the system, *diffing* what is currently managed on a system and what is available, or allowing running arbitrary *remote commands*. These various modes are enabled/disabled by placing/removing files and directories in the `/etc/sysconfig/rhn/allowed-actions/` directory. Due to the default permissions on the `/etc/sysconfig/rhn/` directory, RHN Actions Control will most likely have to be run by someone with root access.

B.1.1. General command line options

There is a `man` page available, as there are for most command line tools, though the use of this tool is simple enough to describe here briefly. Simply decide what RHN scheduled actions should be enabled for use by system administrators. The following options enable the various scheduled action modes:

Option	Description
<code>--enable-deploy</code>	Allow rhncfg-client to deploy files.
<code>--enable-diff</code>	Allow rhncfg-client to diff files.
<code>--enable-upload</code>	Allow rhncfg-client to upload files.
<code>--enable-mtime-upload</code>	Allow rhncfg-client to upload mtime.
<code>--enable-all</code>	Allow rhncfg-client to do everything.
<code>--enable-run</code>	Enable script.run
<code>--disable-deploy</code>	Disable deployment.
<code>--disable-diff</code>	Disable diff
<code>--disable-upload</code>	Disable upload

Option	Description
<code>--disable-mtime-upload</code>	Disable mtime upload
<code>--disable-all</code>	Disable all options
<code>--disable-run</code>	Disable script.run
<code>--report</code>	Report whether the modes are enabled or disabled
<code>-f, --force</code>	Force the operation without asking first
<code>-h, --help</code>	show help message and exit

Table B.1. `rhncfg-control` options

Once a mode is set — and for many, `rhncfg-control --enable-all` is common — your system is now ready for config management through RHN.

B.2. Red Hat Network Configuration Client

As the name implies, the **Red Hat Network Configuration Client** (`rhncfg-client`) is installed and run from an individual client system. From there you may use it to gain knowledge about how RHN deploys configuration files to the client.

The **Red Hat Network Configuration Client** offers these primary modes: `list`, `get`, `channels`, `diff`, and `verify`.

B.2.1. Listing Config Files

To list the configuration files for the machine and the labels of the config channels containing them, issue the command:

```
rhncfg-client list
```

The output resembles the following list:

```
Config Channel File config-channel-17 /etc/example-config.txt config-channel-17 /var/spool/
aalib.rpm config-channel-14 /etc/rhn/rhn.conf
```

These are the configuration files that apply to your system. However, there may be duplicate files present in the other channels. For example, issue the following command:

```
rhncfg-manager list config-channel-14
```

and observe the following output:

```
Files in config channel 'config-channel-14' /etc/example-config.txt /etc/rhn/rhn.conf
```

You may then wonder where the second version of `/etc/example-config.txt` went. The rank of the `/etc/example-config.txt` file in `config-channel-17` was higher than that of the same file in `config-channel-14`. As a result, the version of the configuration file in `config-channel-14` is not deployed for this system, although the file still resides in the channel. The `rhncfg-client` command does not list the file because it will not be deployed on this system.

B.2.2. Getting a Config File

To download the most relevant configuration file for the machine, issue the command:

```
rhncfg-client get /etc/example-config.txt
```

You should see output resembling:

```
Deploying /etc/example-config.txt
```

You may then view the contents of the file with **less** or another pager. Note that the file is selected as the most relevant based upon the rank of the config channel containing it. This is accomplished within the **Configuration** tab of the **System Details** page. Refer to [Section 7.4.2.9, “System Details”](#) for instructions.

B.2.3. Viewing Config Channels

To view the labels and names of the config channels that apply to the system, issue the command:

```
rhncfg-client channels
```

You should see output resembling:

```
Config channels: Label Name ----- ---- config-channel-17 config chan 2 config-channel-14
config chan 1
```

The following table lists the options available for **rhncfg-client get**:

Option	Description
--topdir=TOPDIR	Make all file operations relative to this string.
-h, --help	Show help message and exit

Table B.2. **rhncfg-client get** options

B.2.4. Differentiating between Config Files

To view the differences between the config files deployed on the system and those stored by RHN, issue the command:

```
rhncfg-client diff
```

The output resembles the following:

```
--- /tmp/@3603.0.rhn-cfg-tmp 2004-01-13 14:18:31.000000000 -0500 +++ /etc/example-config.txt
2003-12-16 21:35:32.000000000 -0500 @@ -1,3 +1,5 @@ +additional text
```

In addition, you may include the **--topdir** option to compare config files in RHN with those located in an arbitrary (and unused) location on the client system, like so:

```
[root@ root]# rhncfg-client diff --topdir /home/test/blah/ /usr/bin/diff: /home/test/blah/
etc/example-config.txt: No such file or directory /usr/bin/diff: /home/test/blah/var/spool/
aalib.rpm: No such file or directory
```

B.2.5. Verifying Config Files

To quickly determine if client configuration files are different than those associated with it via RHN, issue the command:

```
rhncfg-client verify
```

The output resembles the following:

```
modified /etc/example-config.txt /var/spool/aalib.rpm
```

The file **example-config.txt** is locally modified, while **aalib.rpm** is not.

The following table lists the options available for **rhncfg-client verify**:

Option	Description
-v, --verbose	Increase the amount of output detail. Displays differences in the mode, owner, and group permissions for the specified config file.
-h, --help	Show help message and exit

Table B.3. **rhncfg-client verify** options

B.3. Red Hat Network Configuration Manager

Unlike the **Red Hat Network Configuration Client**, the **Red Hat Network Configuration Manager** (**rhncfg-manager**) is designed to maintain RHN's central repository of config files and channels, not those located on client systems. This tool offers a command line alternative to the configuration management features within the RHN website, as well as the ability to script some or all of the related maintenance.

It is intended for use by Config Administrators and requires an RHN username and password that has the appropriate permission set. The username may be specified in **/etc/sysconfig/rhn/rhncfg-manager.conf** or in the **[rhncfg-manager]** section of **~/ .rhncfgrc**.

When the **Red Hat Network Configuration Manager** is run as root, it attempts to pull in needed configuration values from the **Red Hat Update Agent**. When run as a user other than root, you may have to make configuration changes within the **~/ .rhncfgrc** file. The session file is cached in **~/ .rhncfg-manager-session** to prevent logging in for every command.

The default timeout for the **Red Hat Network Configuration Manager** is 30 minutes. To alter this, add the **server.session_lifetime** option and new value to the **/etc/rhn/rhn.conf** file on the server running the manager, like so:

```
server.session_lifetime = 120
```

The **Red Hat Network Configuration Manager** offers these primary modes: add, create-channel, diff, diff-revisions, download-channel, get, list, list-channels, remove, remove-channel, revisions, update, and upload-channel.

Each mode offers its own set of options, which can be seen by issuing the following command:

```
rhncfg-manager mode --help
```

Replace *mode* with the name of the mode to be inspected:

```
rhncfg-manager diff-revisions --help
```

You can see such a list of options for the add mode at [Table B.4, “rhncfg-manager add options”](#).

B.3.1. Creating a Config Channel

To create a config channel for your organization, issue the command:

```
rhncfg-manager create-channel channel-label
```

If prompted for your RHN username and password, provide them. The output resembles the following:

```
Red Hat Network username: rhn-user Password: Creating config channel channel-label Config
channel channel-label created
```

Once you have created a config channel, use the remaining modes listed above to populate and maintain that channel.

B.3.2. Adding Files to a Config Channel

To add a file to a config channel, specify the channel label as well as the local file to be uploaded, such as:

```
rhncfg-manager add --channel=channel-label /path/to/file
```

In addition to the required channel label and the path to the file, you may use the available options for modifying the file during its addition. For instance, you may alter the path and file name by including the **--dest-file** option in the command, like:

```
rhncfg-manager add --channel=channel-label --dest-file=/new/path/to/file.txt/path/to/file
```

The output resembles the following:

Appendix B. Command Line Config Management Tools

```
Pushing to channel example-channel Local file >/path/to/file -> remote file /new/path/to/file.txt
```

The following table lists the options available for `rhncfg-manager add`:

Option	Description
<code>-cCHANNEL --channel=CHANNEL</code>	Upload files in this config channel
<code>-dDEST_FILE --dest-file=DEST_FILE</code>	Upload the file as this path
<code>--delim-start=DELIM_START</code>	Start delimiter for variable interpolation
<code>--delim-end=DELIM_END</code>	End delimiter for variable interpolation
<code>-h, --help</code>	show help message and exit

Table B.4. `rhncfg-manager add` options



Note

By default, the maximum file size for configuration files is 128KB. If you need to change that value, find or create the following line in the `/etc/rhn/rhn.conf` file:

```
web.maximum_config_file_size=128
```

Change the value from 128 to whatever limit you want in bytes.

B.3.3. Differentiating between Latest Config Files

To view the differences between the config files on disk and the latest revisions in a channel, issue the command:

```
rhncfg-manager diff --channel=channel-label --dest-file=/path/to/file.txt \ /local/path/to/file
```

You should see output resembling:

```
/tmp/dest_path/example-config.txt /home/test/blah/hello_world.txt --- /tmp/dest_path/example-config.txt config_channel: example-channel revision: 1 +++ /home/test/blah/hello_world.txt 2003-12-14 19:08:59.000000000 -0500 @@ -1 +1 @@ -foo +hello, world
```

The following table lists the options available for `rhncfg-manager diff`:

Option	Description
<code>-cCHANNEL, --channel=CHANNEL</code>	Get file(s) from this config channel
<code>-rREVISION, --revision=REVISION</code>	Use this revision
<code>-dDEST_FILE, --dest-file=DEST_FILE</code>	Upload the file as this path
<code>-tTOPDIR, --topdir=TOPDIR</code>	Make all files relative to this string
<code>-h, --help</code>	Show help message and exit

Table B.5. `rhncfg-manager diff` options

B.3.4. Differentiating between Various Versions

To compare different versions of a file across channels and revisions, use the `-r` flag to indicate which revision of the file should be compared and the `-n` flag to identify the two channels to be checked. Refer to [Section B.3.11, “Determining the Number of File Revisions”](#) for related instructions. Specify only one file name here, since you are comparing the file against another version of itself. For example:

```
rhncfg-manager diff-revisions -n=channel-label1-r=1-n=channel-label2-r=1/path/to/file.txt
```

The output resembles the following:

```
--- /tmp/dest_path/example-config.txt 2004-01-13 14:36:41 \ config channel: example-channel2
revision: 1 --- /tmp/dest_path/example-config.txt 2004-01-13 14:42:42 \ config channel:
example-channel3 revision: 1 @@ -1 +1,20 @@ -foo +blaaaaaaaaaaaaaaaaah +-----BEGIN PGP
SIGNATURE----- +Version: GnuPG v1.0.6 (GNU/Linux) +Comment: For info see http://www.gnupg.org
+ +iD8DBQA9ZY6vse4XmfJPGwgRAsHcAJ9ud9dabUcdscdcqB8AZP7e0Fua0NmKsdhQCe0WHX +VsDTfen2NWdwwPaTM
+S+Cow= +=Ltp2 +-----END PGP SIGNATURE-----
```

The following table lists the options available for `rhncfg-manager diff-revisions`:

Option	Description
<code>-cCHANNEL, --channel=CHANNEL</code>	Use this config channel
<code>-rREVISION, --revision=REVISION</code>	Use this revision
<code>-h, --help</code>	Show help message and exit

Table B.6. `rhncfg-manager diff-revisions` options

B.3.5. Downloading All Files in a Channel

To download all the files in a channel to disk, create a directory and issue the following command:

```
rhncfg-manager download-channel channel-label --topdir .
```

The output resembles the following:

```
Copying /tmp/dest_path/example-config.txt -> \ blah2/tmp/dest_path/example-config.txt
```

The following table lists the options available for `rhncfg-manager download-channel`:

Option	Description
<code>-tTOPDIR, --topdir=TOPDIR</code>	Directory all the file paths are relative to. This option must be set.
<code>-h, --help</code>	Show help message and exit

Table B.7. `rhncfg-manager download-channel` options

B.3.6. Getting the Contents of a File

To direct the contents of a particular file to stdout, issue the command:

```
rhncfg-manager get --channel=channel-label \ /tmp/dest_path/example-config.txt
```

You should see the contents of the file as output.

B.3.7. Listing All Files in a Channel

To list all the files in a channel, issue the command:

```
rhncfg-manager list channel-label
```

You should see output resembling:

```
Files in config channel `example-channel13': /tmp/dest_path/example-config.txt
```

The following table lists the options available for `rhncfg-manager get`:

Option	Description
-cCHANNEL, --channel=CHANNEL	Get file(s) from this config channel
-tTOPDIR, --topdir=TOPDIR	Make all files relative to this string
-rREVISION, --revision=REVISION	Get this file revision
-h, --help	Show help message and exit

Table B.8. `rhncfg-manager get` options

B.3.8. Listing All Config Channels

To list all of your organization's configuration channels, issue the command:

```
rhncfg-manager list-channels
```

The output resembles the following:

```
Available config channels: example-channel example-channel2 example-channel3 config-channel-14 config-channel-17
```

Note that this does not list `local_override` or `server_import` channels.

B.3.9. Removing a File from a Channel

To remove a file from a channel, issue the command:

```
rhncfg-manager remove --channel=channel-label /tmp/dest_path/example-config.txt
```

If prompted for your RHN username and password, provide them. You should see output resembling:


```
Red Hat Network username: rhn-user Password: Removing from config channel example-channel3 /
tmp/dest_path/example-config.txt removed
```

The following table lists the options available for **rhncfg-manager remove**:

Option	Description
-cCHANNEL, --channel=CHANNEL	Remove files from this config channel
-tTOPDIR, --topdir=TOPDIR	Make all files relative to this string
-h, --help	Show help message and exit

Table B.9. **rhncfg-manager remove** options

B.3.10. Deleting a Config Channel

To destroy a config channel in your organization, issue the command:

```
rhncfg-manager remove-channel channel-label
```

The output resembles the following:

```
Removing config channel example-channel Config channel example-channel removed
```

B.3.11. Determining the Number of File Revisions

To find out how many revisions (revisions go from 1 to N where N is an integer greater than 0) of a file/path are in a channel, issue the following command:

```
rhncfg-manager revisions channel-label /tmp/dest_path/example-config.txt
```

The output resembles the following:

```
Analyzing files in config channel example-channel \ /tmp/dest_path/example-config.txt: 1
```

B.3.12. Updating a File in a Channel

To create a new revision of a file in a channel (or add the first revision to that channel if none existed before for the given path), issue the following command:

```
rhncfg-manager update \ --channel=channel-label --dest-file=/path/to/file.txt /local/path/to/
file
```

The output resembles the following:

```
Pushing to channel example-channel: Local file example-channel/tmp/dest_path/example-
config.txt -> \ remote file /tmp/dest_path/example-config.txt
```

The following table lists the options available for **rhncfg-manager update**:

Option	Description
-cCHANNEL, --channel=CHANNEL	Upload files in this config channel
-dDEST_FILE, --dest-file=DEST_FILE	Upload the file as this path
-tTOPDIR, --topdir=TOPDIR	Make all files relative to this string
--delim-start=DELIM_START	Start delimiter for variable interpolation
--delim-end=DELIM_END	End delimiter for variable interpolation
-h, --help	Show help message and exit

Table B.10. **rhncfg-manager update** options

B.3.13. Uploading Multiple Files at Once

To upload multiple files to a config channel from local disk at once, issue the command:

```
rhncfg-manager upload-channel --topdir=topdir channel-label
```

The output resembles the following:

```
Using config channel example-channel14 Uploading /tmp/ola_world.txt from blah4/tmp/ola_world.txt
```

The following table lists the options available for **rhncfg-manager upload-channel**:

Option	Description
-tTOPDIR, --topdir=TOPDIR	Directory all the file paths are relative to
-cCHANNEL, --channel=CHANNEL	List of channels the config info will be uploaded into. Channels delimited by ','. Example: --channel=foo,bar,baz
-h, --help	Show help message and exit

Table B.11. **rhncfg-manager upload-channel** options

Appendix C. RHN API Access

In an effort to provide customers with added flexibility, RHN makes an application programming interface (API) available. This interface can be found by clicking **Help** at the top-right corner of the RHN website, then clicking **API** in the left navigation bar. Or you may go directly to: <https://rhn.redhat.com/rpc/api/>. Use this URL for your XMLRPC server and your browser.

The RHN API is based upon XML-RPC, which allows distinct pieces of software on disparate systems to make remote procedure calls using XML over HTTP. For this reason, any calls you make are expected to meet the constraints of XML-RPC. You can find out more at <http://www.xmlrpc.com/>.

This section bypasses a list of available methods and classes in favor of tips for using the API efficiently. These include steps for determining required values and a sample script that makes some of the calls.

C.1. Using the auth Class and Getting the Session

It is worth noting that you will almost invariably use the auth class first. This class offers a single method, `login`. Use this to establish an RHN session. It requires values for three parameters: username, password, and duration. The first two come directly from your RHN account, while the third is the length of time the session should last in seconds, typically 1200. It returns a session string that can be used in all other methods.

C.2. Obtaining the system_id

Many of the methods require a value for the `system_id` parameter. This is the unique alphanumeric value assigned to each system when registered to RHN. It can be found within the `/etc/sysconfig/rhn/systemid` file on each machine. In addition, you may use the `download_system_id` method within the system class to obtain the value.

C.3. Determining the sid

Several methods require a value for the `sid`, or server ID, parameter. Note that this is different from the `system_id`. You may determine the `sid` of a machine in two different ways. First, you can log into the RHN website, click the name of a system, and view the `sid` at the end of the URL in the location bar. It follows the "=" symbol and is part of a string that resembles the following: "index.pxt?sid=1003486534". Second, you may use the `list_user_systems` method within the system class to obtain a list of systems available to the user that contains the associated `sids`.

C.4. Viewing the cid

Like servers, channels have their own IDs. This value, the `cid`, is a required parameter for some methods, including `set_base_channel` and `set_child_channels`. Also like the `sid`, the `cid` can be obtained through the RHN website. Just click on the name of a channel and view the end of the URL. It follows the "=" symbol, as part of a string that resembles the following: "details.pxt?cid=54".

C.5. Getting the sgid

System groups also have their own IDs. This value, the `sgid`, is a required parameter for the `set_group_membership` method, for instance. Like the `sid` and `cid`, the `sgid` can be obtained through the RHN website. Just click on the name of a system group and view the end of the URL. It

follows the "=" symbol, as part of a string that resembles the following: "details.pxt?sgid=334958". Note that the member parameter within the `set_group_membership` method requires only **yes** or **no** as input to make the association.

C.6. Channel Labels

The architecture of a channel is not always clear from the channel label. Below is a list that shows the correspondence between channel labels and the official title of the architecture they serve.

Channel Label	Platform
channel-i386-sun-solaris	i386 Solaris
channel-ia32	IA-32
channel-ia64	IA-64
channel-sparc	Sparc
channel-alpha	Alpha
channel-s390	IBM S/390
channel-s390x	IBM System z
channel-iSeries	IBM eServer System i
channel-pSeries	IBM eServer System p
channel-x86_64	AMD64 and Intel EM64T
channel-ppc	PPC
channel-sparc-sun-solaris	Sparc Solaris

Table C.1. Channel Labels

This is particularly necessary to know for the `channel.software.create` method.

C.7. Sample API Script

The following sample script depicts how to construct an RHN API client. Review the comments and links for a full discussion of the calls made.

```
#!/usr/bin/perl -w

use strict;
use Frontier::Client;
use Data::Dumper;

#####
# This is a sample script for use of the experimental RHN Management APIs. #
# The API is currently available using XMLRPC only, which is described in #
# depth at:                                                                #
#                                                                           #
# http://www.xmlrpc.com/                                                  #
#                                                                           #
# We use the Frontier modules, available from:                            #
#                                                                           #
# http://theoryx5.uwinnipeg.ca/mod_perl/cpan-search?dist=Frontier-RPC    #
#                                                                           #
#####
```

```
#####  
#   Defining an XMLRPC session.                                     #  
#####  
  
# Define the host first. This will be the FQDN of your satellite system.  
my $HOST = 'satellite.server.yourdomain.com';  
  
# Now we create the client object that will be used throughout the session.  
  
my $client = new Frontier::Client(url => "http://$HOST/rpc/api");  
  
# Next, we execute a login call, which returns a session identifier that will  
# be passed in all subsequent calls. The syntax of this call is described at:  
#  
#   http://$HOST/rpc/api/auth/login/  
  
my $session = $client->call('auth.login', 'username', 'password');  
  
#####  
#   System calls.                                               #  
#####  
  
# This next call returns a list of systems available to the user. The  
# syntax of this call is described at:  
#  
#   http://$HOST/rpc/api/system/list_user_systems/  
#  
# In the code snippet below, we dump data about our systems, and we  
# capture the ID of the first system we find for future operations.  
  
my $systems = $client->call('system.list_user_systems', $session);  
for my $system (@$systems) {  
    print Dumper($system);  
}  
print "\n\nCapturing ID of system @$systems[0]->{name}\n\n";  
my $systemid = @$systems[0]->{id};  
  
# This next call returns a list of packages present on this system. The  
# syntax of this call is described at:  
#  
#   http://$HOST/rpc/api/system/list_packages/  
#  
# This will probably be a pretty long list.  
  
my $packages = $client->call('system.list_packages', $session, $systemid);  
for my $package (@$packages) {  
    print Dumper($package);  
}  
  
# Additional system calls are described at:  
#   http://$HOST/rpc/api/system/
```

Appendix D. Probes

As described in [Section 7.10, “Monitoring — !\[\]\(35e4f762fc1cfea5610d92e2d225d5b4_img.jpg\)”](#), Monitoring-entitled systems can have probes applied to them that constantly confirm their health and full operability. This appendix lists the available probes broken down by command group, such as Apache.

Many probes that monitor internal system aspects (such as the Linux::Disk Usage probe) rather than external aspects (such as the Network Services::SSH probe) require the installation of the Red Hat Network Monitoring Daemon (**rhnmd**). This requirement is noted within the individual probe reference.

Each probe has its own reference in this appendix that identifies required fields (marked with *), default values, and the thresholds that may be set to trigger alerts. Similarly, the beginning of each command group’s section contains information applicable to all probes in that group. [Section D.1, “Probe Guidelines”](#) covers general guidelines; the remaining sections examine individual probes.



Note

Nearly all of the probes use *Transmission Control Protocol* (TCP) as their transport protocol. Exceptions to this are noted within the individual probe references.

D.1. Probe Guidelines

The following general guidelines outline the meaning of each probe state, and provide guidance in setting thresholds for your probes.

The following list provides a brief description of the meaning of each probe state:

Unknown

The probes that cannot collect the metrics needed to determine probe state. Most (though not all) probes enter this state when exceeding their timeout period. Probes in this state may be configured incorrectly, as well.

Pending

The probes whose data has not been received by the RHN Satellite. It is normal for new probes to be in this state. However, if all probes move into this state, your monitoring infrastructure may be failing.

OK

The probes that have run successfully without error. This is the desired state for all probes.

Warning

The probes that have crossed their WARNING thresholds.

Critical

The probes that have crossed their CRITICAL thresholds or reached a critical status by some other means. (Some probes become critical when exceeding their timeout period.)

While adding probes, select meaningful thresholds that, when crossed, notify you and your administrators of problems within your infrastructure. Timeout periods are entered in seconds unless otherwise indicated. Exceptions to these rules are noted within the individual probe references.



Important

Some probes have thresholds based on time. In order for such CRITICAL and WARNING thresholds to work as intended, their values cannot exceed the amount of time allotted to the timeout period. Otherwise, an UNKNOWN status is returned in all instances of extended latency, thereby nullifying the thresholds. For this reason, Red Hat strongly recommends ensuring that timeout periods exceed all timed thresholds.

Remember that Red Hat recommends running your probes without notifications for a time to establish baseline performance for each of your systems. Although the default values provided for probes may suit your needs, every organization has a different environment that may require altering thresholds.

D.2. Apache 1.3.x and 2.0.x

The probes in this section may be applied to instances of the Apache Web server. Although the default values presume you will apply these probes using standard HTTP, you may also use them over secure connections by changing the application protocol to **https** and the port to **443**.

D.2.1. Apache::Processes

The Apache::Processes probe monitors the processes executed on an Apache Web server and collects the following metrics:

- Data Transferred Per Child — Records data transfer information only on individual children. A child process is one that is created from the parent process or another process.
- Data Transferred Per Slot — The cumulative amount of data transferred by a child process that restarts. The number of slots is configured in the **httpd.conf** file using the **MaxRequestsPerChild** setting.

The **ExtendedStatus** directive in the **httpd.conf** file of the Web server must be set to **On** for this probe to function properly.

Field	Value
Application Protocol*	http
Port*	80
Pathname*	/server-status
UserAgent*	NOCpulse-ApacheUptime/1.0
Username	
Password	
Timeout*	15
Critical Maximum Megabytes Transferred Per Child	
Warning Maximum Megabytes Transferred Per Child	
Critical Maximum Megabytes Transferred Per Slot	
Warning Maximum Megabytes Transferred Per Slot	

Table D.1. Apache::Processes settings

D.2.2. Apache::Traffic

The Apache::Traffic probe monitors the requests on an Apache Web server and collects the following metrics:

- Current Requests — The number of requests being processed by the server at probe runtime.
- Request Rate — The accesses to the server per second since the probe last ran.
- Traffic — The kilobytes per second of traffic the server has processed since the probe last ran.

The **ExtendedStatus** directive in the **httpd.conf** file of the Web server must be set to **On** for this probe to function properly.

Field	Value
Application Protocol*	http
Port*	80
Pathname*	/server-status
UserAgent*	NOCpulse-ApacheUptime/1.0
Username	
Password	
Timeout*	15
Critical Maximum Current Requests (number)	
Warning Maximum Current Requests (number)	
Critical Maximum Request Rate (events per second)	
Warning Maximum Request Rate (events per second)	
Critical Maximum Traffic (kilobytes per second)	
Warning Maximum Traffic (kilobytes per second)	

Table D.2. Apache::Traffic settings

D.2.3. Apache::Uptime

The Apache::Uptime probe stores the cumulative time since the Web server was last started. No metrics are collected by this probe, which is designed to help track service level agreements (SLAs).

Field	Value
Application Protocol*	http
Port*	80
Pathname*	/server-status
UserAgent*	NOCpulse-ApacheUptime/1.0
Username	
Password	
Timeout*	15

Table D.3. Apache::Uptime settings

D.3. BEA WebLogic 6.x and higher

The probes in this section (with the exception of JDBC Connection Pool) can be configured to monitor the properties of any BEA WebLogic 6.x and higher server (Administration or Managed) running on a given host, even in a clustered environment. Monitoring of a cluster is achieved by sending all SNMP queries to the Administration Server of the domain and then querying its Managed Servers for individual data.

In order to obtain this higher level of granularity, the **BEA Domain Admin Server** parameter must be used to differentiate between the Administration Server receiving SNMP queries and the Managed Server undergoing the specified probe. If the host to be probed is the Administration Server, then the **BEA Domain Admin Server** parameter can be left blank, and both the SNMP queries and the probe will be sent to it only.

If the host to be probed is a Managed Server, then the IP address of the Administration Server should be provided in the **BEA Domain Admin Server** parameter, and the Managed Server name should be included in the **BEA Server Name** parameter and appended to the end of the **SNMP Community String** field. This causes the SNMP queries to be sent to the Administration Server host, as is required, but redirects the specific probe to the Managed Server host.

It should also be noted that the community string needed for probes run against Managed Server hosts should be in the form of **community_prefix@managed_server_name** in order for the SNMP query to return results for the desired Managed Server. Finally, SNMP must be enabled on each monitored system. SNMP support can be enabled and configured through the WebLogic Console.

Please see the documentation that came with your BEA server or information on the BEA website for more details about BEA's community string naming conventions: <http://e-docs.bea.com/wls/docs70/snmpman/snmpagent.html>

D.3.1. BEA WebLogic::Execute Queue

The BEA WebLogic::Execute Queue probe monitors the WebLogic execute queue and provides the following metrics:

- Idle Execute Threads — The number of execution threads in an idle state.
- Queue Length — The number of requests in the queue.
- Request Rate — The number of requests per second.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP Community String*	public
SNMP Port*	161
SNMP Version*	1
BEA Domain Admin Server	
BEA Server Name*	myserver
Queue Name*	default
Critical Maximum Idle Execute Threads	
Warning Maximum Idle Execute Threads	

Field	Value
Critical Maximum Queue Length	
Warning Maximum Queue Length	
Critical Maximum Request Rate	
Warning Maximum Request Rate	

Table D.4. BEA WebLogic::Execute Queue settings

D.3.2. BEA WebLogic::Heap Free

The BEA WebLogic::Heap Free probe collects the following metric:

- Heap Free — The percentage of free heap space.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP Community String*	public
SNMP Port*	161
SNMP Version*	1
BEA Domain Admin Server	
BEA Server Name*	myserver
Critical Maximum Heap Free	
Warning Maximum Heap Free	
Warning Minimum Heap Free	
Critical Minimum Heap Free	

Table D.5. BEA WebLogic::Heap Free settings

D.3.3. BEA WebLogic::JDBC Connection Pool

The BEA WebLogic::JDBC Connection Pool probe monitors the Java Database Connection (JDBC) pool on a domain Admin Server only (no Managed Servers) and collects the following metrics:

- Connections — The number of connections to the JDBC.
- Connections Rate — The speed at which connections are made to the JDBC, measured in connections per second.
- Waiters — The number of sessions waiting to connect to the JDBC.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP Community String*	public
SNMP Port*	161
SNMP Version*	1
BEA Domain Admin Server	
BEA Server Name*	myserver

Field	Value
JDBC Pool Name*	MyJDBC Connection Pool
Critical Maximum Connections	
Warning Maximum Connections	
Critical Maximum Connection Rate	
Warning Maximum Connection Rate	
Critical Maximum Waiters	
Warning Maximum Waiters	

Table D.6. BEA WebLogic::JDBC Connection Pool settings

D.3.4. BEA WebLogic::Server State

The BEA WebLogic::Server State probe monitors the current state of a BEA Weblogic Web server. If the probe is unable to make a connection to the server, a CRITICAL status results.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP Community String*	public
SNMP Port*	161
SNMP Version*	1
BEA Domain Admin Server	
BEA Server Name*	

Table D.7. BEA WebLogic::Server State settings

D.3.5. BEA WebLogic::Servlet

The BEA WebLogic::Servlet probe monitors the performance of a particular servlet deployed on a WebLogic server and collects the following metrics:

- High Execution Time — The highest amount of time in milliseconds that the servlet takes to execute since the system was started.
- Low Execution Time — The lowest amount of time in milliseconds that the servlet takes to execute since the system was started.
- Execution Time Moving Average — A moving average of the execution time.
- Execution Time Average — A standard average of the execution time.
- Reload Rate — The number of times the specified servlet is reloaded per minute.
- Invocation Rate — The number of times the specified servlet is invoked per minute.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP Community String*	public

Field	Value
SNMP Port*	161
SNMP Version*	1
BEA Domain Admin Server	
BEA Server Name*	myserver
Servlet Name*	
Critical Maximum High Execution Time	
Warning Maximum High Execution Time	
Critical Maximum Execution Time Moving Average	
Warning Maximum Execution Time Moving Average	

Table D.8. BEA WebLogic::Servlet settings

D.4. General

The probes in this section are designed to monitor basic aspects of your systems. When applying them, ensure their timed thresholds do not exceed the amount of time allotted to the timeout period. Otherwise, the probe returns an UNKNOWN status in all instances of extended latency, thereby nullifying the thresholds.

D.4.1. General::Remote Program

The General::Remote Program probe allows you to run any command or script on your system and obtain a status string. Note that the resulting message will be limited to 1024 bytes.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmd**) must be running on the monitored system to execute this probe.

Field	Value
Command*	
OK Exit Status*	0
Warning Exit Status*	1
Critical Exit Status*	2
Timeout	15

Table D.9. General::Remote Program settings

D.4.2. General::Remote Program with Data

The General::Remote Program with Data probe allows you to run any command or script on your system and obtain a value, as well as a status string. To use this probe, you must include XML code in the body of your script. This probe supports the following XML tags:

- `<perldata> </perldata>`
- `<hash> </hash>`
- `<item key = " " > </item>`

The remote program will need to output some iteration of the following code to **STDOUT**:

```
<perldata> <hash> <item
key="data">10</item> <item
key="status_message">status message here</item>
</hash> </perldata>
```

The required value for **data** is the data point to be inserted in the database for time-series trending. The **status_message** is optional and can be whatever text string is desired with a maximum length of 1024 bytes. Remote programs that do not include a **status_message** still report the value and status returned.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**) must be running on the monitored system to execute this probe. XML is case-sensitive. The **data** item key name cannot be changed and it must collect a number as its value.

Field	Value
Command*	
OK Exit Status*	0
Warning Exit Status*	1
Critical Exit Status*	2
Timeout	15

Table D.10. General::Remote Program with Data settings

D.4.3. General::SNMP Check

The General::SNMP Check probe tests your SNMP server by specifying a single object identifier (OID) in dotted notation (such as **1.3.6.1.2.1.1.1.0**) and a threshold associated with the return value. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the SNMP server to answer a connection request.

Requirements — SNMP must be running on the monitored system to perform this probe. Only integers can be used for the threshold values.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP OID*	
SNMP Community String*	public
SNMP Port*	161
SNMP Version*	2
Timeout*	15
Critical Maximum Value	
Warning Maximum Value	
Warning Minimum Value	
Critical Minimum Value	

Table D.11. General::SNMP Check settings

D.4.4. General::TCP Check

The General::TCP Check probe tests your TCP server by verifying that it can connect to a system via the specified port number. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the TCP server to answer a connection request.

The probe passes the string specified in the **Send** field upon making a connection. The probe anticipates a response from the system, which should include the substring specified in the **Expect** field. If the expected string is not found, the probe returns a CRITICAL status.

Field	Value
Send	
Expect	
Port*	1
Timeout*	10
Critical Maximum Latency	
Warning Maximum Latency	

Table D.12. General::TCP Check settings

D.4.5. General::UDP Check

The General::UDP Check probe tests your UDP server by verifying that it can connect to a system via the specified port number. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the UDP server to answer a connection request.

The probe passes the string specified in the **Send** field upon making a connection. The probe anticipates a response from the system, which should include the substring specified in the **Expect** field. If the expected string is not found, the probe returns a CRITICAL status.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
Port*	1
Send	
Expect	
Timeout*	10
Critical Maximum Latency	
Warning Maximum Latency	

Table D.13. General::UDP Check settings

D.4.6. General::Uptime (SNMP)

The General::Uptime (SNMP) probe records the time since the device was last started. It uses the SNMP object identifier (OID) to obtain this value. The only error status it will return is UNKNOWN.

Requirements — SNMP must be running on the monitored system and access to the OID must be enabled to perform this probe.

This probe's transport protocol is User Datagram Protocol (UDP).

Field	Value
SNMP Community String*	public
SNMP Port*	161
SNMP Version*	2
Timeout*	15

Table D.14. General::Uptime (SNMP) settings

D.5. Linux

The probes in this section monitor essential aspects of your Linux systems, from CPU usage to virtual memory. Apply them to mission-critical systems to obtain warnings prior to failure.

Unlike other probe groups, which may or may not require the Red Hat Network Monitoring Daemon, every Linux probe requires that the **rhnm**d daemon be running on the monitored system.

D.5.1. Linux::CPU Usage

The Linux::CPU Usage probe monitors the CPU utilization on a system and collects the following metric:

- CPU Percent Used — The five-second average of the percent of CPU usage at probe execution.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**d) must be running on the monitored system to run this probe.

Field	Value
Timeout*	15
Critical Maximum CPU Percent Used	
Warning Maximum CPU Percent Used	

Table D.15. Linux::CPU Usage settings

D.5.2. Linux::Disk IO Throughput

The Linux::Disk IO Throughput probe monitors a given disk and collects the following metric:

- Read Rate — The amount of data that is read in kilobytes per second.
- Write Rate — The amount of data that is written in kilobytes per second.

To obtain the value for the required **Disk number or disk name** field, run **iostat** on the system to be monitored and see what name has been assigned to the disk you desire. The default value of **0** usually provides statistics from the first hard drive connected directly to the system.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**d) must be running on the monitored system to execute this probe. Also, the **Disk number or disk name** parameter must match the format visible when the **iostat** command is run. If the format is not identical, the configured probe enters an UNKNOWN state.

Field	Value
Disk number or disk name*	0
Timeout*	15
Critical Maximum KB read/second	
Warning Maximum KB read/second	
Warning Minimum KB read/second	
Critical Minimum KB read/second	
Critical Maximum KB written/second	
Warning Maximum KB written/second	
Warning Minimum KB written/second	
Critical Minimum KB written/second	

Table D.16. Linux::Disk IO Throughput settings

D.5.3. Linux::Disk Usage

The Linux::Disk Usage probe monitors the disk space on a specific file system and collects the following metrics:

- File System Used — The percentage of the file system currently in use.
- Space Used — The amount of the file system in megabytes currently in use.
- Space Available — The amount of the file system in megabytes currently available.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmd**) must be running on the monitored system to execute this probe.

Field	Value
File system*	/dev/hda1
Timeout*	15
Critical Maximum File System Percent Used	
Warning Maximum File System Percent Used	
Critical Maximum Space Used	
Warning Maximum Space Used	
Warning Minimum Space Available	
Critical Minimum Space Available	

Table D.17. Linux::Disk Usage settings

D.5.4. Linux::Inodes

The Linux::Inodes probe monitors the specified file system and collects the following metric:

- Inodes — The percentage of inodes currently in use.

An inode is a data structure that holds information about files in a Linux file system. There is an inode for each file, and a file is uniquely identified by the file system on which it resides and its inode number on that system.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
File system*	/
Timeout*	15
Critical Maximum Inodes Percent Used	
Warning Maximum Inodes Percent Used	

Table D.18. Linux::Inodes settings

D.5.5. Linux::Interface Traffic

The Linux::Interface Traffic probe measures the amount of traffic into and out of the specified interface (such as eth0) and collects the following metrics:

- Input Rate — The traffic in bytes per second going into the specified interface.
- Output Rate — The traffic in bytes per second going out of the specified interface.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Interface*	
Timeout*	30
Critical Maximum Input Rate	
Warning Maximum Input Rate	
Warning Minimum Input Rate	
Critical Minimum Input Rate	
Critical Maximum Output Rate	
Warning Maximum Output Rate	
Warning Minimum Output Rate	
Critical Minimum Output Rate	

Table D.19. Linux::Interface Traffic settings

D.5.6. Linux::Load

The Linux::Load probe monitors the CPU of a system and collects the following metric:

- Load — The average load on the system CPU over various periods.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Timeout*	15
Critical CPU Load 1-minute average	

Field	Value
Warning CPU Load 1-minute average	
Critical CPU Load 5-minute average	
Warning CPU Load 5-minute average	
Critical CPU Load 15-minute average	
Warning CPU Load 15-minute average	

Table D.20. Linux::Load settings

D.5.7. Linux::Memory Usage

The Linux::Memory Usage probe monitors the memory on a system and collects the following metric:

- RAM Free — The amount of free random access memory (RAM) in megabytes on a system.

You can also include the reclaimable memory in this metric by entering **yes** or **no** in the **Include reclaimable memory** field.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Include reclaimable memory	no
Timeout*	15
Warning Maximum RAM Free	
Critical Maximum RAM Free	

Table D.21. Linux::Memory Usage settings

D.5.8. Linux::Process Counts by State

The Linux::Process Counts by State probe identifies the number of processes in the following states:

- Blocked — A process that has been switched to the waiting queue and whose state has been switched to **waiting**.
- Defunct — A process that has terminated (either because it has been killed by a signal or because it has called **exit()**) and whose parent process has not yet received notification of its termination by executing some form of the **wait()** system call.
- Stopped — A process that has been stopped before its execution could be completed.
- Sleeping — A process that is in the **Interruptible** sleep state and that can later be reintroduced into memory, resuming execution where it left off.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Timeout*	15
Critical Maximum Blocked Processes	

Field	Value
Warning Maximum Blocked Processes	
Critical Maximum Defunct Processes	
Warning Maximum Defunct Processes	
Critical Maximum Stopped Processes	
Warning Maximum Stopped Processes	
Critical Maximum Sleeping Processes	
Warning Maximum Sleeping Processes	
Critical Maximum Child Processes	
Warning Maximum Child Processes	

Table D.22. Linux::Process Counts by State settings

D.5.9. Linux::Process Count Total

The Linux::Process Count Total probe monitors a system and collects the following metric:

- Process Count — The total number of processes currently running on the system.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**) must be running on the monitored system to execute this probe.

Field	Value
Timeout*	15
Critical Maximum Process Count	
Warning Maximum Process Count	

Table D.23. Linux::Process Count Total settings

D.5.10. Linux::Process Health

The Linux::Process Health probe monitors user-specified processes and collects the following metrics:

- CPU Usage — The CPU usage rate for a given process in milliseconds per second. This metric reports the **time** column of **ps** output, which is the cumulative CPU time used by the process. This makes the metric independent of probe interval, allows sane thresholds to be set, and generates usable graphs (i.e. a sudden spike in CPU usage shows up as a spike in the graph).
- Child Process Groups — The number of child processes spawned from the specified parent process. A child process inherits most of its attributes, such as open files, from its parent.
- Threads — The number of running threads for a given process. A thread is the basic unit of CPU utilization, and consists of a program counter, a register set, and a stack space. A thread is also called a lightweight process.
- Physical Memory Used — The amount of physical memory (or RAM) in kilobytes used by the specified process.
- Virtual Memory Used — The amount of virtual memory in kilobytes used by the specified process, or the size of the process in real memory plus swap.

Specify the process by its command name or process ID. (PID). Entering a PID overrides the entry of a command name. If no command name or PID is entered, the error **Command not found** is displayed and the probe will be set to a CRITICAL state.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Command Name	
Process ID (PID) file	
Timeout*	15
Critical Maximum CPU Usage	
Warning Maximum CPU Usage	
Critical Maximum Child Process Groups	
Warning Maximum Child Process Groups	
Critical Maximum Threads	
Warning Maximum Threads	
Critical Maximum Physical Memory Used	
Warning Maximum Physical Memory Used	
Critical Maximum Virtual Memory Used	
Warning Maximum Virtual Memory Used	

Table D.24. Linux::Process Health settings

D.5.11. Linux::Process Running

The Linux::Process Running probe verifies that the specified process is functioning properly. It counts either processes or process groups, depending on whether the **Count process groups** checkbox is selected.

By default, the checkbox is selected, thereby indicating that the probe should count the number of process group leaders independent of the number of children. This allows you, for example, to verify that two instances of the Apache Web server are running regardless of the (dynamic) number of child processes. If it is not selected, the probe conducts a straightforward count of the number of processes (children and leaders) matching the specified process.

Specify the process by its command name or process ID. (PID). Entering a PID overrides the entry of a command name. If no command name or PID is entered, the error **Command not found** is displayed and the probe enters a CRITICAL state.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Command name	
PID file	
Count process groups	(checked)
Timeout*	15

Field	Value
Critical Maximum Number Running	
Critical Minimum Number Running	

Table D.25. Linux::Process Running settings

D.5.12. Linux::Swap Usage

The Linux::Swap Usage probe monitors the swap partitions running on a system and reports the following metric:

- Swap Free — The percent of swap memory currently free.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**) must be running on the monitored system to execute this probe.

Field	Value
Timeout*	15
Warning Minimum Swap Free	
Critical Minimum Swap Free	

Table D.26. Linux::Swap Usage settings

D.5.13. Linux::TCP Connections by State

The Linux::TCP Connections by State probe identifies the total number of TCP connections, as well as the quantity of each in the following states:

- TIME_WAIT — The socket is waiting after close for remote shutdown transmission so it may handle packets still in the network.
- CLOSE_WAIT — The remote side has been shut down and is now waiting for the socket to close.
- FIN_WAIT — The socket is closed, and the connection is now shutting down.
- ESTABLISHED — The socket has a connection established.
- SYN_RCVD — The connection request has been received from the network.

This probe can be helpful in finding and isolating network traffic to specific IP addresses or examining network connections into the monitored system.

The filter parameters for the probe let you narrow the probe's scope. This probe uses the **netstat -ant** command to retrieve data. The **Local IP address** and **Local port** parameters use values in the **Local Address** column of the output; the **Remote IP address** and **Remote port** parameters use values in the **Foreign Address** column of the output for reporting.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**) must be running on the monitored system to execute this probe.

Field	Value
Local IP address filter pattern list	
Local port number filter	

Field	Value
Remote IP address filter pattern list	
Remote port number filter	
Timeout*	15
Critical Maximum Total Connections	
Warning Maximum Total Connections	
Critical Maximum TIME_WAIT Connections	
Warning Maximum TIME_WAIT Connections	
Critical Maximum CLOSE_WAIT Connections	
Warning Maximum CLOSE_WAIT Connections	
Critical Maximum FIN_WAIT Connections	
Warning Maximum FIN_WAIT Connections	
Critical Maximum ESTABLISHED Connections	
Warning Maximum ESTABLISHED Connections	
Critical Maximum SYN_RCVD Connections	
Warning Maximum SYN_RCVD Connections	

Table D.27. Linux::TCP Connections by State settings

D.5.14. Linux::Users

The Linux::Users probe monitors the users of a system and reports the following metric:

- Users — The number of users currently logged in.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Timeout*	15
Critical Maximum Users	
Warning Maximum Users	

Table D.28. Linux::Users settings

D.5.15. Linux::Virtual Memory

The Linux::Virtual Memory probe monitors the total system memory and collects the following metric:

- Virtual Memory — The percent of total system memory - random access memory (RAM) plus swap - that is free.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
Timeout*	15
Warning Minimum Virtual Memory Free	

Field	Value
Critical Minimum Virtual Memory Free	

Table D.29. Linux::Virtual Memory settings

D.6. LogAgent

The probes in this section monitor the log files on your systems. You can use them to query logs for certain expressions and track the sizes of files. For LogAgent probes to run, the **nocpulse** user must be granted read access to your log files.

Note that data from the first run of these probes is not measured against the thresholds to prevent spurious notifications caused by incomplete metric data. Measurements will begin on the second run.

D.6.1. LogAgent::Log Pattern Match

The LogAgent::Log Pattern Match probe uses regular expressions to match text located within the monitored log file and collects the following metrics:

- Regular Expression Matches — The number of matches that have occurred since the probe last ran.
- Regular Expression Match Rate — The number of matches per minute since the probe last ran.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe. For this probe to run, the **nocpulse** user must be granted read access to your log files.

In addition to the name and location of the log file to be monitored, you must provide a regular expression to be matched against. The expression must be formatted for **egrep**, which is equivalent to **grep -E** and supports extended regular expressions. This is the regular expression set for **egrep**:

```

^ beginning of line
$ end of line
. match one char
* match zero or more chars
[] match one character set, e.g. '[Ff]oo'
[^] match not in set '[^A-F]oo'
+ match one or more of preceding chars
? match zero or one of preceding chars
| or, e.g. a|b
() groups chars, e.g., (foo|bar) or (foo)+
    
```



Warning

Do not include single quotation marks (') within the expression. Doing so causes **egrep** to fail silently and the probe to time out.

Field	Value
Log file*	/var/log/messages
Basic regular expression*	

Field	Value
Timeout*	45
Critical Maximum Matches	
Warning Maximum Matches	
Warning Minimum Matches	
Critical Minimum Matches	
Critical Maximum Match Rate	
Warning Maximum Match Rate	
Warning Minimum Match Rate	
Critical Maximum Match Rate	

Table D.30. LogAgent::Log Pattern Match settings

D.6.2. LogAgent::Log Size

The LogAgent::Log Size probe monitors log file growth and collects the following metrics:

- Size — The size the log file has grown in bytes since the probe last ran.
- Output Rate — The number of bytes per minute the log file has grown since the probe last ran.
- Lines — The number of lines written to the log file since the probe last ran.
- Line Rate — The number of lines written per minute to the log file since the probe last ran.

Requirements — The Red Hat Network Monitoring Daemon (**rhnmd**) must be running on the monitored system to execute this probe. For this probe to run, the **nocpulse** user must be granted read access to your log files.

Field	Value
Log file*	/var/log/messages
Timeout*	20
Critical Maximum Size	
Warning Maximum Size	
Warning Minimum Size	
Critical Minimum Size	
Critical Maximum Output Rate	
Warning Maximum Output Rate	
Warning Minimum Output Rate	
Critical Minimum Output Rate	
Critical Maximum Lines	
Warning Maximum Lines	
Warning Minimum Lines	
Critical Minimum Lines	
Critical Maximum Line Rate	

Field	Value
Warning Maximum Line Rate	
Warning Minimum Line Rate	
Critical Minimum Line Rate	

Table D.31. LogAgent::Log Size settings

D.7. MySQL 3.23 - 3.33

The probes in this section monitor aspects of the MySQL database using the `mysqladmin` binary. No specific user privileges are needed for these probes.

Note that the `mysql-server` package must be installed on the system conducting the monitoring for these probes to complete. Refer to the MySQL Installation section of the *RHN Satellite Installation Guide* for instructions.

D.7.1. MySQL::Database Accessibility

The MySQL::Database Accessibility probe tests connectivity through a database account that has no database privileges. If no connection is made, a CRITICAL status results.

Field	Value
Username*	
Password	
MySQL Port	3306
Database*	mysql
Timeout	15

Table D.32. MySQL::Database Accessibility settings

D.7.2. MySQL::Opened Tables

The MySQL::Opened Tables probe monitors the MySQL server and collects the following metric:

- Opened Tables — The tables that have been opened since the server was started.

Field	Value
Username	
Password	
MySQL Port*	3306
Timeout	15
Critical Maximum Opened Objects	
Warning Maximum Opened Objects	
Warning Minimum Opened Objects	
Critical Minimum Opened Objects	

Table D.33. MySQL::Opened Tables settings

D.7.3. MySQL::Open Tables

The MySQL::Open Tables probe monitors the MySQL server and collects the following metric:

- Open Tables — The number of tables open when the probe runs.

Field	Value
Username	
Password	
MySQL Port*	3306
Timeout	15
Critical Maximum Open Objects	
Warning Maximum Open Objects	
Warning Minimum Open Objects	
Critical Minimum Open Objects	

Table D.34. MySQL::Open Tables settings

D.7.4. MySQL::Query Rate

The MySQL::Query Rate probe monitors the MySQL server and collects the following metric:

- Query Rate — The average number of queries per second per database server.

Field	Value
Username	
Password	
MySQL Port*	3306
Timeout	15
Critical Maximum Query Rate	
Warning Maximum Query Rate	
Warning Minimum Query Rate	
Critical Minimum Query Rate	

Table D.35. MySQL::Query Rate settings

D.7.5. MySQL::Threads Running

The MySQL::Threads Running probe monitors the MySQL server and collects the following metric:

- Threads Running — The total number of running threads within the database.

Field	Value
Username	
Password	
MySQL Port*	3306
Timeout	15
Critical Maximum Threads Running	

Field	Value
Warning Maximum Threads Running	
Warning Minimum Threads Running	
Critical Minimum Threads Running	

Table D.36. MySQL::Threads Running settings

D.8. Network Services

The probes in this section monitor various services integral to a functioning network. When applying them, ensure that their timed thresholds do not exceed the amount of time allotted to the timeout period. Otherwise, an UNKNOWN status is returned in all instances of extended latency, thereby nullifying the thresholds.

D.8.1. Network Services::DNS Lookup

The Network Services::DNS Lookup probe uses the **dig** command to see if it can resolve the system or domain name specified in the **Host or Address to look up** field. It collects the following metric:

- Query Time — The time in milliseconds required to execute the **dig** request.

This is useful in monitoring the status of your DNS servers. To monitor one of your DNS servers, supply a well-known host/domain name, such as a large search engine or corporate Web site.

Field	Value
Host or Address to look up	
Timeout*	10
Critical Maximum Query Time	
Warning Maximum Query Time	

Table D.37. Network Services::DNS Lookup settings

D.8.2. Network Services::FTP

The Network Services::FTP probe uses network sockets to test FTP port availability. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the FTP server to answer a connection request.

This probe supports authentication. Provide a username and password in the appropriate fields to use this feature. The optional **Expect** value is the string to be matched against after a successful connection is made to the FTP server. If the expected string is not found, the probe returns a CRITICAL state.

Field	Value
Expect	FTP
Username	
Password	
FTP Port*	21

Field	Value
Timeout*	10
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.38. Network Services::FTP settings

D.8.3. Network Services::IMAP Mail

The Network Services::IMAP Mail probe determines if it can connect to the IMAP 4 service on the system. Specifying an optional port will override the default port 143. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the IMAP server to answer a connection request.

The required **Expect** value is the string to be matched against after a successful connection is made to the IMAP server. If the expected string is not found, the probe returns a CRITICAL state.

Field	Value
IMAP Port*	143
Expect*	OK
Timeout*	5
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.39. Network Services::IMAP Mail settings

D.8.4. Network Services::Mail Transfer (SMTP)

The Network Services::Mail Transfer (SMTP) probe determines if it can connect to the SMTP port on the system. Specifying an optional port number overrides the default port 25. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the SMTP server to answer a connection request.

Field	Value
SMTP Port*	25
Timeout*	10
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.40. Network Services::Mail Transfer (SMTP) settings

D.8.5. Network Services::Ping

The Network Services::Ping probe determines if the RHN Server can **ping** the monitored system or a specified IP address. It also checks the packet loss and compares the round trip average against the Warning and Critical threshold levels. The required **Packets to send** value allows you to control how many ICMP ECHO packets are sent to the system. This probe collects the following metrics:

Appendix D. Probes

- Round-Trip Average — The time it takes in milliseconds for the ICMP ECHO packet to travel to and from the monitored system.
- Packet Loss — The percent of data lost in transit.

Although optional, the **IP Address** field can be instrumental in collecting metrics for systems that have multiple IP addresses. For instance, if the system is configured with multiple virtual IP addresses or uses Network Address Translation (NAT) to support internal and external IP addresses, this option may be used to check a secondary IP address rather than the primary address associated with the hostname.

Note that this probe conducts the **ping** from an RHN Server and not the monitored system. Populating the IP Address field does not test connectivity between the system and the specified IP address but between the RHN Server and the IP address. Therefore, entering the same IP address for Ping probes on different systems accomplishes precisely the same task. To conduct a **ping** from a monitored system to an individual IP address, use the Remote Ping probe instead. Refer to [Section D.8.7, “Network Services::Remote Ping”](#).

Field	Value
IP Address (defaults to system IP)	
Packets to send*	20
Timeout*	10
Critical Maximum Round-Trip Average	
Warning Maximum Round-Trip Average	
Critical Maximum Packet Loss	
Warning Maximum Packet Loss	

Table D.41. Network Services::Ping settings

D.8.6. Network Services::POP Mail

The Network Services::POP Mail probe determines if it can connect to the POP3 port on the system. A port number must be specified; specifying another port number overrides the default port 110. This probe collects the following metric:

- Remote Service Latency — The time it takes in seconds for the POP server to answer a connection request.

The required **Expect** value is the string to be matched against after a successful connection is made to the POP server. The probe looks for the string in the first line of the response from the system. The default is **+OK**. If the expected string is not found, the probe returns a CRITICAL state.

Field	Value
Port*	110
Expect*	+OK
Timeout*	10
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.42. Network Services::POP Mail settings

D.8.7. Network Services::Remote Ping

The Network Services::Remote Ping probe determines if the monitored system can **ping** a specified IP address. It also monitors the packet loss and compares the round trip average against the Warning and Critical threshold levels. The required **Packets to send** value allows you to control how many ICMP ECHO packets are sent to the address. This probe collects the following metrics:

- Round-Trip Average — The time it takes in milliseconds for the ICMP ECHO packet to travel to and from the IP address.
- Packet Loss — The percent of data lost in transit.

The **IP Address** field identifies the precise address to be pinged. Unlike the similar, optional field in the standard Ping probe, this field is required. The monitored system directs the ping to a third address, rather than to the RHN Server. Since the Remote Ping probe tests connectivity from the monitored system, another IP address must be specified. To conduct pings from the RHN Server to a system or IP address, use the standard Ping probe instead. Refer to [Section D.8.5, “Network Services::Ping”](#).

Requirements — The Red Hat Network Monitoring Daemon (**rhnmmd**) must be running on the monitored system to execute this probe.

Field	Value
IP Address*	
Packets to send*	20
Timeout*	10
Critical Maximum Round-Trip Average	
Warning Maximum Round-Trip Average	
Critical Maximum Packet Loss	
Warning Maximum Packet Loss	

Table D.43. Network Services::Remote Ping settings

D.8.8. Network Services::RPCService

The Network Services::RPCService probe tests the availability of remote procedure call (RPC) programs on a given IP address. It collects the following metric:

- Remote Service Latency — The time it takes in seconds for the RPC server to answer a connection request.

RPC server programs, which provide function calls via that RPC network, register themselves in the RPC network by declaring a program ID and a program name. NFS is an example of a service that works via the RPC mechanism.

Client programs that wish to use the resources of RPC server programs do so by asking the machine on which the server program resides to provide access to RPC functions within the RPC program number or program name. These conversations can occur over either TCP or UDP (but are almost always UDP).

This probe allows you to test simple program availability. You must specify the program name or number, the protocol over which the conversation occurs, and the usual timeout period.

Field	Value
Protocol (TCP/UDP)	udp
Service Name*	nfs
Timeout*	10
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.44. Network Services::RPCService settings

D.8.9. Network Services::Secure Web Server (HTTPS)

The Network Services::Secure Web Server (HTTPS) probe determines the availability of the secure Web server and collects the following metric:

- Remote Service Latency — The time it takes in seconds for the HTTPS server to answer a connection request.

This probe confirms that it can connect to the HTTPS port on the specified host and retrieve the specified URL. If no URL is specified, the probe fetches the root document. The probe looks for a HTTP/1. message from the system unless you alter that value. Specifying another port number overrides the default port of 443.

This probe supports authentication. Provide a username and password in the appropriate fields to use this feature. Unlike most other probes, this probe returns a CRITICAL status if it cannot contact the system within the timeout period.

Field	Value
URL Path	/
Expect Header	HTTP/1
Expect Content	
UserAgent*	NOCpulse-check_http/1.0
Username	
Password	
Timeout*	10
HTTPS Port*	443
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.45. Network Services::Secure Web Server (HTTPS) settings

D.8.10. Network Services::SSH

The Network Services::SSH probe determines the availability of SSH on the specified port and collects the following metric:

- Remote Service Latency — The time it takes in seconds for the SSH server to answer a connection request.

Upon successfully contacting the SSH server and receiving a valid response, the probe displays the protocol and server version information. If the probe receives an invalid response, it displays the message returned from the server and generates a WARNING state.

Field	Value
SSH Port*	22
Timeout*	5
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.46. Network Services::SSH settings

D.8.11. Network Services::Web Server (HTTP)

The Network Services::Web Server (HTTP) probe determines the availability of the Web server and collects the following metric:

- Remote Service Latency — The time it takes in seconds for the HTTP server to answer a connection request.

This probe confirms it can connect to the HTTP port on the specified host and retrieve the specified URL. If no URL is specified, the probe will fetch the root document. The probe looks for a HTTP/1. message from the system, unless you alter that value. Specifying another port number will override the default port of 80. Unlike most other probes, this probe will return a CRITICAL status if it cannot contact the system within the timeout period.

This probe supports authentication. Provide a username and password in the appropriate fields to use this feature. Also, the optional Virtual Host field can be used to monitor a separate documentation set located on the same physical machine presented as a standalone server. If your Web server is not configured to use virtual hosts (which is typically the case), you should leave this field blank. If you do have virtual hosts configured, enter the domain name of the first host here. Add as many probes as necessary to monitor all virtual hosts on the machine.

Field	Value
URL Path	/
Virtual Host	
Expect Header	HTTP/1
Expect Content	
UserAgent*	NOCpulse-check_http/1.0
Username	
Password	
Timeout*	10
HTTP Port*	80
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.47. Network Services::Web Server (HTTP) settings

D.9. Oracle 8i, 9i, and 10g

The probes in this section may be applied to instances of the Oracle database matching the versions supported. Oracle probes require the configuration of the database and associations made by running the following command:

```
$ORACLE_HOME/rdbms/admin/catalog.sql
```

In addition, for these probes to function properly, the Oracle user configured in the probe must have minimum privileges of CONNECT and SELECT_CATALOG_ROLE.

Some Oracle probes are specifically aimed at tuning devices for long-term performance gains, rather than avoiding outages. Therefore, Red Hat recommends scheduling them to occur less frequently, between every hour and every two days. This provides a better statistical representation, de-emphasizing anomalies that can occur at shorter time intervals. This applies to following probes: Buffer Cache, Data Dictionary Cache, Disk Sort Ratio, Library Cache, and Redo Log.

For CRITICAL and WARNING thresholds based upon time to work as intended, their values cannot exceed the amount of time allotted to the timeout period. Otherwise, an UNKNOWN status is returned in all cases of extended latency, thereby nullifying the thresholds. For this reason, Red Hat strongly recommends ensuring that timeout periods exceed all timed thresholds. In this section, this refers specifically to the probe TNS Ping.

Finally, customers using these Oracle probes against a database using Oracle's Multi-Threaded Server (MTS) must contact Red Hat support to have entries added to the RHN Server's /etc/hosts file to ensure that the DNS name is resolved correctly.

D.9.1. Oracle::Active Sessions

The Oracle::Active Sessions probe monitors an Oracle instance and collects the following metrics:

- Active Sessions — The number of active sessions based on the value of **V \$PARAMETER . PROCESSES**.
- Available Sessions — The percentage of active sessions that are available based on the value of **V \$PARAMETER . PROCESSES**.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Timeout*	30
Critical Maximum Active Sessions	
Warning Maximum Active Sessions	
Critical Maximum Available Sessions Used	
Warning Maximum Available Sessions Used	

Table D.48. Oracle::Active Sessions settings

D.9.2. Oracle::Availability

The Oracle::Availability probe determines the availability of the database from the RHN Satellite.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Timeout*	30

Table D.49. Oracle::Availability settings

D.9.3. Oracle::Blocking Sessions

The Oracle::Blocking Sessions probe monitors an Oracle instance and collects the following metric:

- **Blocking Sessions** — The number of sessions preventing other sessions from committing changes to the Oracle database, as determined by the required *Time Blocking* value you provide. Only those sessions that have been blocking for this duration, which is measured in seconds, are counted as blocking sessions.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Time Blocking (seconds)*	20
Timeout*	30
Critical Maximum Blocking Sessions	
Warning Maximum Blocking Sessions	

Table D.50. Oracle::Blocking Sessions settings

D.9.4. Oracle::Buffer Cache

The Oracle::Buffer Cache probe computes the Buffer Cache Hit Ratio so as to optimize the system global area (SGA) Database Buffer Cache size. It collects the following metrics:

- **Db Block Gets** — The number of blocks accessed via single block gets (not through the consistent get mechanism).
- **Consistent Gets** — The number of accesses made to the block buffer to retrieve data in a consistent mode.
- **Physical Reads** — The cumulative number of blocks read from disk.
- **Buffer Cache Hit Ratio** — The rate at which the database goes to the buffer instead of the hard disk to retrieve data. A low ratio suggests more RAM should be added to the system.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port	1521
Timeout*	30
Warning Minimum Buffer Cache Hit Ratio	
Critical Minimum Buffer Cache Hit Ratio	

Table D.51. Oracle::Buffer Cache settings

D.9.5. Oracle::Client Connectivity

The Oracle::Client Connectivity probe determines if the database is up and capable of receiving connections from the monitored system. This probe opens an **rhnm** connection to the system and issues a **sqlplus connect** command on the monitored system.

The **Expected DB name** parameter is the expected value of **V\$DATABASE.NAME**. This value is case-insensitive. A CRITICAL status is returned if this value is not found.

Requirements — The Red Hat Network Monitoring Daemon (**rhnm**) must be running on the monitored system to execute this probe. For this probe to run, the **nocpulse** user must be granted read access to your log files.

Field	Value
Oracle Hostname or IP address*	
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
ORACLE_HOME*	/opt/oracle
Expected DB Name*	
Timeout*	30

Table D.52. Oracle::Client Connectivity settings

D.9.6. Oracle::Data Dictionary Cache

The Oracle::Data Dictionary Cache probe computes the Data Dictionary Cache Hit Ratio so as to optimize the **SHARED_POOL_SIZE** in **init.ora**. It collects the following metrics:

- **Data Dictionary Hit Ratio** — The ratio of cache hits to cache lookup attempts in the data dictionary cache. In other words, the rate at which the database goes to the dictionary instead of the hard disk to retrieve data. A low ratio suggests more RAM should be added to the system.
- **Gets** — The number of blocks accessed via single block gets (not through the consistent get mechanism).
- **Cache Misses** — The number of accesses made to the block buffer to retrieve data in a consistent mode.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Timeout*	30
Warning Minimum Data Dictionary Hit Ratio	
Critical Minimum Data Dictionary Hit Ratio	

Table D.53. Oracle::Data Dictionary Cache settings

D.9.7. Oracle::Disk Sort Ratio

The Oracle::Disk Sort Ratio probe monitors an Oracle database instance and collects the following metric:

- Disk Sort Ratio — The rate of Oracle sorts that were too large to be completed in memory and were instead sorted using a temporary segment.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Timeout*	30
Critical Maximum Disk Sort Ratio	
Warning Maximum Disk Sort Ratio	

Table D.54. Oracle::Disk Sort Ratio settings

D.9.8. Oracle::Idle Sessions

The Oracle::Idle Sessions probe monitors an Oracle instance and collects the following metric:

- Idle Sessions — The number of Oracle sessions that are idle, as determined by the required *Time Idle* value you provide. Only those sessions that have been idle for this duration, which is measured in seconds, are counted as idle sessions.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Time Idle (seconds)*	20
Timeout*	30
Critical Maximum Idle Sessions	

Field	Value
Warning Maximum Idle Sessions	

Table D.55. Oracle::Idle Sessions settings

D.9.9. Oracle::Index Extents

The Oracle::Index Extents probe monitors an Oracle instance and collects the following metric:

- Allocated Extents — The number of allocated extents for any index.
- Available Extents — The percentage of available extents for any index.

The required **Index Name** field contains a default value of % that matches any index name.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Index Owner*	%
Index Name*	%
Timeout*	30
Critical Maximum of Allocated Extents	
Warning Maximum of Allocated Extents	
Critical Maximum of Available Extents	
Warning Maximum of Available Extents	

Table D.56. Oracle::Index Extents settings

D.9.10. Oracle::Library Cache

The Oracle::Library Cache probe computes the Library Cache Miss Ratio so as to optimize the SHARED_POOL_SIZE in `init.ora`. It collects the following metrics:

- Library Cache Miss Ratio — The rate at which a library cache pin miss occurs. This happens when a session executes a statement that it has already parsed but finds that the statement is no longer in the shared pool.
- Executions — The number of times a pin was requested for objects of this namespace.
- Cache Misses — The number of pins of objects with previous pins since the object handle was created that must now retrieve the object from disk.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521

Field	Value
Timeout*	30
Critical Maximum Library Cache Miss Ratio	
Warning Maximum Library Cache Miss Ratio	

Table D.57. Oracle::Library Cache settings

D.9.11. Oracle::Locks

The Oracle::Locks probe monitors an Oracle database instance and collects the following metric:

- Active Locks — The current number of active locks as determined by the value in the v\$llocks table. Database administrators should be aware of high numbers of locks present in a database instance.

Locks are used so that multiple users or processes updating the same data in the database do not conflict. This probe is useful for alerting database administrators when a high number of locks are present in a given instance.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Timeout*	30
Critical Maximum Active Locks	
Warning Maximum Active Locks	

Table D.58. Oracle::Locks settings

D.9.12. Oracle::Redo Log

The Oracle::Redo Log probe monitors an Oracle database instance and collects the following metrics:

- Redo Log Space Request Rate — The average number of redo log space requests per minute since the server has been started.
- Redo Buffer Allocation Retry Rate — The average number of buffer allocation retries per minute since the server was started.

The metrics returned and the thresholds they are measured against are numbers representing the rate of change in events per minute. The rate of change for these metrics should be monitored because fast growth can indicate problems requiring investigation.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Timeout*	30

Field	Value
Critical Maximum Redo Log Space Request Rate	
Warning Maximum Redo Log Space Request Rate	
Critical Maximum Redo Buffer Allocation Retry Rate	
Warning Maximum Redo Buffer Allocation Retry Rate	

Table D.59. Oracle::Redo Log settings

D.9.13. Oracle::Table Extents

The Oracle::Table Extents probe monitors an Oracle database instance and collects the following metrics:

- Allocated Extents-Any Table — The total number of extents for any table.
- Available Extents-Any Table — The percentage of available extents for any table.

In Oracle, table extents allow a table to grow. When a table is full, it is *extended* by an amount of space configured when the table is created. Extents are configured on a per-table basis, with an extent size and a maximum number of extents.

For example, a table that starts with 10 MB of space and that is configured with an extent size of 1 MB and max extents of 10 can grow to a maximum of 20 MB (by being extended by 1 MB ten times). This probe can be configured to alert by (1) the number of allocated extents (e.g. "go critical when the table has been extended 5 or more times"), or (2) the table is extended past a certain percentage of its max extents (e.g. "go critical when the table has exhausted 80% or more of its max extents").

The required **Table Owner** and **Table Name** fields contain a default value of % that matches any table owner or name.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Table Owner*	%
Table Name*	%
Timeout*	30
Critical Maximum Allocated Extents	
Warning Maximum Allocated Extents	
Critical Maximum Available Extents	
Warning Maximum Available Extents	

Table D.60. Oracle::Table Extents settings

D.9.14. Oracle::Tablespace Usage

The Oracle::Tablespace Usage probe monitors an Oracle database instance and collects the following metric:

- Available Space Used — The percentage of available space in each tablespace that has been used.

Tablespace is the shared pool of space in which a set of tables live. This probe alerts the user when the total amount of available space falls below the threshold. Tablespace is measured in bytes, so extents do not factor into it directly (though each extension removes available space from the shared pool).

The required **Tablespace Name** field is case insensitive and contains a default value of % that matches any table name.

Field	Value
Oracle SID*	
Oracle Username*	
Oracle Password*	
Oracle Port*	1521
Tablespace Name*	%
Timeout*	30
Critical Maximum Available Space Used	
Warning Maximum Available Space Used	

Table D.61. Oracle::Tablespace Usage settings

D.9.15. Oracle::TNS Ping

The Oracle::TNS Ping probe determines if an Oracle listener is alive and collects the following metric:

- Remote Service Latency — The time it takes in seconds for the Oracle server to answer a connection request.

Field	Value
TNS Listener Port*	1521
Timeout*	15
Critical Maximum Remote Service Latency	
Warning Maximum Remote Service Latency	

Table D.62. Oracle::TNS Ping settings

D.10. RHN Satellite

The probes in this section may be applied to the RHN Satellite itself to monitor its health and performance. Since these probes run locally, no specific application or transport protocols are required.

D.10.1. RHN Satellite::Disk Space

The RHN Satellite::Disk Space probe monitors the free disk space on a Satellite and collects the following metrics:

- File System Used — The percent of the current file system now in use.
- Space Used — The file size used by the current file system.

- Space Available — The file size available to the current file system.

Field	Value
Device Pathname*	/dev/hda1
Critical Maximum File System Used	
Warning Maximum File System Used	
Critical Maximum Space Used	
Warning Maximum Space Used	
Critical Maximum Space Available	
Warning Maximum Space Available	

Table D.63. RHN Satellite::Disk Space settings

D.10.2. RHN Satellite::Execution Time

The RHN Satellite::Execution Time probe monitors the execution time for probes run from a Satellite and collects the following metric:

- Probe Execution Time Average — The seconds required to fully execute a probe.

Field	Value
Critical Maximum Probe Execution Time Average	
Warning Maximum Probe Execution Time Average	

Table D.64. RHN Satellite::Execution Time settings

D.10.3. RHN Satellite::Interface Traffic

The RHN Satellite::Interface Traffic probe monitors the interface traffic on a Satellite and collects the following metrics:

- Input Rate — The amount of traffic in bytes per second the device receives.
- Output Rate — The amount of traffic in bytes per second the device sends.

Field	Value
Interface*	eth0
Timeout (seconds)*	30
Critical Maximum Input Rate	
Critical Maximum Output Rate	

Table D.65. RHN Satellite::Interface Traffic settings

D.10.4. RHN Satellite::Latency

The RHN Satellite::Latency probe monitors the latency of probes on a Satellite and collects the following metric:

- Probe Latency Average — The lag in seconds between the time a probe becomes ready to run and the time it is actually run. Under normal conditions, this is generally less than a second. When

a Satellite is overloaded (because it has too many probes with respect to their average execution time), the number goes up.

Field	Value
Critical Maximum Probe Latency Average	
Warning Maximum Probe Latency Average	

Table D.66. RHN Satellite::Latency settings

D.10.5. RHN Satellite::Load

The RHN Satellite::Load probe monitors the CPU load on a Satellite and collects the following metric:

- Load — The load average on the CPU for a 1-, 5-, and 15-minute period.

Field	Value
Critical Maximum 1-minute Average	
Warning Maximum 1-minute Average	
Critical Maximum 5-minute Average	
Warning Maximum 5-minute Average	
Critical Maximum 15-minute Average	
Warning Maximum 15-minute Average	

Table D.67. RHN Satellite::Load settings

D.10.6. RHN Satellite::Probe Count

The RHN Satellite::Probe Count probe monitors the number of probes on a Satellite and collects the following metric:

- Probes — The number of individual probes running on a Satellite.

Field	Value
Critical Maximum Probe Count	
Warning Maximum Probe Count	

Table D.68. RHN Satellite::Probe Count settings

D.10.7. RHN Satellite::Process Counts

The RHN Satellite::Process Counts probe monitors the number of processes on a Satellite and collects the following metrics:

- Blocked — The number of processes that have been switched to the waiting queue and waiting state.
- Child — The number of processes spawned by another process already running on the machine.
- Defunct — The number of processes that have terminated (either because they have been killed by a signal or have called `exit()`) and whose parent processes have not yet received notification of their termination by executing some form of the `wait()` system call.

- Stopped — The number of processes that have stopped before their executions could be completed.
- Sleeping — A process that is in the **Interruptible** sleep state and that can later be reintroduced into memory, resuming execution where it left off.

Field	Value
Critical Maximum Blocked Processes	
Warning Maximum Blocked Processes	
Critical Maximum Child Processes	
Warning Maximum Child Processes	
Critical Maximum Defunct Processes	
Warning Maximum Defunct Processes	
Critical Maximum Stopped Processes	
Warning Maximum Stopped Processes	
Critical Maximum Sleeping Processes	
Warning Maximum Sleeping Processes	

Table D.69. RHN Satellite::Process Counts settings

D.10.8. RHN Satellite::Processes

The RHN Satellite::Processes probe monitors the number of processes on a Satellite and collects the following metric:

- Processes — The number of processes running simultaneously on the machine.

Field	Value
Critical Maximum Processes	
Warning Maximum Processes	

Table D.70. RHN Satellite::Processes settings

D.10.9. RHN Satellite::Process Health

The RHN Satellite::Process Health probe monitors customer-specified processes and collects the following metrics:

- CPU Usage — The CPU usage percent for a given process.
- Child Process Groups — The number of child processes spawned from the specified parent process. A child process inherits most of its attributes, such as open files, from its parent.
- Threads — The number of running threads for a given process. A thread is the basic unit of CPU utilization, and consists of a program counter, a register set, and a stack space. A thread is also called a lightweight process.
- Physical Memory Used — The amount of physical memory in kilobytes being used by the specified process.

- **Virtual Memory Used** — The amount of virtual memory in kilobytes being used by the specified process, or the size of the process in real memory plus swap.

Specify the process by its command name or process ID. (PID). Entering a PID overrides the entry of a command name. If no command name or PID is entered, the error **Command not found** is displayed and the probe is set to a CRITICAL state.

Field	Value
Command Name	
Process ID (PID) file	
Timeout*	15
Critical Maximum CPU Usage	
Warning Maximum CPU Usage	
Critical Maximum Child Process Groups	
Warning Maximum Child Process Groups	
Critical Maximum Threads	
Warning Maximum Threads	
Critical Maximum Physical Memory Used	
Warning Maximum Physical Memory Used	
Critical Maximum Virtual Memory Used	
Warning Maximum Virtual Memory Used	

Table D.71. RHN Satellite::Process Health settings

D.10.10. RHN Satellite::Process Running

The RHN Satellite::Process Running probe verifies that the specified process is running. Specify the process by its command name or process ID. (PID). Entering a PID overrides the entry of a command name. A Critical status results if the probe cannot verify the command or PID.

Field	Value
Command Name	
Process ID (PID) file	
Critical Number Running Maximum	
Critical Number Running Minimum	

Table D.72. RHN Satellite::Process Running settings

D.10.11. RHN Satellite::Swap

The RHN Satellite::Swap probe monitors the percent of free swap space available on a Satellite. A CRITICAL status results if the value falls below the Critical threshold. A WARNING status results if the value falls below the Warning threshold.

Field	Value
Critical Minimum Swap Percent Free	
Warning Minimum Swap Percent Free	

Table D.73. RHN Satellite::Swap settings

D.10.12. RHN Satellite::Users

The RHN Satellite::Users probe monitors the number of users currently logged into a Satellite. A CRITICAL status results if the value exceeds the Critical threshold. A WARNING status results if the value exceeds the Warning threshold.

Field	Value
Critical Maximum Users	
Warning Maximum Users	

Table D.74. RHN Satellite::Users settings

Glossary

Action	A task that is scheduled by a system administrator using Red Hat Network to be performed on one or more client systems. For example, an action can be scheduled to update the kernel packages on all the systems within a selected group.
Activation Key	RHN Management and Provisioning customers can generate activation keys through the RHN website. Each unique key can then be used to register a Red Hat system, entitle the system to RHN, subscribe the system to specific channels, and subscribe the system to RHN system groups through the command line utility <code>rhnreg_ks</code> from the <code>rhn_register</code> package.
Base Channel	A base channel is a type of Channel that consists of a list of packages based on a specific architecture and Red Hat release. For example, all the packages in Red Hat Enterprise Linux AS 3 for the x86 architecture make a base channel.
Bug Fix Alert	An Errata Alert that pertains to a bug fix.
Bugzilla	Bugzilla is an online application (http://www.redhat.com/bugzilla ¹) that allows users to communicate directly with the developers. From Bugzilla, users can submit bug reports and feature requests for Red Hat Enterprise Linux and related open source packages.
Channel	A channel is a list of packages. Channels are used to choose packages to be installed from client systems. Every client system must be subscribed to one Base Channel and can be subscribed to one or more Child Channel .
Child Channel	A child channel is a Channel associated with a Base Channel but contains extra packages.
Client System	See Registered System .
Digital Certificate	A client component in XML format that is stored in the <code>/etc/sysconfig/rhn/systemid</code> file on registered systems. Red Hat Network verifies this certificate to authenticate the registered system before each connection. This certificate is issued by Red Hat and passed to the system as part of the registration process. It includes unique information about the registered system to avoid fraudulent use.
Email Notification	Similar to an Errata Alert , except the information is delivered via email. If the email notifications option is selected, notifications are sent for every Red Hat Network Errata Alert . The email includes the type of Errata Alert, summary of the Errata, description of the Errata, and a list of which systems are affected by the report.
Enhancement Alert	An Errata Alert that pertains to a package enhancement request.

Entitled Server	A server that is subscribed to an RHN service level. Because the server is entitled, the RHN website can be used to manage its packages.
Errata	<p>Information published by Red Hat describing security fixes, bug fixes, and package enhancements for Red Hat Enterprise Linux. The information includes the topics of the Errata, Bugzilla bug IDs, relevant releases/architectures, solutions including required RPMs, and MD5 checksums for verification. Errata are also available at http://www.redhat.com/errata/. Each RHN <i>Errata Alert</i> is based on the Red Hat Enterprise Linux Errata List.</p> <p>Security issues and bug fixes are submitted by Red Hat engineers as well as the Linux community through Bugzilla which generates a bug report for each issue. Red Hat engineering evaluates the reports, resolves the bug, and generates new RPM packages. After the Red Hat quality assurance team tests new packages they are placed on the Red Hat Public File Server and on the Red Hat Network Server and an Errata is generated.</p>
Errata Alert	RHN Errata Alert that updated packages based on Red Hat Errata are available for one or more systems within an organization. There are three types of Errata Alerts: Security Alerts, Bug Fix Alerts, and Enhancement Alerts.
Management	One of the RHN service level offerings. It has more features than the Update service level, including user management, system groups, and enhanced system details.
Notification Method	An email address to which RHN Monitoring messages will be sent.
Package	All software in Red Hat Enterprise Linux is divided into software packages. Software updates are released in the form of RPM packages that can be installed on a Red Hat Enterprise Linux system.
Probe	A set of criteria that is either a template or a set of values assigned to a system that is used to measure the performance of a system.
Probe State	The measure of a probe's adherence to its defined criteria. States include: OK, Warning, Critical, Pending, Unknown
Probe Suite	collection or group of RHN Monitoring Probes.
Provisioning	One of the RHN service level offerings. It has more features than the Management service level, including kickstarting, reconfiguring, tracking, and reverting systems.
Red Hat Network Daemon	The RHN client daemon (rhnsd) that periodically polls Red Hat Network for scheduled actions.
Red Hat Network Registration Client	The RHN client application (rhn_register) that collects information about the client system, creates a <i>System Profile</i> and <i>Digital Certificate</i> , establishes a connection with the Red Hat Network servers, and registers the system with Red Hat Network.

Red Hat Update Agent	The RHN client application (up2date) that allows users to retrieve and install all updated packages for the client system on which the application is run. Use the Red Hat Update Agent Configuration Tool to configure its preferences, including whether to install the packages after they are downloaded.
Registered System	A system that is registered with Red Hat Network. Also known as a client system.
RPM	A software package manager that was developed by Red Hat Inc.. It can be used to build, install, query, verify, update, and uninstall software packages. All software updates from RHN are delivered in RPM format.
RPM Database	Each Red Hat Enterprise Linux system has an RPM database that stores information about all the RPM packages installed on the system. This information includes the version of the package, which files were installed with the package, a brief description of the package, the installation date, and more.
RPM Update	Red Hat Network option to deliver the RPM packages based on the Errata Alert list to a client system without user intervention. If this feature is selected, packages are delivered through the Red Hat Network Daemon running on the client system.
Satellite Administrator	Satellite Administrator are sets of users that have the highest level of control over an organization's Red Hat Network account. Members of this group can add users, systems, and system groups to the organization as well as remove them. An Satellite Administrator can also give users administrative privileges to system groups. An RHN organization must have at least one member of the Satellite Administrator group.
Security Alert	An Errata Alert that pertains to system security.
Service Level	A Red Hat Network subscription service. Different service levels offer different features of RHN. There are three paid service levels currently available: RHN Update, RHN Management, and RHN Provisioning.
Sibling	Siblings are virtual guests running on the same host. Virtual guests that run on separate hosts are not siblings.
Software Manager	The name of the first Service Level offering for Red Hat Network. Software Manager is now known as RHN Update .
System Directory	The System Directory section of Red Hat Network allows an organization to divide its client systems into system groups. Only members of the Satellite Administrator group can add systems to the organization.
System ID	A unique string of characters and numbers that identifies a registered system. It is stored in the system's Digital Certificate .

Glossary

System Profile	Hardware and software information about the client system. It is created during the registration process. The software information is a list of RPM packages and their versions installed on the client system. The System Profile is used to determine every Errata Alert relevant to each client system.
System Set Manager	Interface that allows users to perform actions on multiple systems. Actions include applying Errata Updates, upgrading packages, and adding/removing systems to/from system groups.
Update	One of the RHN service level offerings. Update was formerly called Basic. Update offers the same services as the Basic subscription did, plus more new features.
Virtual Guest	Any of the virtual instances running on the virtual host, under the control of the hypervisor. Also referred to as domain U or domU.
Virtual Host	The physical system that supports the hypervisor and all guest systems. The virtual host may also be referred to as domain 0, or dom0.
Yellowdog Updater Modified (yum)	The Yellowdog Updater Modified is the Red Hat Network client application (yum) that allows users to retrieve and install new or updated packages for the client system on which the application is run.

Appendix E. Revision History

Revision 1.0 Fri Feb 27 2009

Index

A

- account
 - deactivate, 64
- action
 - completed systems, 133
 - details, 133
 - failed systems, 133
 - in progress systems, 133
- activation key, 95
 - deleting, 97
 - disabling, 97
 - editing, 97
- activation keys
 - creating, editing, and deleting, 95
 - multiple use, 98
 - registration, 46
 - using, 46
- addresses
 - change, 63
- Apache
 - probes, 244
 - Processes, 244
 - Traffic, 245
 - Uptime, 245
- application programming interface
 - API, 239

B

- base channel, 117

C

- changing email address, 136
- changing password, 136
- channel
 - configuration
 - create, 126
- Channel List, 117
- Channels
 - Software and Configuration Files, 117
- channels, 117
 - all, 118
 - base, 117
 - child, 117
 - errata, 120
 - list of, 117
 - My, 119
 - packages, 120
 - Popular, 119

- Red Hat, 119
- Shared, 119
- Channels and Packages
 - Channel List, 117
- child channel, 117
- client applications
 - obtaining, 5
 - redirecting, 206
- client systems
 - configuring, 206
 - registering, 207
 - updating, 208
- Cobbler, 189
- cobbler, 189
- config management
 - system preparation, 125
- configuration
 - actions, 124
 - channel
 - create, 126
 - files, 124
 - Schedule, 124
- Configuration Management
 - command line tools, 229
- create
 - configuration
 - channel, 126
- custom information
 - about systems, 76

D

- deactivate
 - user, 135
- delete
 - user (RHN Satellite only), 135
- deleting a system, 71
- Digital Certificate, 5

E

- email address
 - change, 64
 - changing, 136
- entitlement
 - with activation key, 95
- Errata, 112
 - Advanced Search, 115
 - All Errata, 113
 - apply applicable, 77
 - Relevant Errata, 113
- Errata Alert Icons
 - explanation of, 59

Errata notifications
 automatic updates, 4
Errata Updates
 applying, 114
 searching, 115
 viewing details, 114
 viewing list of all errata, 113
 viewing list of applicable errata, 113
EUS (see Extended Update Support)
Extended Update Support, 117

G

General
 probes, 249
 Remote Program, 249
 Remote Program with Data, 249
 SNMP Check, 250
 TCP Check, 251
 UDP Check, 251
 Uptime (SNMP), 251
getting started, 5
GNU Privacy Guard, 5

H

hardware profile
 Updating on server, 76
Help Desk, 146
HTTP Proxy, 52

I

initialization script
 /etc/init.d/rhnsd, 49
 /etc/rc.d/init.d/rhnsd, 49

K

kickstart
 explained, 101
Koan, 189
koan, 189

L

Linux
 CPU Usage, 252
 Disk IO Throughput, 252
 Disk Usage, 253
 Inodes, 253
 Interface Traffic, 254
 Load, 254
 Memory Usage, 255
 probes

 nocpulse, 252
 Process Count Total, 256
 Process Counts by State, 255
 Process Health, 256
 Process Running, 257
 Swap Usage, 258
 TCP Connections by State, 258
 Users, 259
 Virtual Memory, 259
List Navigation
 explanation of, 60
LogAgent
 Log Pattern Match, 260
 Log Size, 261
 probes
 nocpulse, 260

M

macros
 within configuration Files
 interpolation, 128
Management
 service level, 2
manual installation
 System Profile, 40
Monitoring, 139
 All, 141
 Critical, 140
 Current State, 141
 General Config, 145
 introduction, 151
 Notification, 141
 OK, 141
 Pending, 141
 prerequisites, 151
 Scout Config Push, 144
 service level, 4
 Status, 139
 Unknown, 141
 Warning, 140
monitoring
 list of probes, 243
MySQL, 154
 Database Accessibility, 262
 Open Tables, 263
 Opened Tables, 262
 probes, 262
 Query Rate, 263
 Threads Running, 263
mysql package, 154

N

- navigation, 55
- Network Services
 - DNS Lookup, 264
 - FTP, 264
 - IMAP Mail, 265
 - Mail Transfer (SMTP), 265
 - Ping, 265
 - POP Mail, 266
 - probes, 264
 - Remote Ping, 267
 - RPCService, 267
 - Secure Web Server (HTTPS), 268
 - SSH, 268
 - Web Server (HTTP), 269

notes

- about systems, 76

Notification

- filter, 145

Notifications

- Monitoring, 155

notifications

- creating methods, 155
- deleting methods, 157
- filtering, 157
- receiving, 155
- redirecting, 156

ntsysv, 49

O

Oracle

- Active Sessions, 270
- Availability, 271
- Blocking Sessions, 271
- Buffer Cache, 271
- Client Connectivity, 272
- Data Dictionary Cache, 272
- Disk Sort Ratio, 273
- Idle Sessions, 273
- Index Extents, 274
- Library Cache, 274
- Locks, 275
- probes, 270
- Redo Log, 275
- Table Extents, 276
- Tablespace Usage, 276
- TNS Ping, 277

Overview, 62

- Account Deactivation, 64
- Addresses, 63
- Email, 64

Help, 146

Your Account, 63

Your Preferences, 64

overview of website, 56

P

package installation

- scheduled, 4

package list

- Updating on server, 40, 77

Package Updater (pup)

- complete description, 15

packages

- filter, 121

password

- change, 63

port 22, 153

port 4545, 151

preferences

- change, 64
- language, 65
- locale, 65

probe

- guidelines, 243

probe list

Apache

- Processes, 244
- Traffic, 245
- Uptime, 245

General

- Remote Program, 249
- Remote Program with Data, 249
- SNMP Check, 250
- TCP Check, 251
- UDP Check, 251
- Uptime (SNMP), 251

Linux

- CPU Usage, 252
- Disk IO Throughput, 252
- Disk Usage, 253
- Inodes, 253
- Interface Traffic, 254
- Load, 254
- Memory Usage, 255
- Process Count Total, 256
- Process Counts by State, 255
- Process Health, 256
- Process Running, 257
- Swap Usage, 258
- TCP Connections by State, 258
- Users, 259

- Virtual Memory, 259
- LogAgent
 - Log Pattern Match, 260
 - Log Size, 261
- MySQL
 - Database Accessibility, 262
 - Open Tables, 263
 - Opened Tables, 262
 - Query Rate, 263
 - Threads Running, 263
- Network Services
 - DNS Lookup, 264
 - FTP, 264
 - IMAP Mail, 265
 - Mail Transfer (SMTP), 265
 - Ping, 265
 - POP Mail, 266
 - Remote Ping, 267
 - RPCService, 267
 - Secure Web Server (HTTPS), 268
 - SSH, 268
 - Web Server (HTTP), 269
- Oracle
 - Active Sessions, 270
 - Availability, 271
 - Blocking Sessions, 271
 - Buffer Cache, 271
 - Client Connectivity, 272
 - Data Dictionary Cache, 272
 - Disk Sort Ratio, 273
 - Idle Sessions, 273
 - Index Extents, 274
 - Library Cache, 274
 - Locks, 275
 - Redo Log, 275
 - Table Extents, 276
 - Tablespace Usage, 276
 - TNS Ping, 277
- RHN Satellite
 - Disk Space, 277
 - Execution Time, 278
 - Interface Traffic, 278
 - Latency, 278
 - Load, 279
 - Probe Count, 279
 - Process Counts, 279
 - Process Health, 280
 - Process Running, 281
 - Processes, 280
 - Swap, 281
 - Users, 282

- WebLogic
 - Execute Queue, 246
 - Heap Free, 247
 - JDBC Connection Pool, 247
 - Server State, 248
 - Servlet, 248
- Probes
 - Monitoring, 157
- probes
 - Apache, 244
 - General, 249
 - Linux, 252
 - LogAgent
 - nocpulse, 260
 - managing, 158
 - MySQL, 262
 - Network Services, 264
 - on the RHN Server, 159
 - Oracle, 270
 - RHN Satellite, 277
 - thresholds, 158
 - WebLogic, 246
- Provisioning
 - service level, 3
- proxy server
 - with Red Hat Network Alert Notification Tool, 52
 - with Red Hat Network Registration Client, 214
 - with Red Hat Update Agent, 41

Q

- quality assurance
 - overview, 5
- Quick Search
 - explanation of, 59

R

- reactivating
 - systems, 75
- Red Hat Enterprise Linux 2.1
 - requiring the Red Hat Network Registration Client, xi, 21
- Red Hat Enterprise Linux 5
 - rhn_register, 7
- Red Hat Network
 - an introduction to, 1
 - components
 - primary, 1
- Red Hat Network Actions Control
 - rhn-actions-control, 229
- Red Hat Network Alert Notification Tool

- adding to panel, 51
- applying Errata Updates, 54
- configuring, 51
- icons, 53
- launching RHN website, 54
- requirements, 51
- with a proxy server, 52
- Red Hat Network Configuration Client
 - rhncfg-client, 230
- Red Hat Network Configuration Manager
 - rhncfg-manager, 232
- Red Hat Network Daemon, 49
 - configuring, 49
 - disabling, 49
 - initial description, 1
 - troubleshooting, 49
 - using to apply Errata Updates, 114
 - viewing status, 49
- Red Hat Network Monitoring Daemon (rhnmmd) monitoring daemon, 151
 - installation, 152
 - probes requiring the daemon, 152
 - SSH key installation, 154
 - using sshd instead, 153
- Red Hat Network packages
 - comparison, 6
- Red Hat Network Registration Client
 - initial description, 2
- Red Hat packages
 - for UNIX, 203
 - installing, 203
- Red Hat Update Agent, 54
 - Command Line Arguments, 36
 - configuration, 41
 - UNIX Command Line Arguments, 211
 - with a proxy server, 41
- Red Hat Update Agent (up2date)
 - activation keys, 46
 - command line options, 37
 - command line version, 36, 45
 - complete description, 21
 - configuration tool, 41
 - configuring general settings, 41
 - configuring package exceptions, 44
 - configuring retrieval and installation, 42
 - excluding packages, 44
 - graphical options, 22
 - initial description, 1
 - installing GPG keys, 39
 - log file, 41
 - registering with, 24
 - starting, 21
 - synchronizing system profile, 40
- reference guide
 - bug reporting, xii
 - conventions, xi
 - introduction to the, xi
- registering
 - with activation keys, 46
- Registration, 213
 - as part of an organization, 220
 - Configuration, 213
 - Email notification, 218
 - Hardware System Profile, 220
 - Password, 218
 - RPM Package List, 222
 - Software System Profile, 222
 - System Profile, 218, 220
 - text mode, 227
 - through the Web, 61
 - username, 218
 - with a proxy server, 214
 - with activation key, 95
- remote commands
 - enabling, 211
 - issuing, 212
- RHN Satellite
 - Disk Space, 277
 - Execution Time, 278
 - Interface Traffic, 278
 - Latency, 278
 - Load, 279
 - Probe Count, 279
 - probes, 277
 - Process Counts, 279
 - Process Health, 280
 - Process Running, 281
 - Processes, 280
 - Swap, 281
 - Users, 282
- RHN Tools channel, 153
- RHN website, 54
 - initial description, 1
- rhncatalog
 - troubleshooting with, 159
- rhncatalog
 - options, 160
 - troubleshooting with, 160
- rhnmmd daemon, 153
- rhnmreg_ks, 95
- rhnsd, 49
- rhncatalog_register (see Registration)

complete description, 7

S

Satellite Administrator, 136

Schedule, 131

Scheduled Actions

Action Details, 133

Actions List, 132

Archived Actions, 132

Completed Actions, 132

Failed Actions, 132

Pending Actions, 131

Scout Config Push, 151

Secure Sockets Layer, 5

security

overview, 5

service levels

Management, 2

Monitoring, 4

Provisioning, 3

Update, 2

Software

Channel List

Channel Details, 120

Package Search, 122

software

searching, 122

software channels

details, 120

SSH, 153

SSH key, 154

sshd, 153

SSL

setting up, 206

SSL certificates

deploying, 206

SSL expiration errors

connection

certificate verification, 5

system group, 84

adding and removing, 85

creating, 85

deleting, 86

editing details, 86

list of, 84

viewing details, 86

system group list

status, 85

System Groups

assigning and removing, 83

System Group List, 84

system groups

joining and leaving, 83

system list, 67

System Profile, 220

Custom Information, 76

Notes, 76

Reactivation, 75

Updating hardware profile, 76

Updating package list, 40, 77

Updating Properties, 74

System Set Manager, 87

systems

deleting, 71

entitling, 65

overview, 67

searching, 93

viewing a list of, 67

viewing details for, 70

Systems

Advanced Search, 93

Entitlements, 65

System Details, 70

System List, 67

Systems Overview, 67

systems list

status, 67

Systems Selected

explanation of, 59

T

Troubleshooting

Monitoring, 159

U

UNIX variants (see supported)

Update

service level, 2

updating

via command line, 210

via website, 210

user

deactivate, 135

delete (RHN Satellite only), 135

user account, 218

user roles, 136

users, 133

changing email address, 136

changing password, 136

roles, 136

V

variables

macros

in configuration files, 128

W

WebLogic

Execute Queue, 246

Heap Free, 247

JDBC Connection Pool, 247

probes, 246

Server State, 248

Servlet, 248

website, 55

activation keys, 95

All Errata, 113

Channel List, 117

Channels, 117

custom system information, 98

Errata, 112

Erratum Search, 115

Help, 146

language, 65

locale, 65

logging in, 60

Monitoring, 139

navigation bar, 55

overview, 55

Overview, 62

Relevant Errata, 113

Schedule, 131

Software Channel Details, 120

Software Search, 122

stored profiles, 98

System Details, 70

System Entitlements, 65

System Group List, 84

System Groups, 84

System List, 67

System Search, 93

Systems, 67

Systems Overview, 67

Users, 133

Your Account, 63

