



Best Practices for Deploying and Managing Linux with Red Hat Network

Abstract

This technical whitepaper provides a best practices overview for companies deploying and managing their open source environment through Red Hat Network (RHN).

Table of Contents

Red Hat Network Overview	2
Deploying and Managing Open Source Solutions	5
Part I: Setting Up Your Environment	5
Part II: Registering and Tasking Your Systems	11
Kickstart with RHN	12
Part III: Managing Your Systems	13
Appendices:	
1. Getting more from RHN—The API Access Layer	22
2. Running RHN in Highly Secure Environments	23
3. Key Glossary Terms	24



Red Hat Network Overview

Customers today demand much more from their technology than just bits—they need full IT solutions that include deployment, patch, monitoring, and configuration to solve their customers' problems, reduce costs and complexity, increase productivity, and enhance security. These tools need to be tightly integrated with content, based on industry standards, and easy to integrate with the existing environment.

The Red Hat Network solution allows customers to choose the level of services and architectural models required depending on IT needs. RHN is integrated with Red Hat® Enterprise Linux® and other Red Hat offerings to ensure customers are able to manage systems effectively and with minimum complexity.

Service Entitlements

The first question customers must answer is what kind of service entitlements they want their systems to have. Red Hat Network currently consists of four service modules: Update, Management, Provisioning, and Monitoring¹. Customers purchase entitlements to these services on an annual per-system subscription basis.

Update Module—Included with every Red Hat Enterprise Linux subscription. Update Module is the entry-level offering for RHN. It allows you to easily maintain single systems and includes functionality such as a graphical user interface, priority notification, errata information, RPM dependency checking, and auto update.

Management Module—Management Module allows you to manage your entire Red Hat Enterprise Linux or Sun® Solaris™ infrastructure. Designed for enterprise scalability, the Management Module features systems grouping, role-based administration for policies and permissions, scheduled actions, and higher-end functionality with Satellite Server such as third-party channels, custom channels, local package caching, and off-network capability.

¹ For a comparison of the different service entitlements, visit www.redhat.com/software/rhn/.

Provisioning Module—The Provisioning Module enables you to manage the complete life cycle of your Linux infrastructure. The Provisioning Module includes features such as OS provisioning (from bare metal boxes or previously deployed boxes), configuration management, multi-state rollback, scheduled remote actions, Kickstart configuration tools, and RPM-based application provisioning.

Monitoring Module—The Monitoring Module allows you to maintain the availability of your applications on Red Hat Enterprise Linux with performance monitoring. The module includes monitoring probes, probe suites, notifications, and reporting.

RHN service entitlements are priced on a per system (or per node) basis. This best practices overview assumes the purchase of all entitlements for all examples provided.

Architectural Model

The second question customers must answer is what is their preferred Architectural Model. RHN offers two basic models:

Hosted Model—The customer uses a back-end (database, application, and web server) hosted by Red Hat. Each system connects individually to RHN across the Internet. This model is most effective when managing single systems or a small number of systems.

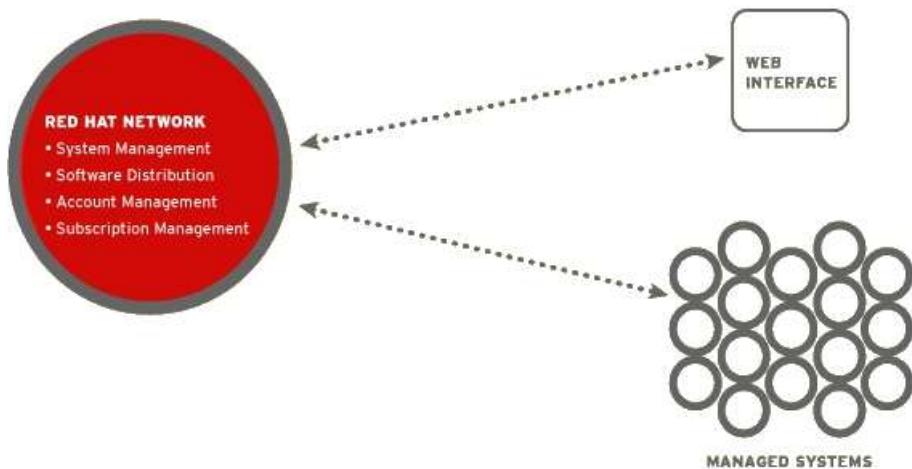


Figure 1. Hosted Architectural Model

Satellite Model—The entire RHN solution is placed on the customer's local network. Because RHN is local, it can be optimized and configured for each customer, providing additional functionality and security. Satellite Server can run in a connected (to Red Hat) or disconnected (completely offline) mode. Purchase of a Satellite Server includes Red Hat Enterprise Linux, the RHN Satellite Server software, 24 x 7 worldwide support, and an embedded Oracle database (optional).

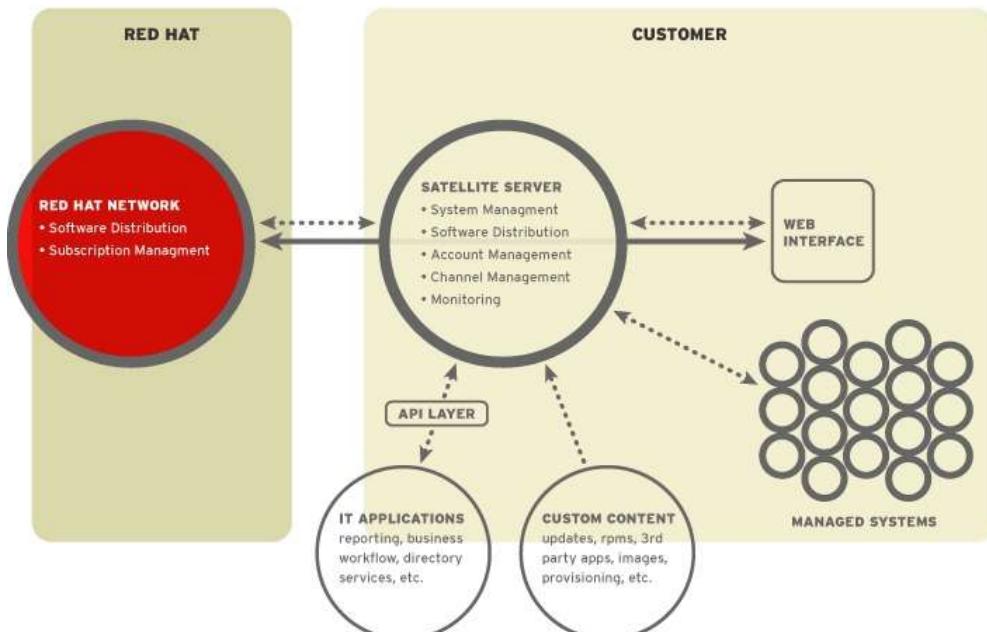


Figure 2. Satellite Architectural Model

The Hosted model is the default option when customers purchase an RHN entitlement. A Satellite Server, which is capable of serving thousands of systems, must be purchased separately. Contact your local sales representative for more information regarding scalability with a Satellite Server.

Red Hat also offers another option: Proxy Server. The Proxy Server can be run in either a Hosted or Satellite architecture and functions as an intelligent caching box on the customer's local network. The Proxy Server connects directly to RHN servers in a Hosted environment and directly to the Satellite Server in a Satellite environment.

This paper assumes the purchase of the Satellite Server (and the use of Proxy Servers where appropriate) for all examples.²

2 For specific pricing information, visit www.redhat.com/software/rhn/purchase/ or contact a Red Hat sales representative toll free (US only) at 1-866-273-3428 x45606.

Deploying and Managing Open Source Solutions

This whitepaper walks through a typical scenario most system administrators might find when deploying and managing their Linux systems. Much of the functionality described can be used in alternate ways.

The use case assumes that all systems are racked, powered, and networked. It also assumes that the system administrator is responsible for a dynamic environment, including environments where:

- New servers are being added on a “net new” basis
- Large scale migrations are taking place (such as UNIX or Windows to Linux)
- Redeployments or upgrades are being performed
- Both central and remote locations are being managed

Part I: Setting Up Your Environment

Install Satellite

The first step to configuring your RHN environment is to set up your Satellite Server. This system connects back to Red Hat Network (unless you run disconnected) and serve as the central repository and hub of connection for all of your client systems.

To optimize your installation, Red Hat recommends purchasing a one-week Professional Services consulting package. A Red Hat professional will come to your site to and train your staff on troubleshooting, impending changes, and enhanced functionality available for your Satellite Server.

If you choose to install the Satellite Server on your own, you still receive 24x7x365 installation and production support. For a list of technical and hardware requirements needed to install Satellite Server, visit
www.redhat.com/software/rhn/requirements/.

Identify Content

After installing the Satellite Server, customers must identify the content that they wish to manage. This content will consist of the following:

- *Red Hat Enterprise Linux base channels*—A base channel provides content for Red Hat Enterprise Linux. Base channels include Red Hat Enterprise Linux AS, ES, WS, and Desktop, Red Hat Application Server, etc.
- *Red Hat Enterprise Linux core build(s)*—Many customers choose to develop a core Linux build for their environment. For more information regarding core builds, speak with your local salesperson.
- *Custom content*—Any custom content (applications or otherwise) that you would like to distribute, install, and manage on your systems via RHN.
- *Third-party applications or content*—Third-party applications you wish to manage with RHN.
- *Solaris content*—Solaris packages, patches, and patch sets can be distributed.
- *Configuration files*—All text-based configuration files can be stored and managed using RHN.

Red Hat Network uses the RPM package format to manage all of this content with the exception of Solaris content and configuration files. Customers who package custom and third-party applications in this format can update, manage, and provision content throughout their enterprise environment just as they would Red Hat Enterprise Linux content. In addition, Red Hat Network provides tools to assist you in pushing your RPM content into Red Hat Network.

Build Repositories

Once the content has been identified, the next step is to build and populate different repositories.

Assign all content to a base or child channel—RHN uses base channels to represent the main or parent channel to which a system belongs. This channel takes precedence over all other channels to which the system is subscribed. Child channels are sub-channels or second-tier channels. A system can only have a single base channel but can have as many child channels as desired. RHN allows your company to establish permissions for channels as well as systems, ensuring only the proper users have access to content and systems as shown in Figure 3.

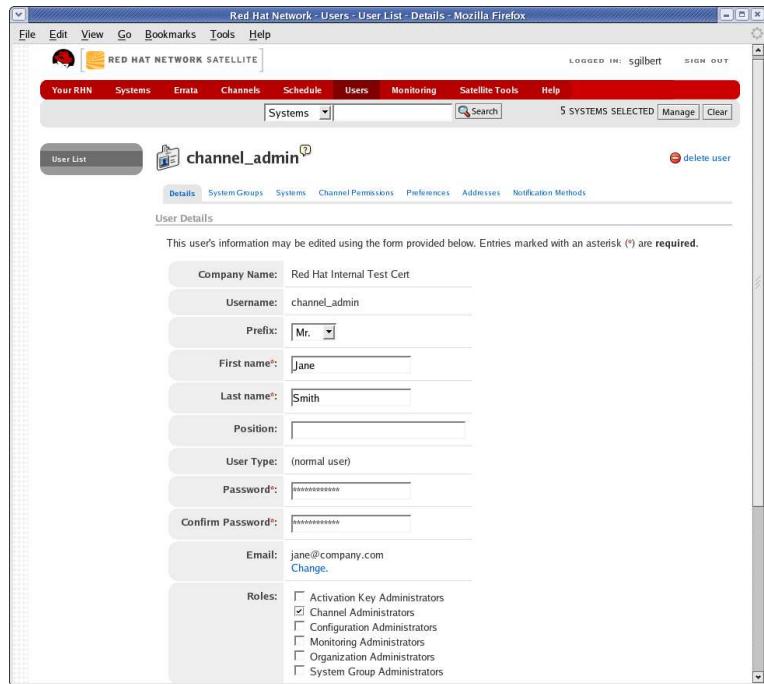


Figure 3. System and Channel Permissions per User

Assign channel permissions—Red Hat Network allows you to assign permissions to distinct individuals or groups. Only people with permission to access a channel can upload new content or make changes to the existing content.

Push content to channels—The next step is to push the content into channels through a simple command line tool called “push.” This process takes all RPM files and sends them into channels. You may automate this process with nightly/weekly scripts or other options appropriate for your environment.

Stage Content

When a new patch or errata is received, the last thing any system administrator wants to do is blindly apply that package to all their production systems without testing it. To facilitate the testing process, RHN provides the ability to develop staged environments. While your particular environment may vary, a typical testing environment includes the following stages:

- *Red Hat base channel*—This channel that receives new packages from Red Hat. This channel regularly receives new content and is updated at an interval set by the customer, typically on an hourly or daily basis.
- *Development*—This channel receives selected packages from the base channel, but only those packages that the organization wants to apply to their systems. Developers can then test and configure packages in the channel.
- *Testing and QA*—Once the developers are finished developing packages, they are pushed to the testing and QA channels. The QA organization will then install these packages on test machines subscribed to the “Testing & QA channel” as their parent channel. After testing these packages, the QA department will push the packages that passed QA to the production stage. Packages that failed QA will either stay in QA for further testing or be sent back to development.
- *Production*—Once the production stage receives packages, they can be installed on production systems. These systems are registered to the production channel as their parent channel. By breaking the channels into a staged environment, RHN makes it easier for companies to manage the flow and testing of their content to production systems.

A few features are used to successfully implement this process:

- *Clone and manage channels*—Provides the ability to clone and then rename an entire channel. For example, customers might have a channel called “Red Hat Enterprise Linux AS 4 Development” and wish to clone that channel to create “Red Hat Enterprise Linux 4 AS Test/QA.” This step is usually done in the beginning when first establishing the channels.

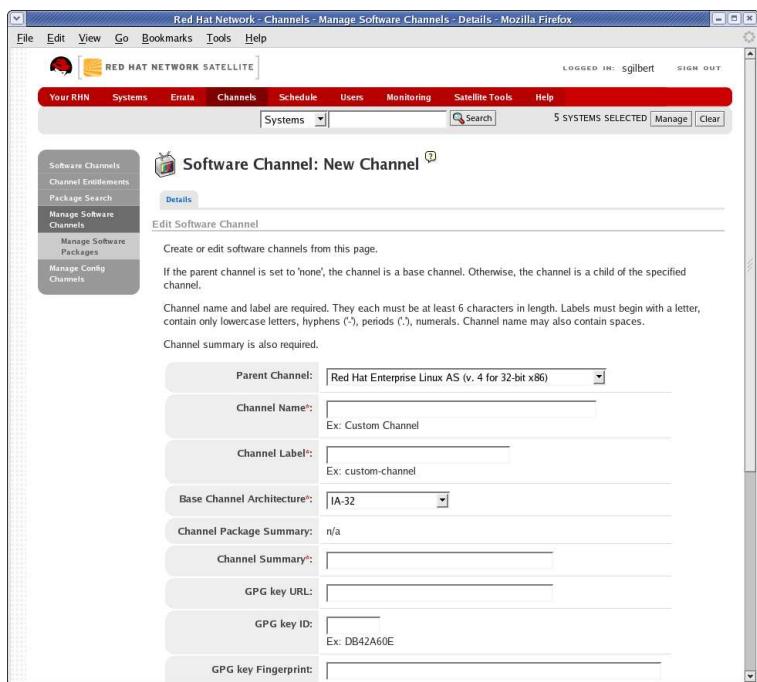


Figure 4. Channel Creation

- *Clone and manage errata*—This provides the ability to clone errata between the channels. This is used when you would like to move an errata between the two channels but do not wish to clone the entire channel(s). RHN allows you to clone and replicate errata individually or in groups.
- *Auto update*—Some customers choose to auto update their systems, especially when using staged environments where any new content will be rigorously tested before making it to the production channel. When using auto-update, the user can configure the client to receive updates on a regular time interval (established by the user), and then install all changes that have been added to the system's parent and/or child channels.

Define System Groups and Permissions

With the content staged and built, it then becomes necessary to understand how to manage the system infrastructure. RHN allows you to group systems together so that you can manage the entire group as easily as you would manage a single system. This process includes the following steps:

- *Define Groups*—RHN understands that you have different

ways of seeing your infrastructure. You might want to view by hardware type, function, location, or some other means.³ RHN allows you to group systems in a many-to-many format. This means that not only can a group have multiple systems assigned to it, a system can also be assigned to multiple groups. For example, a database server from XYZ Vendor located in Raleigh, NC could belong to the “XYZ Vendor,” “Database,” and “Raleigh” groups simultaneously.

<input type="checkbox"/>	Status	Monitoring	Group Name	Systems	Use in SSM
<input type="checkbox"/>	✓		Example Corp Development DB Servers	2	Use Group
<input type="checkbox"/>	✓		Example Corp Development File Servers	2	Use Group
<input type="checkbox"/>	✓		Example Corp Production DB Servers	4	Use Group
<input type="checkbox"/>	!		Example Corp Production File Servers	3	Use Group
<input type="checkbox"/>	!	!	idc-group	3	Use Group
<input type="checkbox"/>	!	!	jamo-group	3	Use Group
<input type="checkbox"/>	✓		jfautley	0	Use Group
<input type="checkbox"/>	!	!	matt's-group	5	Use Group
<input type="checkbox"/>	!	!	Motive Demo Group	3	Use Group
<input type="checkbox"/>	!	!	se-group	1	Use Group
<input type="checkbox"/>	!	!	sgilbert-app	5	Use Group
<input type="checkbox"/>	!	!	sgilbert-web	6	Use Group

Figure 5. Systems Grouping

- **Assign System Permissions**—Just like channel permissions, system permissions allow you to restrict access to different systems or groups of systems. This helps to ensure that a centralized security process is in place and only those with authorization are allowed to make changes to a system.

Build Activation Keys

Activation keys are the secret to fast and easy registration of your systems. An activation key is a key that when applied to a system upon registration, automatically assigns that system to a predefined set of channels, group(s), policies, and permissions. The use of an activation key allows any user (as defined by the

³ You can also apply personalized asset tags to each machine.

organization administrator) to assign systems to distinct groups in the organization as well as streamline the process for getting the system properly deployed and into management. The use of activation keys will be further explored in the next section, but for now it is enough to know that they exist and should be mapped out as part of the process of setting up your environment.

Part II: Registering and Tasking Your Systems

1. *System already running Red Hat Enterprise Linux.* If your system already has Red Hat Enterprise Linux installed, register the system with RHN by running the register command. When you register, you can also give the system an activation key to automatically assign it the proper permissions, group(s), and channels (content and configurations). If you are registering multiple systems, you can write a simple script to implement the action on all systems.
2. *System not running Red Hat Enterprise Linux–DVD Install.* Install Red Hat Enterprise Linux via the installation DVD or other form of media. The process is then the same as above.
3. *System not running Red Hat Enterprise Linux–PXE Boot Install.* Set up a DHCP and PXE server that uses a Kickstart script hosted on RHN to provide images for the OS. During this process, the Kickstart script can also assign the necessary activation key and complete the your system's registration.

Once the system(s) has been registered with RHN, you are ready to task them. You can task them in the following ways:

1. *Matching another system.* Point your newly registered system at an existing system and ask RHN to replicate the desired characteristics on the new system. Then run up2date on the new system, and RHN will completely update your box so that it mirrors the original one.
2. *Matching a stored profile.* Like the process above, this is done using a system image rather than a live system. Some companies prefer to bring all of their systems to a known “generic” or “master” state at registration, which allows them a

consistent baseline from which to fully task their systems.

3. *Scheduling a series of actions to customize your system(s).*

Used alone or in conjunction with steps 1 and 2, you can schedule actions to completely task your system. Change activation key defaults (or set them), packages, configuration channels, add child channels, set permissions, etc., then update your system to bring it to the fully updated state and it is ready to go.

Kickstart with RHN

Kickstart is a way to automate installation of Red Hat Enterprise Linux on your systems. This is accomplished by creating a file (ks.cfg) that contains responses to all the questions asked by the installation program during interactive installation. Once the file is created, it can be copied onto a normal Red Hat Enterprise Linux boot disk or saved in Red Hat Network. When this disk or disk image is used to boot a computer, the booting sequence finds the Kickstart file and automatically installs based on the values in the file.

This type of installation is useful for several reasons. First and foremost, it is much faster and easier than manually entering all the information. Second, the Kickstart file makes it easier to execute identical installations. Third, installation can be done over a network, meaning that several nodes can be installed simultaneously. Finally, the %post section can automate many configuration details that would normally have to be executed by hand after the installation is complete.

Red Hat Network adds value to the process by providing you with the ability to:

- *Create multiple Kickstart files.* RHN allows you to create your own Kickstart file through a simple GUI interface that automates the creation process. RHN can also create a Kickstart script from stored profiles. When run, the script will create a clone of the profiled system.

- *Remotely administer Kickstart files to systems.* RHN provides remote administration to systems, including the ability to reprovision systems, saving you the hassle of having to use a boot disk at each individual system.
- *Combine your Kickstart actions with other RHN actions for a complete automated deployment.* You can insert commands into the %post section of the Kickstart script or schedule actions to occur after the Kickstart script runs. Through the efficient use of Kickstart with RHN and other RHN commands, you can control your environment in a completely remote and efficient manner.

Part III: Managing Your Systems

Now that you have deployed and tasked all of your systems, you can begin using and managing them. This section outlines how simple patching, managing, (re)deploying, and monitoring your systems can be with RHN.

Patching your Systems

The following steps keep your systems updated and easily patched:

- *Obtain Notification*—RHN notifies you of new errata for your systems. For individual users, the applet on your GUI will flash when updates are available. For users managing multiple systems through the web interface, checking the channels, systems, or groups in RHN displays those that have received updated errata. Administrators also receive email with information about any new errata.
- *Understand Errata Information*—Errata information is included in email sent to users and is presented on the web interface. RHN provides this information so that users get a complete understanding of what the patch is and why it is being applied. Also included are notes about configuration specifications or other information that IT professionals may find useful.

- *Evaluate the patches through the staged environments*—At this point, your Satellite Server will have already synced with updated content on RHN unless you are running in a disconnected mode. The timing and interval of this synchronization process is established by the Satellite Org Admin. As discussed in the *Stage Content* section above, you can now use the errata cloning and management functionality to move errata through the different staged environments in your infrastructure.
- *Customize errata as necessary*—Errata management can also be used to make customized changes to an errata as it moves from stage to stage. For example, you might want to apply an errata to two different groups—one in Tokyo, the other in Atlanta. If these errata contain different configuration instructions, you can easily specify those instructions in the errata and make sure that each group receives a custom errata specific for their environment.

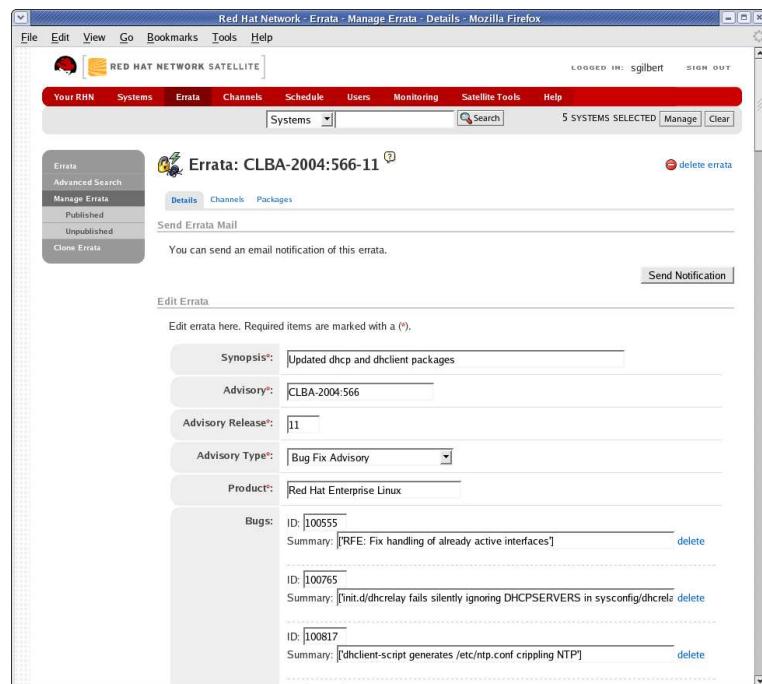


Figure 6. Customizing your Errata

- *Define Groups*—Once the errata has reached the channel(s) where it is ready for distribution, RHN provides various ways to update your systems. You can choose an entire predefined group, a single system, or customize a temporary group. To customize your group, use the RHN system search

functionality (Figure 7) and System Set Manager (Figure 8). These features allow you to search for systems with various characteristics (packages, hardware, drivers, custom tags, location, etc.), and then click on the groups that you would like to apply actions to. System Set Manager allows you to merge and combine search results in many ways.

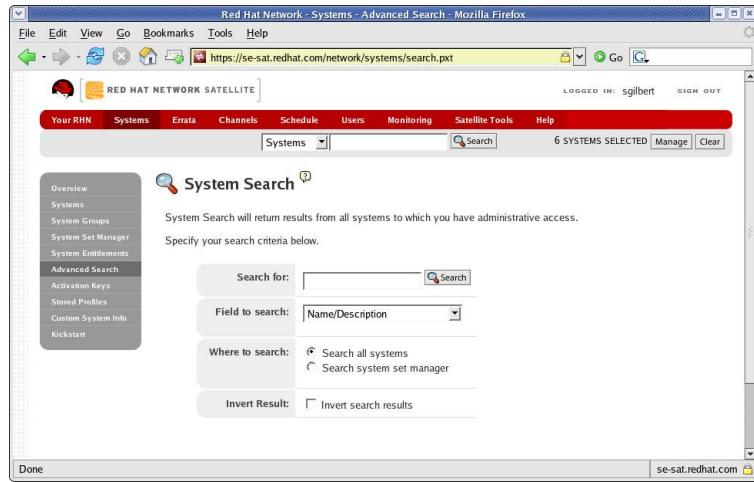


Figure 7. System Search

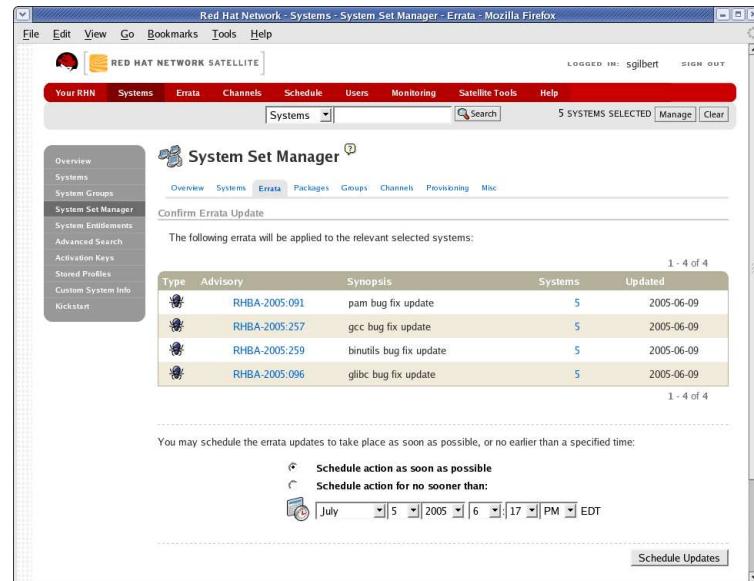


Figure 8. System Set Manager

- **Schedule necessary actions**—With your systems and groups now identified (either through predefined rules or through the System Set Manager), you are now ready to schedule actions such as patches and configuration changes or even complete a re-provisioning of the systems. Pick the packages to

update, apply these actions against your selected systems, schedule the action to take place at a time convenient to your environment, and then let the action take place. You should note also that you can schedule multiple actions against a system at once.

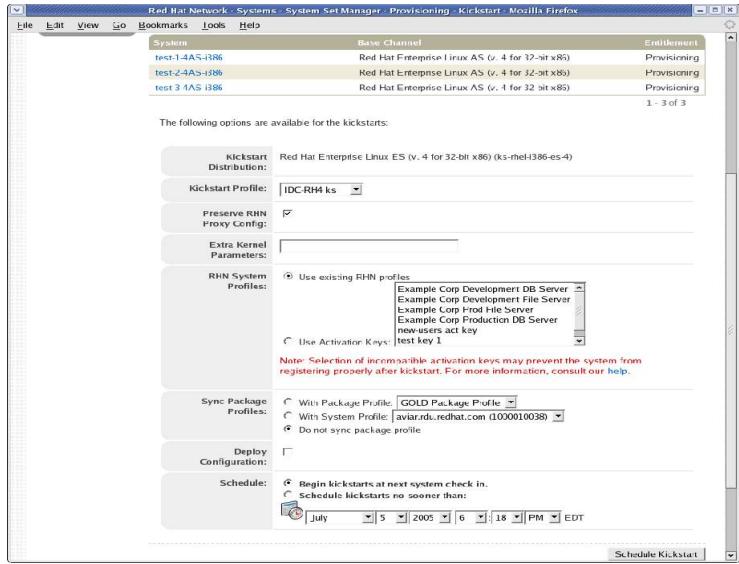


Figure 9. Provision your Systems

- *Schedule any arbitrary actions or reboot*—Once you have scheduled your actions, you can also schedule additional arbitrary actions or necessary reboots for a machine. RHN allows you to configure the clients to give system administrators root level access to a system. This functionality allows administrators to centrally schedule necessary actions against boxes in remote or distributed locations.

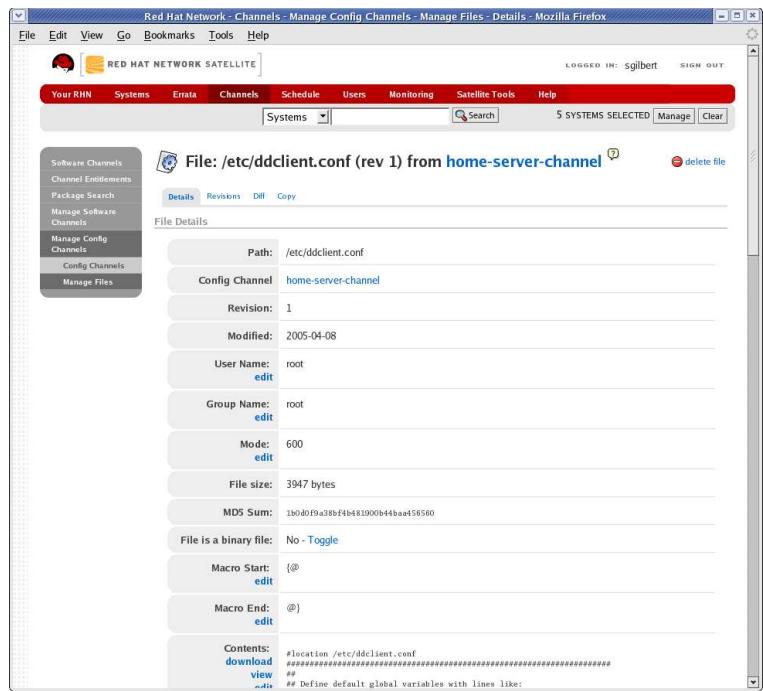


Figure 10. Provisioning Configuration Files

- *Failed actions notifications*—When the time arrives for the scheduled action(s), RHN executes the pre-arranged tasks, including full dependency checking against the packages being applied to your systems. In the event that any system cannot be properly managed or there is an error in the process, RHN issues a scheduled failed action report to the administrator.

Rollback and Recovery

To fully appreciate how RHN provides rollback and recovery, it is necessary to give an overview of how RHN records snapshots of your systems. When a system is given a Provisioning entitlement, RHN immediately begins storing snapshot profiles of that system whenever any change is made. Those changes can include updated packages, changed permissions, added or subtracted channels, new configuration files, or additions to new groups. Each change is recorded, time stamped, and stored in the central database.

By storing these snapshot profiles as shown in Figure 11, a user can compare two such profiles against each other and direct the machine to take on a desired state. In the event of a rollback, the

user is comparing the current state of the machine against a previous state and making the necessary changes. When cloning another image or machine, the user compares the machine that will be changed against the desired machine or image and asks the machine to change to the desired state.

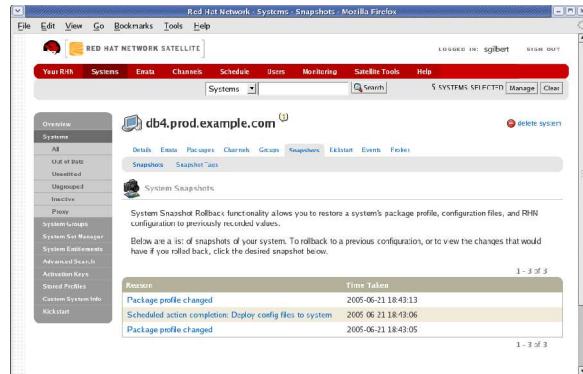


Figure 11. System Snapshots

Red Hat Network chose to follow this method of performing rollback for two reasons. First, in cases where multiple states had to be rolled back, the snapshot method (or “multi state”) is far more effective than rolling back many iterative states through RPM. Rolling back through many states risks mismanaging configurations or other errors repeatedly. Second, this method allows users to easily recover their systems in the event that there is a client failure. Since the snapshot is stored centrally in the database, it is a simple matter of bringing a system back up and then pointing it to the image of the latest “known good” state.

It should be noted that when a system is rolled back (or pointed to another image to clone), it is not just package changes that occur. These changes include permissions, group or channel changes, configuration files, and other actions that may have been made against the system. There may be cases when it is cleaner to reinstall the OS completely. In these cases, RHN uses logic to discern if re-provisioning the system is a better process than making all the individual changes.

Re-provision your Systems

Red Hat Network can be an important tool for enabling a flexible infrastructure. Many times, systems are deployed for a particular function, such as an application server or web server. Business needs, such as changing projects or customer demand, may require that the system be re-purposed for a different use. Rather than manually adjusting the packages, applications, and configurations on the system, Red Hat Network allows you to choose a profile and activation key for the type of system needed and RHN will quickly handle re-provisioning the system to meet the new requirement.

For some customers, it may make sense to wipe systems clean and reinstall the image to make sure that no other files or changes have been applied outside the RHN application. In this event, RHN stores the state of your system, and then re-provisions the system accordingly. Settings are then restored, packages and configuration files reapplied, and data reloaded.

Note that re-provisioning is only for servers that do not serve as data storage repositories. If your server is used as a storage device, the data must either be backed up to other servers and/or put into RPM format and loaded into an RHN channel.

Monitoring your Systems

Before you bring your systems online into production environments, you should include a monitoring solution to ensure availability and performance. Once you have identified systems that need monitoring functionality, you can use Red Hat Network Monitoring Module.

- *Create probes*—You can start by creating monitoring probes for your systems. Probes can monitor the system, network functionality, or applications. Dozens of pre-built probes are available, including many for applications from Oracle, MySQL, Apache, and BEA. You can also create custom probes using the tools included with Red Hat Network.

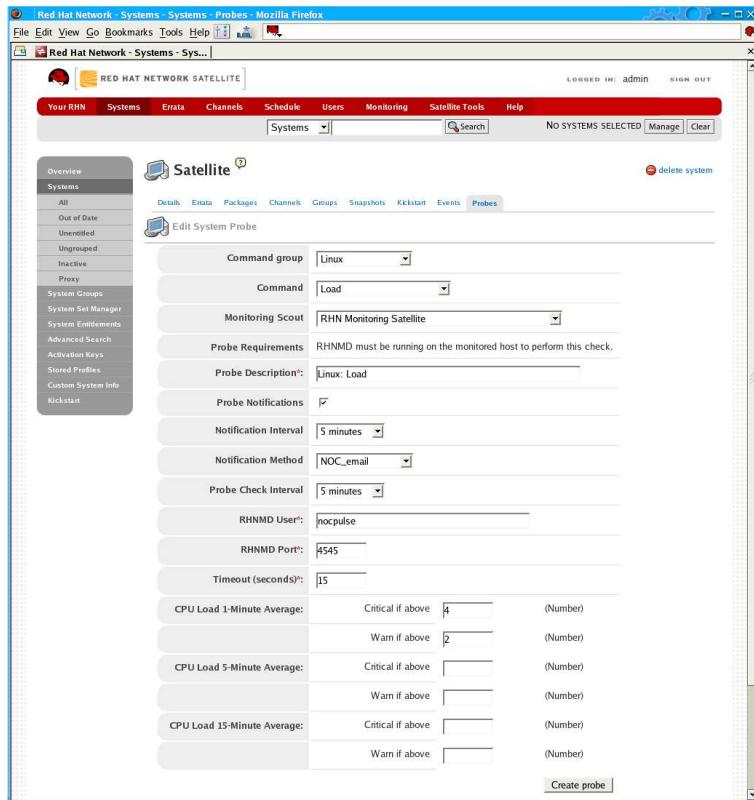


Figure 12. Create and Configure Probes

- *Configure probes*—Each probe can be configured for warning and critical performance thresholds. When these thresholds are reached, email or pager notifications can be sent to people identified in the probe.
- *Define and deploy probe suites*—You will often deploy the same probes across your systems, particularly across systems of a similar type. With Red Hat Network, you can create groups, or suites, of fully configured probes. These suites are then deployed to a given system, or group of systems, all at once rather than adding them one at a time. Once deployed, you can globally adjust the probe configurations for all systems that received the suite. You can also de-couple deployed probes from a probe suite if system-level tweaks are needed.

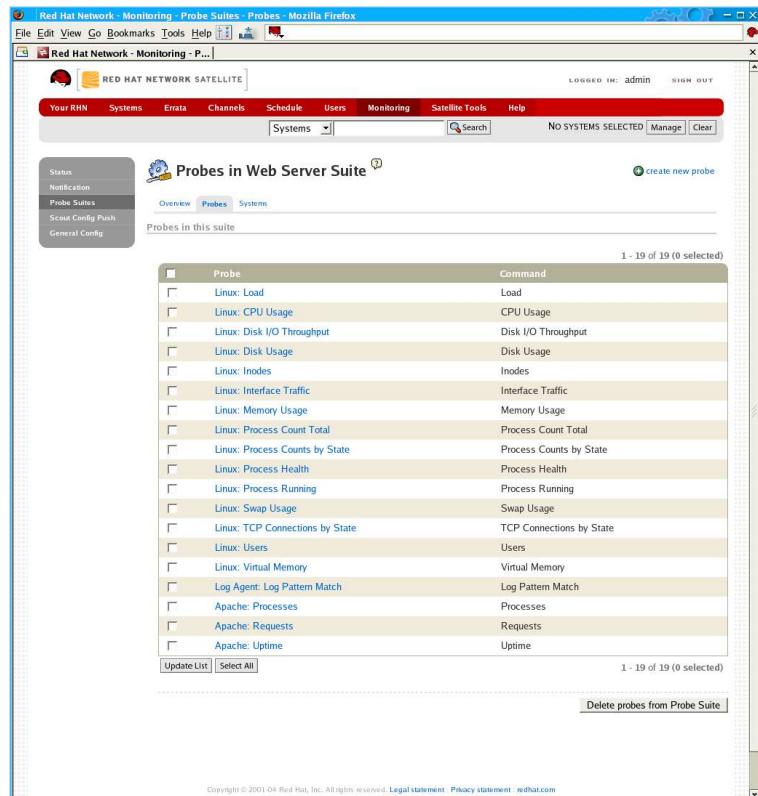


Figure 13. Probe Suites

- *View reports*—Beyond receiving alerts when systems exceed thresholds, it may be helpful to look at performance over a period of time. With Red Hat Network, you can create graphs of a probe's performance over any period of time, as well as view the raw data from the information collected by the Monitoring Module.

Useful Links

Red Hat Network homepage
rhn.redhat.com

Red Hat Network product information
www.redhat.com/software/rhn/

Worldwide contact information
www.redhat.com/about/corporate/woffices/

APPENDIX 1: Getting More from RHN—The API Access Layer

Red Hat knows that customers value choice, and we also know that every customer's systems management needs will be a little different depending on the exact solution that is being designed. To accommodate this, Red Hat Network is built with flexibility in mind and features a full set of APIs to ensure easy integration with your environment. Some potential additional functionality available in conjunction with Red Hat Network APIs are:

Automation—Create scripts that let you perform actions more quickly than navigating the RHN GUI.

Third party integration—Integrate actions from RHN with other third party tools to provide a more robust solution.

Custom application integration—Use RHN as a compliment to your existing processes and solutions. RHN does not require that you replace your existing environment or that you use only our products in your solution.

For more information regarding the APIs currently available, go to <https://rhn.redhat.com/rpc/api/>.

Lastly, Red Hat takes recommendations on new calls to the API layer with each release. If you are interested in learning more about this functionality and/or have a recommendation for Red Hat Network, consult your salesperson or sales engineer.

APPENDIX 2:

Running RHN in Highly Secure Environments

Keeping your Satellite Server connected to the central RHN servers via the Internet provides your company with an immediate and automated stream of Red Hat content. For increased security, however, some customers opt to run in a disconnected or completely off-network mode. Typically, Red Hat sees two kinds of deployments:

1. *Disconnected*—In the disconnected mode, customers sync their Satellite Servers to the Internet only for predefined time periods and only long enough to pull down the necessary content changes from the central RHN servers. Essentially, this process is the same as being always connected except that you will only receive updates at pre-determined times.
2. *Off Network*—Satellite Server offers the capability to take an infrastructure completely off the network. Customers can pull down ISOs and packages in one of two ways:
 - Have physical media shipped to them.
 - Pull down packages/ISOs to a connected system, and apply these packages to physical media. Packages can then be installed from the physical media onto the Satellite Server directly.

To understand more about how you can use Red Hat Network in highly secure environments, talk with your salesperson or sales engineer.

APPENDIX 3:

Key Glossary Terms

Definitions of features and functionality available in Red Hat Network and referenced in this whitepaper

Activation Keys—A unique RHN generated key that can be used by an administrator to register a system to RHN, entitle the system, subscribe the system to selected channels, and then assign the system to predetermined groups and permissions. This process saves time and allows new systems to be deployed into production immediately.

API Access Layer—Application Program Interface. RHN provides an API layer that allows users to easily interact and integrate with RHN to allow RHN to fit into their existing environments as well as augment the functionality of RHN.

Arbitrary Actions—Actions that can be executed against specific systems in conjunction with other actions executed by RHN. For example, an administrator may wish to execute a reboot (or other) command on a specific system after updating that system. Arbitrary actions allows the administrator to schedule that command to take place after the update occurs.

Auto Update—The ability to have a system automatically update itself of any new packages that have been added to the channels to which that system is subscribed.

Bare Metal Provisioning (w/ PXE)—The ability to provision a system without a pre-installed operating system by using PXE Boot in conjunction with RHN Satellite Server. This functionality allows for rapid deployment of new servers into your production environment.

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.

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(s).

Channel Cloning and Management—The ability to clone a channel and manage the deployment of channels into your environment. Channel cloning and management is used to duplicate channels or to set up staged environments.

Channel Permissions—The ability within RHN to assign permissions to different users for access to different channels. This feature ensures that only users who have been granted access can manage defined channels.

Configuration Channels—Channels that function as repositories for configuration files.

Configuration Management—The ability for RHN to make remote changes to configuration files through the RHN interface. RHN can make these configuration changes to any text-based file in the managed system's file space.

Custom Channels—Channels that function as repositories for custom content. As long as content is properly packaged in the RPM format, administrators can use RHN to manage custom content throughout their environment.

Delta-based Actions—The functionality is used by RHN when performing system cloning and/or rollback. RHN looks at the differences (or deltas) between the current and desired state of the system and then executes the necessary actions to provision your system to the desired state.

Dependency Checking—The process undertaken by RHN through RPM to ensure that when an update is applied to a system, the system has all the necessary dependency packages needed to make the update. If the system does not have all of the necessary dependent packages (or if they are of earlier versions not yet updated), RHN will include those (updated) packages in the update.

Disconnected Satellite—See “Off Network Capability”

Email Notification—An alert sent by RHN via email about new errata, failed actions, or other requests set by the user regarding the state of their system(s).

Errata—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 Cloning and Management—The ability to clone and manage (create and make changes to) the deployment of errata into your environment. Errata cloning and management is used to create custom errata or to pass errata through staged environments.

Kickstart—Kickstart is a method of automating the installation of Red Hat Enterprise Linux onto a computer. This is accomplished by creating a file (ks.cfg) that contains responses to all the questions that would be asked by the installation program during interactive installation. RHN uses Kickstart to provision systems.

Local Package Caching—The process of storing or caching content locally on a Proxy or Satellite Server for faster downloads and easier distribution.

Off-Network Capability—The ability for RHN Satellite Server to run in a completely disconnected or off-network environment, thereby ensuring the highest level of security. To sync the Satellite Server, it is necessary to physically provide media to the Satellite Server. This media can either be downloaded and created by a system that does have access to the Internet or by media sent to the customer from Red Hat.

Org Admin—Organization Administrators 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 Organization Administrator can also give users administrative privileges to system groups. An RHN organization must have at least one member of the Organization Administrator group.

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.

Package Profile Comparison—The ability for RHN to compare sets of packages between two systems or between a system and an existing image. This allows the user to audit the packages of a system or compare against another system. This functionality is also used by RHN to determine necessary changes to a system when provisioning that system.

Provisioning—The act of providing a system with all of the necessary components to effectively deploy that server. RHN uses Kickstart functionality and activation keys to provision a system with all necessary components: permissions, channel subscriptions, groups, operating system, additional packages, configuration changes, and any arbitrary actions defined by the user.

RPM (Red Hat Package Manager)—A software package manager that was developed by Red Hat. It can be used to build, install, query, verify, update, and un-install software packages. All software updates from RHN are delivered in RPM format.

RPM-based Application Provisioning—The ability for RHN to provision applications that are packaged in the proper RPM format. This is done by putting the application packages in a custom content channel and then using RHN to distribute to the different systems.

Scheduled Actions—The ability to schedule an action to occur within a predefined interval. Scheduled actions can be used to affect changes in a determined sequence of events or to select a distinct time period for an action to occur.

State Image Snapshot—RHN records snapshots of your system whenever there is a change in the state. These snapshots are then stored to create a profile of your system that can be used to rollback your system or in the event of disaster recovery.

System Cloning—The ability to clone a system through the use of RHN provisioning functionality.

System Grouping—The ability for RHN to group multiple individual systems together so that they can be managed as a single entity. This allows administrators to effectively manage an entire group of systems as easily as they could manage a single system.

System Profile—Hardware and software information about the client system. It is created during the registration process and regularly updated by RHN. 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 relevant to each client system.

System Permissions—The ability within RHN to assign permissions to different users for access to different systems. This feature ensures that only users who have been granted access can manage defined systems.

System Search—The ability to search through managed systems. RHN allows you to search by packages, hardware characteristics, DMI information, network information, defined asset tags, and more.

System Set—Interface that allows users to create temporary groups and perform actions on multiple systems. Actions include applying Errata Updates, upgrading packages, and adding/removing systems to/from system groups.