

A horizontal decorative bar consisting of a solid red segment on the left, followed by four segments of abstract, colorful patterns in blue, orange, green, and purple.

Migrating to Red Hat Enterprise Linux 4: Upgrading to the latest Red Hat release

By Donald Fischer

Abstract

Red Hat Enterprise Linux subscribers may choose to deploy any of the supported versions of the product family, and may upgrade to a new major release on their own schedule without incurring additional subscription cost. This whitepaper provides guidance on how to evaluate, plan, and execute a migration to Red Hat Enterprise Linux 4 from a prior Red Hat Enterprise Linux release or legacy Red Hat Linux distribution.

February 2005



Table of contents

Why migrate to Red Hat Enterprise Linux 4?.....	3
Types of migrations	3
Administrator-guided migration from previous Red Hat release.....	3
Automated migration from previous Red Hat release.....	4
Migration from another platform.....	4
Choosing an upgrade strategy.....	4
Assessing your current deployment.....	5
Evaluate RPM-managed files.....	5
Evaluate non-RPM managed files.....	6
Custom drivers and kernel modules.....	7
Planning your migration.....	7
Step 1: Define a repeatable customization change set.....	7
Step 2: Educate user community in advance.....	8
Step 3: Implement a controlled test migration.....	8
Step 4: Choose migration timing.....	8
Implementing your migration.....	9
Back up current systems.....	9
Roll out migration.....	9
Capture issues for future planning.....	9
Additional references.....	9



Why migrate to Red Hat Enterprise Linux 4?

There are a number of advantages for migrating from older Red Hat releases or other platforms to Red Hat Enterprise Linux 4, the latest release of Red Hat's enterprise computing platform. Among the leading reasons to migrate:

- *Latest and greatest features:* The open source development community continues to innovate at a breakneck pace, and Red Hat Enterprise Linux 4 incorporates the latest mature open source technologies for use in commercial environments. For example, Red Hat Enterprise Linux 4 includes:
 - Security Enhanced Linux (SELinux), a fine-grained access control mechanism that dramatically improves system security
 - State-of-the-art desktop environment based on GNOME 2.8
 - New storage management and virtualization technologies based on Logical Volume Manager 2 (LVM2)
- *Latest hardware support:* Red Hat Enterprise Linux 4 supports the latest server, workstation, and client hardware platforms and devices.
- *Application compatibility:* Users of previous Red Hat Enterprise Linux releases can be confident that their standards-conforming applications will continue to run on Red Hat Enterprise Linux 4, which includes compatibility runtime environments for applications designed for both Red Hat Enterprise Linux v.3 and Red Hat Enterprise Linux v.2.1. More detailed information on application compatibility is available in the Red Hat whitepaper *Red Hat Enterprise Linux 4 Application Compatibility*.
- *Security and maintenance lifetime:* Red Hat provides support and maintenance for Red Hat Enterprise Linux release for seven years from the major version release date. For example, Red Hat Enterprise Linux 4, released in February 2005, will have support and maintenance through at least February 2012. Customers who have deployed on earlier Red Hat Enterprise Linux versions should consider migrating to the latest release to maximize the support and maintenance lifetime of their deployments. Customers who still have legacy Red Hat Linux deployments should consider migrating to Red Hat Enterprise Linux to ensure a supported platform with guaranteed maintenance and security patch availability.

For more information on features in Red Hat Enterprise Linux 4, see *Red Hat Enterprise Linux 4: the defining milestone in the evolution of the enterprise platform* available at www.redhat.com/software/rhel.

Types of migrations

Administrator-guided migration from previous Red Hat release

The most common type of migration from earlier Red Hat releases to the latest Red Hat Enterprise Linux release is an administrator-guided migration. In this type of migration, the systems administrator evaluates the currently deployed system, backs up the deployed configuration data, user data, and software applications to an external system, performs a fresh install of Red Hat Enterprise Linux 4, and finally restores the required configuration data, user



data, and software applications. Red Hat generally recommends administrator-guided migrations for commercial deployments, as this method provides the highest assurances of a successful migration.

Automated migration from previous Red Hat release

Red Hat Enterprise Linux also includes support for an automated migration from earlier Red Hat Enterprise Linux releases. Automated migration is implemented by Anaconda, the Red Hat Enterprise Linux 4 installer. Red Hat does not recommend automated migration, as it provides a lower assurance of successful migration than administrator-guided migrations. Red Hat technical support will only provide assistance with automated migration from the previous major release of Red Hat Enterprise Linux. The Red Hat Enterprise Linux installer will warn the user when upgrading from any release earlier than Red Hat Enterprise Linux 3, but still allow the user to proceed at their discretion.

Migration from another platform

In addition to migrations from earlier Red Hat releases, many administrators seek to migrate from other operating systems, including UNIX, Microsoft Windows, mainframe operating environments, or other Linux distributions. Migrations from non-Red Hat platforms vary in complexity and are not covered in detail in this whitepaper.

For more information on migrating from non-Linux platforms, see *Solaris 10 and Linux 2.6: Understanding Two Strategies for Enterprise Operating Systems* and *Unix to Linux Migration: An introduction* available at www.redhat.com/solutions/info/whitepapers.

Choosing an upgrade strategy

Red Hat generally recommends the administrator-guided migration approach for commercial deployments of Red Hat Enterprise Linux. Administrator-guided migrations offer a number of advantages over automated migrations:

- *Ensures known good state:* The administrator-guided migration approach ensures the system is in a known good state at the conclusion of the migration. Deployed systems follow a natural tendency towards entropy—even when managed by a single system administrator, systems tend to accumulate configuration changes and customizations that become obsolete. By installing a new system and individually scrutinizing changes made to the default installation, system administrators can ensure that the resultant system is in a known good state at the conclusion of the migration. This is often not the case for the deployed systems before the migration.
- *Ensures a repeatable deployment:* Distilling the changes required to turn a fresh Red Hat Enterprise Linux installation into the desired configuration pays dividends beyond the one-time migration from an old release to a new release. For example, saving and/or automating the configuration process allows for quick provisioning of identically-configured systems in the case of hardware failure or increased load demand which requires more than one identically configured server. Red Hat Network, described below, provides



an ideal tool set for creating repeatable deployments.

- *Avoids incomplete migrations:* The Red Hat Enterprise Linux automated upgrade capability handles only the upgrade of system components distributed with the Red Hat base operating system. Automated upgrades may impact non-standards-conforming third-party applications in unspecified ways, some of which may not be found until application runtime. By establishing a clean baseline and then re-installing third-party applications using their native installation tool on a fresh system, application install-time dependency checking can be employed to ensure a complete runtime environment is available for all third-party software.

If these concerns do not apply to your deployment, you may still want to consider an automated upgrade installation using the Red Hat Enterprise Linux installer. Details on how to perform an automated upgrade installation are provided in Appendix A of the *Red Hat Enterprise Linux Installation Guide*, available on the Red Hat Enterprise Linux documentation CD or on the web at www.redhat.com/docs.

Because of the substantial changes in the open source distribution between major releases, Red Hat recommends against using the automated upgrade installation method to upgrade a system from Red Hat Enterprise Linux v.2.1 or Red Hat Linux (9 or earlier releases) to Red Hat Enterprise Linux 4. The administrator-guided migration method will deliver substantially better results in these cases.

The remainder of this document assumes the use of the recommended administrator-guided migration method.

Assessing your current deployment

The first step in implementing a migration to Red Hat Enterprise Linux 4 is conducting an analysis of your current deployed systems. For this section, we assume a migration from a recent release of Red Hat Enterprise Linux or legacy versions of Red Hat Linux and the use of the administrator-guided migration method.

Evaluate RPM-managed files

Red Hat Enterprise Linux systems store configuration information in a mix of static text files and area-specific data files. The majority of configuration items that system administrators typically need to capture for re-application in a new system installation are found in text files in the `/etc` portion of the file system hierarchy.

For packages managed with the RPM Package Manager, which includes all of the packages shipped with Red Hat Enterprise Linux and most other Red Hat products, there is an automated way to list configuration files that have changed. RPM provides a capability for package creators to specify which files contain configuration information, and also provides a way to determine which files have changed since installation by comparing checksums of each file's contents and various other attributes of the file such as ownership and permissions. Combining these two features, it is possible to generate a list of RPM-managed configuration files that have been modified from the default version placed on the file system when the RPM was installed.



Issuing the command 'rpm -Va' on a Red Hat system will generate a list of modified files for all packages on the system. Because this command verifies all files in all packages installed on the system, it will usually take a while to complete. It will also generate a large amount of output, so it is best to redirect the output to a file that you can then analyze (for example, use 'rpm -Va > /tmp/rpmva-output.txt' .)

The 'rpm -Va' command will produce output only in cases where the current file content or attributes do not match those from the originally installed RPM package. For each file that differs, a line will be printed in the following format:

```
SM5DLUGT c <filename>
```

Components of the output:

- S – the size of the file differs
- M – the file's mode differs
- 5 – the MD5 checksum of the file contents differs
- D – the file's major/minor numbers differ
- L – the file's symbolic link contents differ
- U – the file's owner differs
- G – the file's group differs
- T – the file's modification time differs
- c – appears only if the file was marked as a configuration file by the package creator
- filename – the name of the file as installed on the file system

By analyzing the output of the 'rpm -Va' command with particular focus on the configuration files indicated with a 'c' between the filename, administrators can quickly compile a list of changes and customizations that may need to be re-applied in a fresh installation to create a similarly-configured system.

Note that by design, many configuration files are expected to be modified after installation. For example, the system-config network tool provided with Red Hat Enterprise Linux provides a user interface for setting network configuration that will change various network configuration files. It is still useful for administrators to use RPM to generate a list of configuration files changed since installation because those changes may need to be re-applied to the fresh system either manually or by using the provided graphical configuration tools when migrating.

More detailed information on how to use RPM to verify installed files is available in Chapter 6 of Ed Bailey's *Maximum RPM*, available at www.rpm.org/max-rpm

Evaluate non-RPM managed files

In addition to RPM-managed files, deployed systems always contain a set of files that are not managed by RPM. These files include files added to the system ad-hoc by system administrators, third-party software packages that not packaged with RPM, system and user data files, and additional system and user configuration files created at runtime rather than at the time of



installation.

Non-RPM managed files that may need to be migrated to a new system installation are most likely to be found in the following parts of the file system hierarchy:

- /etc – Host-specific system configuration files
- /home – User home directories (data files and configuration files specific to individual users)
- /root – Home directory for the root user (system administrator account)
- /opt – Third-party add-on application software packages
- /usr/local – Locally installed software packages
- /srv – Data for services provided by the system

Because file system layout and administration conventions vary from site to site, it is important to inspect the full file system for locally installed files and directories that may need to be migrated to a new installation. It is always a good idea to do a full system backup of the existing file system before re-installing.

For improved migration support, application developers should consider packaging their software using RPM and should ensure that their application packaging conforms to Red Hat's file system conventions. For more information, refer to chapter 3 "File System Structure" in the the *Red Hat Enterprise Linux 4 Reference Guide*, available at www.redhat.com/docs and the Red Hat whitepaper *Red Hat Enterprise Linux 4 Application Compatibility*.

Custom drivers and kernel modules

Third-party or locally configured kernel modules and drivers require particular attention in planning a migration. Because the Linux kernel does not have a defined kernel module Application Binary Interface (ABI) that is preserved across releases, any third-party software that depends on kernel modules is likely to require at least recompilation to work with new major releases of Red Hat Enterprise Linux. In many cases, kernel module source code may need to be reworked to match changed implementations in the kernel itself. These changes do not apply to standards-compliant, user-level applications which are designed to conform to the Application Binary Interfaces. For more information on ABI, see the Red Hat whitepaper *Red Hat Enterprise Linux 4 Application Compatibility*.

Planning your migration

After assessing your current system deployments, the next step is to plan your migration and test your migrated configuration with your user community.

Step 1: Define a repeatable customization change set

Based on the information gathered in the assessment phase, the system administrator should be able to define a concrete set of configuration changes and third-party software installation that will be required in the new Red Hat Enterprise Linux 4 environment.



Ideally, the application of these changes should be automated. In addition to making the migration deployment quicker and more accurate, automating customization steps enables the replacement of deployed systems and addition of more identically configured systems in the future to meet additional capacity demand.

Red Hat Network's provisioning capabilities provide an ideal set of tools for specifying and deploying initial system configuration and post-installation configuration changes. The Red Hat Network provisioning and configuration management tools make use of operating system level building blocks such as the Anaconda installer's Kickstart functionality for automated system configuration. Red Hat Network provisioning also allows configuration file changes, customized for each system, to be deployed, tracked, and maintained using the RPM package manager framework. For more details on Red Hat Network provisioning capabilities, see *RHN Reference Guide*, available at rhn.redhat.com/help.

While Red Hat Network provisioning provides an excellent set of tools, other approaches can also be used to consistently apply changes to multiple platforms. These solutions may range from custom-developer scripts to third-party management tools.

Step 2: Educate user community in advance

An essential component of any platform migration is communication with the impacted user community. Users who will be impacted by a migration should be notified well in advance so that they can prepare for changes in their working environment. Ideally, the user community should be involved throughout the migration planning, test, and implementation process to ensure the resulting environment continues to meet their existing needs while adding new capabilities.

Step 3: Implement a controlled test migration

Once user needs have been identified and a repeatable change set is in place using Red Hat Network provisioning or other technologies, system administrators should conduct one or more test migrations that simulate the final environment. This can initially be accomplished by taking a snapshot copy of the software environment from a production machine and migrating the copy in a test environment repeatedly until a successful, repeatable migration is achieved. If possible, the next step should involve migrating a subset of the user community who are willing to participate in a test migration. This test community can help to flesh out runtime migration issues that may not be obvious in a controlled environment.

Step 4: Choose migration timing

The final step in planning a migration is to choose the timeline for the migration. Depending on the operating requirements of the particular environment and user community, it may be desirable to migrate one part of the user base at a time in a "rolling upgrade," or to migrate the entire environment at once in a "big bang," generally associated with a published system downtime.



Implementing your migration

Once the migration has been planned and extensively tested, the actual rollout of the migration should be straightforward.

Back up current systems

A critical step in the migration process is to do a complete backup of all existing systems. At all times, there should be a plan in place that allows data to be recovered from pre-migration system images, and as a worst case scenario, a plan to roll back to pre-migration system images all together.

System and user data and third-party applications that need to be carried forward to the new environment should be copied to an off-system location.

Roll out migration

With repeatable change sets defined and tested and system and user data copied from the target system to an off-system location, the rollout consists of repeating the following steps for each target:

- Take target system out of service
- Install Red Hat Enterprise Linux 4 in desired configuration (using standard system installer or automated Red Hat Network Provisioning or Kickstart installation tools)
- Apply change sets to local system configuration (using Red Hat Network Provisioning or other configuration management tools)
- Restore system and user data from off-system location
- Restore target system to service

Capture issues for future planning

Even with extensive planning and testing, some issues will typically arise in the final migration process. Be sure to capture as completely as possible the administrator issues that arise during the migration and user community feedback during and after the migration. This information may then be fed back into the planning process for future platform migrations.

Additional references

Additional information on migrations and upgrades is available from the following sources:

- System upgrades are described in Appendix A of the *Red Hat Enterprise Linux Installation Guide*, available at www.redhat.com/docs
- Red Hat Enterprise Linux configuration tools, configuration file locations, and file system structure are described in the *Red Hat Enterprise Linux Reference Guide*, available at www.redhat.com/docs
- Red Hat Network provisioning and configuration management capabilities are described in detail in the *Red Hat Network Reference Guide*, available at rhn.redhat.com/help/
- Application compatibility guidelines for Red Hat Enterprise Linux, including



cross-major release compatibility, are described in the whitepaper *Red Hat Enterprise Linux 4 Application Compatibility*, available at www.redhat.com

- For more information on the RPM package manager, including details on RPM's verification capabilities and how RPM handles package upgrades, consult Ed Bailey's book *Maximum RPM*, available at www.rpm.org/max-rpm

For more information on Red Hat, visit www.redhat.com or contact us at 1-888-REDHAT1(US and Canada) / +1-919-754-3700 (international).