



# Red Hat

RH133

Enterprise Linux  
System Administration

RH133-RHEL4-1-20050308

Red Hat Europe, 10 Alan Turing Road,  
Guildford, Surrey, GU2 7YF,  
United Kingdom

Tel: + (44) 1483 300169

FAX: + (44) 1483 574944

Handwritten text at the top of the page, possibly a title or header.

A horizontal line of text, possibly a separator or a specific entry.

Handwritten text at the bottom of the page, possibly a signature or footer.

**RH133**

**Red Hat Enterprise Linux System Administration**

RH133-RHEL4-1-20050308



1. The first part of the document discusses the importance of maintaining accurate records of all transactions and activities.

2. It is essential to ensure that all data is entered correctly and consistently to avoid any discrepancies or errors.

3. Regular audits and reviews should be conducted to verify the accuracy and integrity of the information.

4. The final section provides a summary of the key findings and recommendations for future improvements.

# Table of Contents

## RH133 Red Hat Enterprise Linux System Administration

### RH133 UNIT 1 — Installation

Objectives	1-2
Agenda	1-3
Initial Installation	1-4
Hardware Overview	1-5
CPU and Memory	1-6
Preparing to Install	1-7
Multiboot systems	1-8
Device Node Examples	1-9
The RHEL Installer	1-10
Installer Features	1-11
RHEL Installation Overview	1-12
Partitioning Hard Drives	1-13
Sample Partition Structure	1-14
Configuring File Systems	1-15
Software RAID	1-16
LVM: Logical Volume Manager	1-17
Network Configuration	1-18
Firewall Setup	1-19
Security Enhanced Linux	1-20
SELinux Installation Options and Control	1-21
Package Selection	1-22
Validating the Installation	1-23
noprobe Mode and Driver Disks	1-24
Post-Install Configuration	1-25
End of Unit 1	1-26
<b>Lab: Installation</b>	

### RH133 UNIT 2 — System Initialization and Services

Objectives	2-2
Agenda	2-3
Boot Sequence Overview	2-4
BIOS Initialization	2-5
Boot Loader Components	2-6
GRUB and grub.conf	2-7
Starting the Boot Process: GRUB	2-8
Kernel Initialization	2-9
init Initialization	2-10
Run levels	2-11
/etc/rc.d/rc.sysinit	2-12
/etc/rc.d/rc	2-13
Daemon Processes	2-14
System V run levels	2-15
/etc/rc.d/rc.local	2-16
Virtual Consoles	2-17
Controlling Services	2-18
System Shutdown	2-19
System Reboot	2-20
End of Unit 2	2-21
<b>Lab: Managing Startup</b>	

**RH133 UNIT 3 — Kernel Services and Configuration**

Objectives	3-2
Agenda	3-3
Kernel Modules	3-4
Kernel Module Configuration	3-5
The <code>/proc</code> filesystem	3-6
The <code>/proc</code> filesystem (continued)	3-7
<code>/proc/sys</code> configuration with <code>sysctl</code>	3-8
General Hardware Resources	3-9
System Bus Support	3-10
Hotswappable Bus Support	3-11
System Monitoring and Process Control	3-12
End of Unit 3	3-13
<b>Lab: Configuring kernel parameters</b>	

**RH133 UNIT 4 — Filesystem Management**

Objectives	4-2
Agenda	4-3
System Initialization: Device Recognition	4-4
Disk Partitioning	4-5
Managing Partitions	4-6
Managing Data: Filesystem Creation	4-7
Journaling for ext2 filesystems: ext3	4-8
Managing Data: <code>mount</code>	4-9
Managing Data: <code>mount</code> options	4-10
Managing Data: Unmounting Filesystems	4-11
Managing Data: File System Labels	4-12
Managing Data: <code>mount</code> , by example	4-13
Managing Data: Connecting Network Resources	4-14
Managing Data: <code>/etc/fstab</code>	4-15
Managing Data: The Auto-Mounter	4-16
ext2/ext3 Filesystem Attributes	4-17
Virtual Memory Files	4-18
Filesystem Maintenance	4-19
Filesystem Maintenance (cont.)	4-20
Adding a Drive	4-21
End of Unit 2	4-22
<b>Lab: Filesystem Management</b>	

**RH133 UNIT 5 — Network Configuration**

Objectives	5-2
Agenda	5-3
Device Recognition	5-4
Network Interfaces	5-5
<code>mii-tool</code>	5-6
<code>ifconfig</code>	5-7
<code>ifup/ifdown</code>	5-8
Interface Configuration Files	5-9
Configuration Utilities	5-10
Binding Multiple IP Addresses	5-11
DHCP/BOOTP	5-12
Global Network Parameters	5-13
Default Route	5-14

Static Routes	5-15
Name Resolution	5-16
DNS Client Configuration	5-17
DNS Utilities	5-18
Network Diagnostics	5-19
End of Unit 5	5-20
<b>Lab: Static Network Settings</b>	

**RH133 UNIT 6 — RPM and Kickstart**

Objectives	6-2
Agenda	6-3
The RPM Way	6-4
RPM Package Manager	6-5
Installing and Removing Software	6-6
Updating a Kernel RPM	6-7
RPM Queries	6-8
RPM Verification	6-9
Other RPM Utilities and Features	6-10
Automatic Dependency Resolution	6-11
Red Hat Network (RHN)	6-12
RHN in the Enterprise	6-13
RHN Registration	6-14
The <code>up2date</code> utility	6-15
Remote Administration	6-16
Network Installation Server	6-17
Using Kickstart to Automate Installation	6-18
Kickstart: Commands section	6-19
Kickstart: <code>%packages</code>	6-20
Kickstart: <code>%pre</code> , <code>%post</code>	6-21
End of Unit 7	6-22
<b>Lab: RPM and Kickstart</b>	

**RH133 UNIT 7 — User Administration**

Objectives	7-2
Agenda	7-3
User Policy Considerations	7-4
The User Account Database - <code>/etc/passwd</code>	7-5
Adding a New User Account	7-6
User Private Groups	7-7
Group Administration	7-8
Modifying/Deleting Accounts	7-9
Password Aging Policies	7-10
Login Shell Scripts	7-11
Non Login Shell Scripts	7-12
Switching Accounts	7-13
<code>sudo</code>	7-14
Network Users	7-15
Authentication Configuration	7-16
Example: NIS Configuration	7-17
Example: LDAP Configuration	7-18
File Ownership	7-19
Linux File Permissions	7-20
SUID / SGID Executables	7-21
The Sticky Bit	7-22
The <code>setgid</code> Access Mode	7-23

Default File Permissions	7-24
Access Control Lists (ACLs)	7-25
SELinux	7-26
Controlling SELinux	7-27
SELinux Contexts	7-28
Troubleshooting SELinux	7-29
End of Unit 7	7-30
<b>Lab: User and Group Administration</b>	

**RH133 UNIT 8 — Printing and Administration Tools**

Objectives	8-2
Agenda	8-3
CUPS Overview	8-4
CUPS Configuration Files	8-5
CUPS Queue Management	8-6
cron	8-7
Controlling Access to cron	8-8
System crontab Files	8-9
System Cron Job: tmpwatch	8-10
System Cron Job: logrotate	8-11
System Cron Job: logwatch	8-12
System Logging	8-13
syslog Configuration	8-14
Tape Drives	8-15
Using tar/star	8-16
Using dump/restore	8-17
Using cpio	8-18
Remote Backups	8-19
Other Backup Software	8-20
End of Unit 8	8-21
<b>Lab: Printing and Admin Tools</b>	

**RH133 UNIT 9 — The X Window System**

Objectives	9-2
Agenda	9-3
XOrg: The X11 Server	9-4
XOrg Server Design	9-5
XOrg Server Configuration	9-6
XOrg Modularity	9-7
Server and Client Relationship	9-8
XOrg in runlevel 3	9-9
XOrg in runlevel 5	9-10
Configuration Utilities	9-11
Remote X Sessions	9-12
End of Unit 9	9-13
<b>Lab: The X Window System</b>	

**RH133 UNIT 10 — Advanced Filesystem Management**

Objectives	10-2
Agenda	10-3
Software RAID Configuration	10-4



Software RAID Recovery	10-5
Converting LVM1 to LVM2	10-6
Creating Logical Volumes	10-7
Resizing Logical Volumes	10-8
The Linux Quota System	10-9
The Linux Quota System (cont )	10-10
End of Unit 10	10-11
<b>Lab: Logical Volumes, RAID and Quotas</b>	
<b>RH133 UNIT 11 — Troubleshooting</b>	
Objectives	11-2
Agenda	11-3
Troubleshooting	11-4
Things to Check: X	11-5
Things to Check: Networking	11-6
Order of the Boot Process	11-7
Filesystem Corruption	11-8
Filesystem Recovery	11-9
Recovery Run-levels	11-10
Rescue Environment	11-11
Rescue Environment Utilities	11-12
Rescue Environment Details	11-13
End of Unit 11	11-14
<b>Lab: System Rescue and Troubleshooting</b>	



*[Faint, illegible text at the bottom of the page, possibly bleed-through from the reverse side.]*

# Welcome !

## RH133

### Red Hat Enterprise Linux System Administration

Rev RH133-RHEL4-0

Copyright © 2004 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 293 2664 or +1 (919) 754 3700.

# Welcome to RH133

Please let us know if you have any special needs while at our training facility.

Rev RH133-RHEL4-0

Copyright © 2004 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (319) 794 3700.

## Phone and network availability

Please only make calls during breaks. Your instructor will show you which phone to use.

Network access and analogue phone lines may be available; your instructor will provide information about these facilities

Please turn pagers to silent and cell phones off during class

## Restrooms

Your instructor will notify you of the location of these facilities.

## Lunch and breaks

Your instructor will notify you of the areas to which you have access for lunch and for breaks.

## In case of Emergency

Please let us know if anything comes up that will prevent you from attending

## Access

Each facility has its own opening and closing times. Your instructor will provide you with this information.

# Participant Introductions

- Please introduce yourself to the rest of the class.

Rev RH133-RHEL4-0

Copyright © 2004 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (916) 754 3700.

www.ck12.org

# Introduction

## RH133 Red Hat Enterprise Linux System Administration

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (913) 754 3700.

# Copyright

- The contents of this course and all its modules and related materials, including handouts to audience members, are Copyright © 2003 Red Hat, Inc.
- No part of this publication may be stored in a retrieval system, transmitted or reproduced in any way, including, but not limited to, photocopy, photograph, magnetic, electronic or other record, without the prior written permission of Red Hat, Inc.
- This curriculum contains proprietary information which is for the exclusive use of customers of Red Hat, Inc., and is not to be shared with personnel other than those in attendance at this course.
- This instructional program, including all material provided herein, is supplied without any guarantees from Red Hat, Inc. Red Hat, Inc. assumes no liability for damages or legal action arising from the use or misuse of contents or details contained herein.
- If you believe Red Hat training materials are being used, copied, or otherwise improperly distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll-free (USA) +1 866 626 2994 or +1 (919) 754 3700.

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat  2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.



# Red Hat Enterprise Linux

- Enterprise-targeted operating system
- Focused on mature open source technology
- 12 to 18 month release cycle
  - Certified with leading OEM and ISV products
- Purchased with one year Red Hat Network subscription and support contract
  - Support available for five years after release
  - Up to 24x7 coverage plans available

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2994 or +1 (919) 754 3700.

## About Red Hat Enterprise Linux

The Red Hat Enterprise Linux product family is designed specifically for organizations planning to use Linux in production settings. All products in the Red Hat Enterprise Linux family are built on the same software foundation, and maintain the highest level of ABI/API compatibility across releases and errata. Extensive support services are available: a one year support contract and Update Module entitlement to Red Hat Network are included with purchase. Various Service Level Agreements are available which may provide up to 24x7 coverage with guaranteed one hour response time. Support will be available for up to five years after a particular release.

Red Hat Enterprise Linux is released on a twelve to eighteen month cycle. It is based on code developed by the open source community and adds performance enhancements, intensive testing, and certification on products produced by top independent software and hardware vendors such as Dell, IBM, Fujitsu, BEA, and Oracle. The longer release cycle allows vendors and enterprise users to focus on a common, stable platform and to effectively plan migration and upgrade cycles. Red Hat Enterprise Linux provides a high degree of standardization through its support for seven processor architectures (Intel x86-compatible, Intel Itanium 2, AMD64, IBM PowerPC on eServer iSeries and eServer pSeries, and IBM mainframe on eServer zSeries and S/390).

*Red Hat Enterprise Linux AS:* the top-of-the-line Red Hat Enterprise Linux solution, this product supports the largest x86-compatible servers and is available with the highest levels of support.

*Red Hat Enterprise Linux ES:* for entry-level or mid-range departmental servers. Red Hat Enterprise Linux ES provides the same core capabilities as AS, for systems with up to two physical CPUs and up to 8 GB of main memory.

*Red Hat Enterprise Linux WS:* the desktop/client partner for Red Hat Enterprise Linux AS and Red Hat Enterprise Linux ES on x86-compatible systems. Based on the same development environment and same software core as the server products, Red Hat Enterprise Linux WS does not include some network server applications. It is ideal for desktop deployments or use as a compute node in a HPC cluster environment.

# Red Hat Network

- A comprehensive software delivery, system management, and monitoring framework
  - **Update Module**, included with Red Hat Enterprise Linux, provides software updates
  - **Management Module** adds more scalable management capabilities for large deployments
  - **Provisioning Module** provides bare metal installation, configuration management, and multi-state configuration rollback capabilities

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 260 2984 or +1 (910) 764 3700.

## About Red Hat Network

Red Hat Network is a complete systems management platform. It is a framework of modules for easy software updates, systems management, and monitoring, built on open standards. There are currently three modules in Red Hat Network; the Update Module, the Management Module, and the Provisioning Module.

The Update Module is included with all subscriptions to Red Hat Enterprise Linux. It allows for easy software updates to all your Red Hat Enterprise Linux systems.

The Management Module is an enhanced version of the Update Module, which adds additional functionality tailored for large organizations. These enhancements include system grouping and set management, multiple organizational administrators, and package profile comparison among others. In addition, with RHN Proxy Server or Satellite Server, local package caching and management capabilities become available.

The Provisioning Module provides mechanisms to provision and manage the configuration of Red Hat Enterprise Linux systems throughout their entire life cycle. It supports bare metal and existing state provisioning, storage and editing of Kickstart files in RHN, configuration file management and deployment, multi-state rollback and snapshot based recovery, and RPM-based application provisioning. If used with RHN Satellite Server, support is added for PXE boot bare-metal provisioning, an integrated network installation tree, and configuration management profiles.

# Red Hat Applications

- Optional layered products which enhance the standard Red Hat Enterprise Linux system
  - Red Hat Cluster Suite
  - Red Hat Content Management System
  - Red Hat Developer Suite
  - Red Hat Portal Server

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



5

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 255 2994 or +1 (919) 754 3100.

## Red Hat Applications

Red Hat provides a set of optional layered products that can be used to enhance the standard Red Hat Enterprise Linux operating system. Red Hat provides full maintenance and support services for these open source middleware and application layer products. Current offerings include:

*Red Hat Cluster Suite*: this product provides high availability clustering features. Both network load balancing clusters and two to eight node high availability application clusters are supported. Originally part of Red Hat Enterprise Linux AS, this product has been enhanced and is now available as a separate layered product for both AS and ES based systems.

*Red Hat Content Management System*: a complete workflow-based engine to manage content creation and delivery for an intranet, extranet, or Internet web site.

*Red Hat Developer Suite*: a fully featured Integrated Development Environment (IDE) for application developers based on the open source Eclipse project. Plugins for C/C++, Java, RPM, and profiling are included, and additional plugins will be provided as they become available.

*Red Hat Portal Server*: a servlet-based framework to aggregate local and remote content along with applications into an easy-to-configure web interface. Customizable templates allow the enterprise, a specific department, or the end user to provide information with the look-and-feel which is desired.

# The Fedora Project

- Red Hat-sponsored open source project
- Focused on latest open source technology
  - Rapid four to six month release cycle
  - Available as free download from the Internet
- An open, community-supported proving ground for technologies which may be used in upcoming enterprise products
  - Red Hat does not provide formal support

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## About the Fedora Project

The Fedora Project is a community supported open source project sponsored by Red Hat intended to provide a rapidly evolving, technology-driven Linux distribution with an open, highly scalable development and distribution model. It is designed to be an incubator and test bed for new technologies which may be used in later Red Hat enterprise products.

The basic Fedora Core distribution will be available for free download from the Internet.

The Fedora Project will produce releases on a short four to six month release cycle, to bring the latest innovations of open source technology to the community. This may make it attractive for power users and developers who want access to cutting-edge technology and can handle the risks of adopting rapidly changing new technology. Red Hat does not provide formal support services for the Fedora Project.

# Audience and Prerequisites

- Audience: Linux or UNIX users who understand the basics of Red Hat Enterprise Linux, that desire further technical training to begin the process of becoming a system administrator
- Prerequisites: Experience with Linux or UNIX desktop productivity and command-line tools

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2994 or +1 (919) 754 3700.

## Audience for RH133

The Red Hat Linux System Administration course is designed for users with Linux or UNIX experience who want to start building skills in system administration on Red Hat Linux, to a level of competence where they are able to configure and attach a workstation to an existing network

## Prerequisites for RH133 include knowledge in the following areas:

- File and directory operations
- Understanding users and groups
- Standard I/O and pipes
- String processing
- Managing processes
- Using the bash shell
- Using the Red Hat Linux graphical environment
- Sending e-mail and using printing
- Use of the vi text editor

# Classroom Network

	Names	IP Addresses
Our Network	example.com	192.168.0.0/24
Our Server	server1.example.com	192.168.0.254
Our Stations	stationX.example.com	192.168.0.X
Evil Outside Network	cracker.org	192.168.1.0/24
Evil Outside Server	server1.cracker.org	192.168.1.254
Evil Outside Stations	stationX.cracker.org	192.168.1.X
Trusted	trusted.cracker.org	192.168.1.21

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 296 2984 or +1 (915) 764 3700.

# Unit 1

## Installation

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

# UNIT 1: Objectives

- Upon completion of this unit you should be able to:
  - Use Red Hat resources to identify supported hardware
  - Describe how Linux accesses devices
  - Install Red Hat Enterprise Linux
  - Perform basic post-install configuration

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2984 or +1 (919) 754 3700.



# UNIT 1: Agenda

- Supported Linux hardware
- Linux and hardware access
- Installing Linux
- Post-install configuration

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 263 2994 or +1 (918) 754 3700.

# Initial Installation

- Please turn to Lab One Sequence One
  - Perform Sequence One
  - Complete Installation described in Sequence Two

Rev. RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

Please turn to the section following Unit One, Lab One Complete the BIOS setup indicated in Sequence One and perform an installation as per the instructions in Sequence Two After you have completed the installation, you will have a machine running Red Hat Enterprise Linux. This installation is minimal, but following Unit One, you will complete Sequence Three of the lab which is a more featured installation

# Hardware Overview

- Kernel Support
  - Core Support: CPU, Memory, Process Management, Interrupt/Exception Handling etc.
  - Dynamically Loadable Kernel Modules
    - Device Drivers
    - Additional Functionality
- User Mode Access to kernel facilities
  - System Calls and Signals
  - Filesystem Device Nodes
  - Network Interfaces

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



5

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2994 or +1 (919) 754 3700.

Mediating access to hardware is one of the primary roles of any operating system. The Linux kernel provides core facilities for accessing base components, such as the system's CPU, memory, console, and PCI bus. Usually, the detection and configuration of these components is automatic.

Support for peripheral devices is generally implemented through kernel device modules. The kernel must coordinate low level resources among the various drivers, such as interrupt lines (IRQ's), ioports, and more generalized iomapped memory and direct memory access. Most device drivers can either be statically compiled into the core kernel image, or implemented as a dynamically loaded kernel module.

Kernel modules may also offer additional functionality such as kernel level packet filtering, a type of firewall.

In Linux, as in Unix, kernel facilities are accessed using what is called 'user mode access'. Access to kernel functions such as file and process creation are done by making system calls. Signals allow communication between running processes and can be sent by a user to a process with the kernel acting as the messenger. Access to most devices is achieved through file system device nodes. Utilities can access devices in a uniform manner, without knowing the device driver's implementation details. In Unix, "everything is a file". The one notable exception to this rule is networking devices. They generally are not accessed through a device node but instead are accessed through a "network interface" abstraction. Keep in mind that the creation of the Unix operating system predates networking by a decade.

# CPU and Memory

- Seven Supported Architectures: x86, Itanium2, AMD64/EM64T, S/390, zSeries, iSeries, pSeries
- CPU support on x86
  - Technical support for more than 2 physical CPUs only on AS variant (may use Hyper-Threading)
  - Up to 32 physical CPUs with SMP or hugemem kernel
- Memory support on x86
  - Technical support for more than 16 GB on AS or WS
  - Standard i686/athlon kernel: 4 GB
  - SMP i686/athlon kernel: 16 GB
  - hugemem SMP kernel: 64 GB

Rev RH133 RHEL4.1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (319) 754 3700.

Red Hat Enterprise Linux is available for the Intel x86, Intel Itanium2, AMD64/EM64T, IBM eServer zSeries, IBM eServer iSeries, IBM eServer pSeries, and IBM S/390 architectures. This manual, associated course, and the RHCT and RHCE certifications cover the Intel x86 architecture only.

The official technical support provided by Red Hat will vary depending on the variant of Red Hat Enterprise Linux that you purchased. On the Intel x86 architecture, technical support for more than two physical CPUs is available only with the AS variant. The two physical CPUs may both use Hyper-Threading Technology, allowing more than two logical processors. The standard kernel supports one processor. Both the smp and the hugemem kernels support up to 32 processors (logical Hyper-Threaded processors do not count toward this number)

On the Intel x86 architecture, official technical support for more than 16 GB of RAM is available with the AS and WS variants. The standard uniprocessor kernel supports up to 4 GB of RAM. The smp kernel is similar to the bigmem kernel from RHEL 2.1. It includes PAE support and supports up to 16 GB of RAM. Due to limitations of the 32-bit architecture, a single process can only address 4 GB of that address space. Furthermore, with these kernels only 3 GB are available as per process user space to the program, as 1 GB is reserved for direct use by the kernel.

The new hugemem kernel supports up to 64 GB of RAM. In addition, almost the entire 4 GB address space is available to the program as user space. The kernel may also directly use a 4 GB memory space. However, this kernel will incur a small amount of additional overhead when switching from user space to kernel space.

## Preparing to Install

- Read the RELEASE-NOTES file on the first CD or at <http://www.redhat.com>
- Check Hardware Compatibility
  - Red Hat Supported Hardware List
    - Hardware certified by Red Hat
    - Hardware compatible with Red Hat Linux
  - XFree86 supported video cards

\*XOrg

Row RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2894 or +1 (919) 754 3700.

### Release Notes

The RELEASE-NOTES file contains valuable information concerning your release of Red Hat Enterprise Linux. In addition, it will contain changes that you should be aware of from previous releases. It is a valuable resource that should always be read prior to installing a new version of Red Hat Enterprise Linux.

### The Red Hat Hardware Compatibility List

The Red Hat Hardware Compatibility list, at <http://hardware.redhat.com/hcl>, contains information about hardware that has been tested by Red Hat. Hardware on this list should be easily supported, and support for these devices is included with standard Red Hat support plans.

### The XOrg Project

Red Hat Enterprise Linux 4 ships with the XOrg version 6.8 of the X Window System. The XOrg Project maintains a list of currently supported video cards at <http://xorg.freedesktop.org>. Often configuration information for the newest video cards can be found at this site.

# Multiboot Systems

- Red Hat Enterprise Linux and the GRUB bootloader can coexist with other operating systems, including the following:
  - Windows NT/2000/XP/2003
  - DOS, Windows 3 x/9x/ME
  - NetBSD, FreeBSD, and other open systems
- Two major issues arise when implementing multiboot systems:
  - Partitioning and the boot process

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2854 or +1 (319) 754 3700.

## Partitioning Issues

In order to run RHEL, it is necessary to create Linux swap and native partitions. It is usually advisable to install the other operating systems first. While RHEL will not try to delete other operating systems, other operating systems are not always so courteous. When installing the other operating system unpartitioned space must be left for the Linux partitions.

Sometimes another operating system already exists on a system and occupies all the available disk space. In these cases, there are two options:

- Back up the existing operating system(s) and files, repartition the drive leaving space for RHEL, then reinstall the existing operating system(s) from the backup
- Back up the existing operating system(s) and files, then use the third-party tool Partition Magic to resize the existing partitions to make space (the backup is a safety measure, and is not required for the repartitioning itself)

## Boot Process Issues

If a system will be booting multiple operating systems it will need a boot loader that is capable of booting multiple operating systems. Boot floppies are also an option, though not a particularly convenient one. In general, boot process configuration falls into one of two categories:

- GRUB is the primary boot loader and will launch Linux and other operating systems (or their boot loaders): use this approach with DOS, Windows 3 x, Windows 9x/ME, and Windows NT/2000/XP/2003
- A boot loader such as System Commander or NTLDR is already on the system and will launch GRUB as a secondary boot loader

# Device Node Examples

- Block Devices
  - `/dev/hda` - IDE drive
  - `/dev/sda` - SCSI Drive
  - `/dev/fd0` - floppy drive
- Character Devices
  - `/dev/tty[0-6]` - virtual consoles
  - `/dev/st0` - SCSI tape drive
- Symbolic links

Rev RH133 RHEL4-I

Copyright © 2005 Red Hat, Inc.



9

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2894 or +1 (915) 754 3700.

## Block Devices:

<code>hd[a-t]</code>	IDE drives
<code>sd[a-z]+</code>	SCSI drives
<code>fd[0-7]</code>	standard floppy drives
<code>md[0-31]</code>	software RAID metadisks
<code>loop[0-15]</code>	loopback devices
<code>ram[0-19]</code>	ramdisks

## Character Devices:

<code>tty[0-31]</code>	virtual consoles
<code>ttyS[0-9]+</code>	serial ports
<code>lp[0-3]</code>	parallel ports
<code>null</code>	infinite sink ( the bit bucket )
<code>zero</code>	infinite source of zeros
<code>[u]random</code>	sources of random information
<code>fb[0-31]</code>	framebuffer devices

## Symbolic Links:

<code>/dev/cdrom</code>	-->	<code>/dev/hd[a-t], /dev/sd[a-z]+</code>
<code>/dev/modem</code>	-->	<code>/dev/ttyS[0-9]+</code>
<code>/dev/pilot</code>	-->	<code>/dev/ttyS[0-9]+</code>

# The RHEL Installer

- First Stage Installer Images
  - `diskboot.img` - VFAT filesystem image for bootable media larger than a floppy
    - floppy installation is no longer supported
  - `boot.iso` - ISO9660 bootable CD image
  - `pxeboot` directory
- Second Stage Installer
  - graphical or textual
  - can be invoked in `noprobe` or Kickstart mode

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat 10

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## First Stage Installer

There are three different versions of the first stage installer available. Which method you choose depends on resources available to your system and your network configuration.

The first method, `boot.iso` is a ISO9660 filesystem for use if your system supports booting from a CD-ROM. You might choose this option when you do not wish to perform a CD based install, but you need to boot from a CD. Booting from `boot.iso` is the same as passing the `askmethod` argument to the installer when booting from CD 1. You can create a bootable CD using the `cdrecord` command. For instance:

```
cdrecord dev=/dev/hdc boot.iso
```

The second method, `diskboot.img`, is a VFAT filesystem designed to be used with USB pendrives, or similar media. This method requires that your BIOS support booting from a USB drive. You will need to use the `dd` command to move this image to your media. For instance:

```
dd < diskboot.img > /dev/sda
```

The third method, Pre-boot Execution Environment (PXE) provides for a diskless installation. Instructions for setting up a PXE environment are in the file `/usr/share/doc/syslinux-2.11/pxelinux.doc`. Further discussion of the PXE method is beyond the scope of this course.

## Second Stage Installer

The second stage installer, once located and loaded by the first stage, drives the remainder of the installation process.



## Installer Features

- noprobe and Kickstart modes available
- mediacheck tests media integrity
- Multiple Interfaces:
  - Graphical
    - Starts X server and a GUI installer
    - Works with hard drive, CDROM, NFS installation
    - Graphical is the default
  - Text
    - Menu-based terminal interface
    - Works with all installation methods

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2984 or +1 (919) 754 3700.

noprobe mode allows and requires complete control over all installation parameters. Kickstart mode permits automated installation.

The graphical interface makes installation easy and intuitive. The graphical interface can be started in *lowres* mode, which means it uses lower screen resolution settings for the installation.

The text based installer supports all installation methods, including FTP and HTTP. It is also useful when the installer has difficulty managing your display adapter. While this is uncommon, it can be particularly useful on laptops that have proprietary display adapters.

# RHEL Installation Overview

- Language, keyboard and mouse selection
- Media selection if applicable
- Disk partitioning
- Bootloader configuration
- Network and firewall configuration
- Authentication setup
- Package selection
- X server configuration

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2994 or +1 (916) 754 3700.

The second stage installer may either be the newer graphical installation program or the traditional text-based one. The first three installation steps will ask you to select the installation language, keyboard, and mouse type

# Partitioning Hard Drives

- Hard drives are divided into *partitions*
- Partitions normally contain file systems
  - *Primary, extended, and logical* partitions
  - The default filesystem type is *ext3*
  - Multiple partitions may be assembled into a larger virtual partition: software RAID and LVM
- Filesystems are accessed via a *mount point*, which is a designated directory in the file system hierarchy.

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 288 2994 or +1 (910) 754 3700.

## Disks and partitions

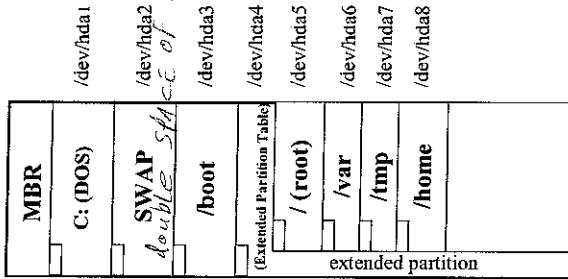
Disks are normally divided up into one or more *partitions*, each of which normally contains a file system, or swap space for virtual memory. This is useful because file systems on different partitions are independent from each other. If one file system fills up, other file systems on the disk may still have space available.

On the x86 architecture, there is a standard disk partitioning format which is used by most operating systems. The first four partitions on the disk are called *primary* partitions. If more than four partitions are needed, one of the primary partitions may be converted into a special *extended* partition that contains one or more *logical* partitions. A primary partition or a logical partition may contain a file system or swap space.

Multiple partitions may be assembled into a single virtual partition by using advanced techniques. Software RAID is used to provide redundancy, improve performance, or create partitions bigger than a single disk. Multiple partitions from different disks are assembled into a *RAID device*, a disk array which is treated like a normal partition. LVM is used to assign one or more partitions to a volume group, which can be used to create virtual partitions called *logical volumes*. These logical volumes are easier to resize than normal partitions, can have snapshots of their state taken at a particular point in time, and have other special features. (Be careful to note that a logical volume and a logical partition are two different things.)

The inverted tree of the file system hierarchy is divided into one or more file systems which are stored on devices. A *mount point* is a designated directory in the file system hierarchy that is used to access a particular file system. When a partition, RAID device, or logical volume is associated with ("mounted on") the mount point, files and directories on that device's file system are accessible under that directory.

# Sample Partition Structure



Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

# Configuring File Systems

- Must select mount points, partition sizes, and file system types in the installer
  - Can set up manually or automatically
- There are many layouts which may be used
  - / must include /etc, /lib, /bin, /sbin, /dev
  - Swap space is typically 2x physical RAM *data*
  - Typical mount points: /boot, /home, /usr, /var, /tmp, /usr/local, /opt *APPS*

*data base, web*  
*terminal, /etc, /var, etc*  
*drives, static, 3rd party*

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2894 or +1 (919) 754 3700.

## Configuring the file system hierarchy

At installation time, you must divide up your disks into partitions of various sizes and identify whether those partitions should be formatted with a file system (typically ext3), used as swap space, or used as a RAID or LVM partition. If the partition contains a file system, it must also be assigned a mount point. You can have the installer automatically make these decisions, or you can make them manually.

If you choose automatic configuration, you still have some input in the partitioning process. You can ask to review and modify the selections manually after the installer makes its decisions. You can select which drives to use for the installation. You can also indicate if the installer should delete all existing partitions, delete all Linux partitions from previous installations, or leave all existing partitions alone (using unallocated disk space for new partitions).

You have a great deal of freedom in how you may manually configure your file system hierarchy. You must have a file system mounted on /. You typically should have about twice your RAM in swap space. Any one swap partition should be no more than 2 GB in size. The /etc, /lib, /bin, /sbin, and /dev directories may not be on separate file systems; they must be part of the / file system or the system will not boot properly.

It's common to have a /boot file system about 100 MB in size at the front of the disk, to hold files needed by the BIOS at boot time (such as the kernel and parts of the boot loader). This helps to avoid problems with old BIOS code. One limitation on /boot is that most boot loaders expect it to be on a normal disk partition or RAID 1 device.

The /var directory holds files that change frequently. This includes log files, the mail spool, and temporary space for software updates from Red Hat Network. If /var is on a separate partition, it should probably be at least 1 GB in size.

Depending on how much software you choose to install, /usr will probably need between 350 MB and 5 GB of space. The /tmp directory should have a decent amount of free space for temporary files written by programs. The rest of / requires a few hundred megabytes. This does not take into account any space needed for personal user files or software not included with RHEL.

# Software RAID

- Redundant Array of Inexpensive Disks
  - Multiple partitions on different disks combined into one RAID device
  - Fault tolerance, larger disk size, performance
- Install-time RAID levels:
  - RAID 0: striping (no redundancy) *large disk*
  - RAID 1: mirroring *redundancy + speed for reading*
  - RAID 5: striping with distributed parity

Row RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 295 2884 or +1 (818) 754 3700.

## Using software RAID

RAID is an acronym for “Redundant Array of Inexpensive Disks”. With software RAID, the operating system combines multiple RAID partitions on different disks into a single RAID device (Linux also supports hardware RAID using special disk controllers or external storage devices. These devices usually look like normal disks or disk partitions to the installer.)

The installer allows you to set up software RAID devices. You first create RAID partitions by creating a partition normally with a file system type of “software RAID”. Then you click the “RAID” button, and create a RAID device from the RAID partitions. Like a normal partition, for the assembled RAID device you will need to select a mount point and a file system type, but you will also need to assign a RAID device name (such as `/dev/md0`) and select what RAID level to use. Each RAID level has different advantages and disadvantages.

RAID level 0 is called “striping”, and requires at least two RAID partitions. The resulting RAID device is a virtual partition the size of all the member RAID partitions added together. RAID 0 allows creation of file systems bigger than any one disk, and has high performance for reads and writes. However, it is not truly a redundant array; if any disk in the RAID device fails, the file system on the RAID device is destroyed.

RAID level 1 is called “mirroring”, and also requires two RAID partitions. The resulting RAID device is a virtual partition the size of the smallest of the member RAID partitions. All RAID partitions which are members of the RAID device contain identical data. If any disk in the RAID device fails, the RAID device continues to function without losing data. This is useful for fault tolerance, but is costly in terms of disk space. Performance for reads and writes is good.

RAID level 5 is called “striping with parity”, and requires at least three RAID partitions of the same size. Like RAID 0, this RAID level allows creation of file systems bigger than any one disk. However, additional parity data is also stored on the RAID device which can be used to preserve file system data even if a single disk in the RAID device fails. Therefore RAID 5 can survive single disk failures, but at the cost of some storage efficiency. Read performance is good, but write performance is slower due to the parity updates.

# LVM: Logical Volume Manager

- Manages storage on one or more partitions as virtual partitions, or *logical volumes*
  - Real partitions are *physical volumes* and are assigned to a *volume group* (a virtual disk)
  - Disk space in the volume group is divided into *extents* which are assigned to a logical volume
- Easy to resize logical volumes
  - Add a physical volume to the volume group and assign the new extents to the logical volume

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2094 or +1 (619) 754 3700.

## Introduction to the Logical Volume Manager

A logical volume manager may be used to create virtual partitions called *logical volumes* from one or more disk partitions or RAID devices. Each partition, or *physical volume*, is assigned to a virtual disk called a *volume group*. Multiple physical volumes may be assigned to the same volume group, and a volume group may be partitioned into multiple logical volumes.

Each volume group divides its pool of disk space into *extents* of identical size. The size of an extent is set for a particular volume group when that volume group is first created. An extent is typically between 1 MB and 64 MB in size. Extents may then be assigned to a new or existing logical volume in the volume group. Currently, a single logical volume may contain at most 65534 extents, so larger extents allow larger logical volumes.

The logical volume manager provides no redundancy by itself. If a single physical volume fails, any logical volume which is assigned extents from that volume will also fail.

LVM provides flexible disk management. For example, it is easier to resize logical volumes than it is to resize normal disk partitions. New physical volumes may be added to a volume group, or existing logical volumes can be reduced in size, providing additional extents. Those extents can then be assigned to any logical volume in the volume group. The file system on the logical volume being resized must also support resizing. The standard ext3 file system currently supports off-line resizing.

To create a logical volume in the installer, you first need to create a normal disk partition with a file system type of “physical volume (LVM)”. Then, click the “LVM” button to create a new volume group. Give the volume group a name and set the extent size. Then create the new logical volume, assigning a logical volume name, mount point, file system type, and size.

# Network Configuration

- Can configure each NIC independently
  - DHCP or static IP configuration
  - Determine if automatically activated on boot

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 288 2884 or +1 (919) 754 3700.

There are several options available when configuring network interface cards under RHEL, and each card can be configured individually. You may choose between manually assigning an IP address or having the system contact a DHCP server at startup for its network configuration. You may also select whether or not the interface should be automatically activated at boot time.



## Firewall Setup

- Installer can set up a kernel mode stateful packet filter
- Choice of two settings: "Enabled" and "No Firewall"
- "Trusted Devices" can bypass the firewall
- Can allow access to arbitrary services

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

The installer will now prompt you to select a default firewall configuration for your local machine. This enables you to block remote machines from accessing network services on your machine.

You will be presented with two choices for the firewall, "Enable firewall" and "No Firewall"

"Trusted Devices" allows you to select certain network interfaces as "trusted". All network traffic from a trusted device will bypass the firewall.

"Allow Incoming" allows you to let remote machines access particular services through the firewall. Some common services are listed. You may specify additional services to allow in the "Other:" dialog box. This box takes a series of port:protocol pairs separated by commas. For "port" you may use either the name from `/etc/services` or the port number; for example, both `imap:tcp` and `517:udp` are acceptable definitions.

Firewall rules are written to `/etc/sysconfig/iptables` and invoked at boot up by the `/etc/rc.d/init.d/iptables` script.

# Security Enhanced Linux

- Access control determines what actions processes can perform on what objects
  - Discretionary Access Control (traditional Linux)
    - Users control permissions on objects
  - Mandatory Access Control (SELinux)
    - System policy restricts permissions which can be granted

*stop after hacker has gained*

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 293 2594 or +1 (919) 754 3700.

Access control mechanisms enforce security restrictions. They control what rights processes have to access objects like files, directories, and network sockets. Linux traditionally leaves control over permissions to the owner of an object or to root.

SELinux introduces Mandatory Access Control to Linux. With MAC, the system administrator can create a mandatory policy that limits what access a particular process may be granted to an object. Processes run in a domain, and objects are assigned types. A particular domain may have access to an object limited or denied based on that object's SELinux type. In addition, normal access permissions still apply.

The mandatory policy even applies to processes running as the root user. If a process owned by root is not running in a domain that has access to a file of a particular SELinux type, SELinux may still deny access. Once a process starts running in a particular domain, the policy may restrict it from changing to a different domain that would have access to the file or other object. It is possible to carefully confine network services to limit the effects of a compromise, even if the service is running as root.

# SELinux Installation Options

- Installation options:
  - Disabled
  - Warn (Permissive)
  - Active (default) (Enforcing)

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

During installation, SELinux is automatically activated. This is done on the same screen as firewalling. There are three options to choose from:

- Disabled: This turns enforcing off, which means that labeling and domains are not set up. This is the most efficient way of running your machine, but less secure.
- Warn: This option sets up policies and logging, so you can monitor what is happening in the machine, without actually running SELinux. This enables the possibility of writing new rules for testing purposes.
- Active: SELinux is now enforced, but it will only affect certain daemons. When active is chosen, select demons will be confined by SELinux permissions.

To change between enforcing and permissive mode, you can do that either at boot time, at runtime or you can make it permanent:

- During boot, add "enforcing=1" to the kernel line to turn on, "enforcing=0" turns it off.
- At run time "setenforce 1" turns SELinux into enforcing mode, "setenforce 0" turns it permissive.
- To make it permanent either edit GRUB or `sysconfig /boot/grub/grub.conf`: edit the file and add "enforcing=1" to the kernel line to turn on, "enforcing=0" turns it off OR `/etc/sysconfig/selinux`: this file is well documented to help you choose the right option.

To fine tune your security settings, you can also use "system-config-securitylevel"

# Package Selection

- Package Selection
  - universally (“everything”)
  - by predefined components
    - defined in `RedHat/base/comps.xml`
  - Individually

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (918) 764 3700.

Decisions about what software to install are not crucial as packages can always be installed later using the *rpm* facility.

# Validating the Installation

- Virtual consoles during installation
- Post-boot validation
  - `dmesg` and `/var/log/dmesg`
  - `/var/log/messages`
  - `/root/install.log`
- GRUB drops to a prompt if there is a problem loading files

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 288 2994 or +1 (919) 754 3700.

Five virtual consoles are provided during the installation process. The consoles can be accessed using Alt-F1 through Alt-F5. In order to exit the graphical installer and view these consoles it is necessary to hold down both Ctrl and Alt

- Alt-F1: The installer program in text mode
- Alt-F2: A bash shell (second stage installer only)
- Alt-F3: A log of installer messages (*anaconda*)
- Alt-F4: A log of kernel messages
- Alt-F5: stdout from mke2fs and grub commands (*format HD*)
- Alt-F7: The installer program in graphical mode

After installation is complete and upon rebooting the system, there are several places that can be checked for installation and configuration information

## Default Log files

- `/var/log/dmesg`: contents of the kernel buffer at the end of `/etc/rc.d/rc.sysinit`
- `/var/log/messages`: output from the `syslogd` system logging daemon
- `/root/install.log`: logging information from the installer program

## noprobe Mode and Driver Disks

- Method for supporting hardware newer than the install program
- Used at install time for less common hardware
- Prompt for Driver Disk
  - When run in noprobe mode
  - When started with: linux dd
  - When no PCI devices are detected

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2894 or +1 (919) 754 3700.

### The Need for noprobe Mode

**noprobe** mode lets you manually specify which drivers to try and load. Additionally, **noprobe** mode allows you to pass parameters to drivers such as IRQ and I/O port. This is quite useful for hardware configured to use non-standard resources.

### Uses for Driver Disks

A driver disk adds support for hardware that is not otherwise supported by the installation program. The driver disk could be produced by Red Hat, it could be a disk you make yourself, or it could be a disk that a hardware vendor includes with a piece of hardware. Red Hat provides four different driver disks with RHEL, all of which may be found on binary CD 1 /images: bootdisk.img drvblock.img, drvnet.img, pcmciadd.img.

There is really no need to use a driver disk unless you need a particular device in order to install RHEL. You will most likely use a driver disk for SCSI adapters and NICs and PCMCIA devices, as those are really the only devices which are used during the installation that might require driver disk support. If an unsupported device is not needed to install RHEL on your system, continue with a regular installation and then add support for the new piece of hardware once the installation is complete.

### Obtaining a driver disk

Your best option for finding driver disk information is on Red Hat's website at <http://www.redhat.com/support/errata/> under the section called Bug Fixes.


If you find a driver disk that is appropriate for your device support needs, create a boot disk using that file. Once you have created your driver disk, boot your system using the diskette as a boot disk and enter either `linux noprobe` or `linux dd` at the boot: prompt.

# Post-Install Configuration

- Setup Agent (`firstboot`)
  - Configure X Window System if necessary
  - Set date and time
  - Register with Red Hat Network and get updated RPMs
  - Install additional RPMs or Red Hat documentation from CDROM
  - Setup users
- `system-config-*` configuration tools

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



25

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2994 or +1 (910) 754 3700.

Often one may want to change one or more of the configuration options selected at install time. One reason might be the addition of new hardware to the system. RHEL includes utilities to modify nearly all options chosen at install time.

## setup

The setup utility is a console-based front end to a number of configuration utilities. These utilities can also be run directly.

## system-config-\*

There are many configuration tools provided by Red Hat whose commands all start with the string `system-config-`. Several examples of tools in this suite would be:

`system-config-display`  
`system-config-printer-{gui,tui}`  
`system-config-date`

## firstboot

RHEL also runs a graphical program called `firstboot` if the system is booted into run-level 5 after the installation. `firstboot` offers several configuration functions such as setting the date and time, installing additional software, or registering for Red Hat Network.

# End of Unit 1

- Questions and answers
- Summary
  - How to identify supported hardware
  - How Linux accesses devices
  - Linux Installation process
  - Post-install configuration basics

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2664 or +1 (619) 754 3700.

## Important files and directories covered in this Unit:

`/dev/*`  
`/usr/src/linux-2.6/Documentation/devices.txt`  
`/images/{bootdisk,drvblock,drvnet,pcmciaadd} img (on CDI)`  
`/images/boot.iso (on CDI)`  
`/dosutils/tawrite.exe (on CDI)`  
`/var/log/messages & /var/log/dmesg & /root/install log`

## Important commands covered in this Unit:

`firstboot, system-config-*`



# Lab 1

## Hardware and Installation

---

### Sequence 1: Preparing the Computer

#### Tasks:

Boot your system with the Red Hat Enterprise Linux CD (disk 1) in the CDROM drive.

Enter the BIOS setup during boot -- ask your instructor if you are not sure how to do this

Set your system's boot order to A, CDROM, C.

Modify any other settings as recommended by the instructor

Save and exit the BIOS setup.

**Sequence 2: Installing Red Hat Enterprise Linux in Graphical Mode****Tasks:**

## Installation STEP-BY-STEP

1. Bootup system using CD.
2. Press Enter at the *boot:* prompt.
3. Choose the appropriate language (English)
4. Press enter on the *OK* prompt
5. Choose the appropriate keyboard (US)
6. Press enter on the *OK* prompt
7. Choose **NFS image** for the installation method
8. Configure TCP/IP. Select **Use dynamic IP configuration (BOOTP/DHCP)**.
9. Press enter on the *OK* prompt.
10. Enter the appropriate information for an **NFS** installation:

**NFS method:**

NFS server name: 192.168.0.254

NFS mount point: /var/ftp/pub

11. At this point Anaconda (the installer) will retrieve the necessary installation image and will probe the system for it's monitor and mouse type and will finally present you with the welcome screen. Click *Next*
12. Choose the appropriate mouse for your system (ask the instructor if you need assistance). Click *Next*
13. Manually partition your system using *diskdruid*. using the following partitioning scheme (delete any pre-existing partitions):  

/boot	100M	
/	400M	
/usr	1256M	
swap	512M	<u>Note: swap is a File System Type not a Mount Point</u>
/var	400M	
14. Format all partitions, but do NOT check for bad blocks unless you wish to spend hours on this lab
15. Use the default Boot Loader settings unless the instructor advises otherwise; do not create a Boot Loader password
16. Choose DHCP for networking and activate on boot

17. Choose *enable Firewall* and allow ssh. Leave SELinux at the default state *Active*.
18. Select the appropriate language support
19. Set the time zone as appropriate for your location; implement UTC if the instructor suggests it
20. Set the root password to *redhat*. (It is not a good password, but please use it anyway)  
Select *Customize the set of packages to be installed* and click *Next*.
21. Unselect **ALL** selections except the *X Window System* and click *Next*. We will be doing a more complete installation in the next part of the lab, so we want to keep this first installation lean.  
  
**NOTE:** The total install size should be approximately 1082MB. If it is larger you have not unselected all packages.
22. You should now be at the *About to Install* screen. Click *Next* to begin.
23. Track the progress of filesystem formatting by switching to `tty5` (`Ctrl-Alt-F5` will take you there; `Alt-F7` will return you to the installer)
24. After the reboot following the installation, complete the initial set up tool, do not register the machine with Red Hat Network. Select "Tell Me Why" followed by "Remind Me Later"

Once you have completed the installation and the newly-installed system has booted, log in as root and examine the following:

- `/var/log/messages`
- `/var/log/dmesg`

This lean installation provides the spartan `twm` window manager. In the next sequence, you will install packages that will provide more functionality and a more attractive environment

**Sequence 3: Installing Red Hat Enterprise Linux in Text Mode****Tasks:**

Now that you have a fully-functional Red Hat Enterprise Linux system, it is time to break it and start over. Before you break it, make sure you have a copy of cd #1 from the install cd set or see your instructor for a copy of the appropriate boot media required to do installations in the classroom

Next, trash your system and reboot using the following command that corresponds to your hardware (IDE or SCSI):

```
cat /var/log/messages > /dev/hda; reboot
cat /var/log/messages > /dev/sda; reboot
```

Once the system goes down insert cd #1 or boot media provided by the instructor and when it comes back up, perform an installation according to the following guidelines. (Note: Because you have overwritten the standard partition table, the installer will warn you that it could not find a suitable partition table, and that it must be initialized.)

**Installation STEP-BY-STEP**

1. Bootup system using CD.
2. Type 'linux text' at the *boot* prompt.
3. Choose the appropriate language (English).
4. Press enter on the *OK* prompt.
5. Choose the appropriate keyboard (US).
6. Press enter on the *OK* prompt.
7. Choose the appropriate installation method (*FTP* or *HTTP*):
8. Configure TCP/IP. **Select "Use dynamic IP configuration (BOOTP/DHCP)"**
9. Press enter on the *OK* prompt
10. Enter the appropriate information for an *FTP* or *HTTP* installation:

**FTP method:**

FTP site name: 192.168.0.254  
Red Hat Directory: pub/

**HTTP method:**

Web site name: server1.example.com  
Red Hat directory: pub/

11. At this point Anaconda (the installer) will retrieve the necessary installation image and will probe the system for it's monitor and mouse type and will finally present you with the welcome screen
12. Choose the appropriate mouse for your system (ask the instructor if you need assistance) Click *Next*
13. Partition your system using `diskdruid`. using the following partitioning scheme (delete any pre-existing partitions ):

```
100M      /boot
2000M     /
512M      (swap) Note: swap is a File System Type not a Mount Point
          3 x 256M RAID0 for /home mount point
```

14. Boot loader, time zone, graphics, and firewall should follow the defaults unless you are instructed to do otherwise
15. Choose ***Disable Firewall*** and ***Active*** for the SELinux setting.
16. Set the language as appropriate
17. Set the root password to ***redhat***
18. Install the default packages.

Note: it is critical that you follow guidance here for the partitioning scheme else some of the follow on exercises may produce unexpected results

10/10/10

10/10/10



# UNIT 2

## System Initialization and Services

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 298 2994 or +1 (619) 754 3700.

## UNIT 2: Objectives

- Upon completion of this unit you should be able to:
  - Describe BIOS functions with respect to the boot process
  - Describe the functions of and configure the boot loader
  - List the functions performed by the kernel during boot
  - State the functions of `init`
  - Use `inittab` to configure `init`
  - List and describe the System V run levels
  - Configure `init` scripts manually and with tools
  - Shutdown and reboot a system into any run level

RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 268 2994 or +1 (910) 754 3700.



## UNIT 2: Agenda

- BIOS boot time responsibilities
- bootloader responsibilities
- kernel boot time responsibilities
- init boot time responsibilities
- System V run levels
- Boot scripts
- Shutdown and reboot

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (918) 754 3700.

# Boot Sequence Overview

- BIOS Initialization
- Boot Loader
- Kernel initialization
- `init` starts and enters desired run level by executing:
  - `/etc/rc.d/rc.sysinit`
  - `/etc/rc.d/rc` and `/etc/rc.d/rc?.d/`
  - `/etc/rc.d/rc.local`
  - X Display Manager if appropriate

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



4

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2904 or +1 (919) 754 3700.

Each major step in a Linux system's boot sequence - BIOS initialization, Boot loader, kernel initialization, and `init` startup - is covered in the upcoming pages.

# BIOS Initialization

- Peripherals detected
- Boot device selected
- First sector of boot device read and executed

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



5

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2994 or +1 (919) 754 3700.

## BIOS - Starting the boot process

The BIOS (Basic Input/Output System) is the interface between the hardware and software on a very basic level. The BIOS provides the basic set of instructions used by the operating system. A successful boot depends on the BIOS, which in fact provides the lowest level of interface to peripheral devices and controls.

The BIOS will first run a power on self test (POST), then it will look for peripherals and a device to boot from. The hardware configuration information is permanently stored in a small area (usually 64 bytes) of CMOS (Complementary Metal Oxide Semiconductor), most commonly referred to as simply "the CMOS." The CMOS is powered by a small battery located in your motherboard. This battery allows the CMOS to retain its settings even when the computer is turned off and disconnected from power.

At the end of the POST, a boot device is selected from the list of detected boot devices. Any modern BIOS will allow you to set the desired order of preference for the boot device from a list. Boot devices could include: the floppy drive, hard drive, CDROM, network-interface, Zip drive or other removable media).

The BIOS reads and executes the first physical sector of the chosen boot media on the system. Usually this is contained in the first 512 bytes of the hard disk.

# Boot Loader Components

- Boot Loader
  - 1st Stage - small, resides in MBR or boot sector
  - 2nd Stage - loaded from boot partition
- Minimum specifications for Linux:
  - Label, kernel location, OS root filesystem and location of the initial ramdisk (`initrd`)
- Minimum specification for other OS:
  - boot device, label

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2964 or +1 (919) 754 3700.

The boot loader is responsible for loading and starting your Linux operating system (or possibly other operating systems) when the computer is started up.

The boot loader is generally invoked in one of two ways:

- BIOS passes control to an initial program loader (IPL) installed within a drive's Master Boot Record
- BIOS passes control to another boot loader, which passes control to an IPL installed within a partition's boot sector.

In either case, the IPL (initial program loader) must exist within a very small space, no larger than 446 bytes. Therefore, the IPL for GRUB is merely a first stage, whose sole task is to locate and load a second stage boot loader, which does most of the work to boot the system

There are two possible ways to configure boot loaders:

- primary boot loader: Install the first stage of your Linux boot loader into the Master Boot Record. The boot loader must be configured to pass control to any other desired operating systems.
- secondary boot loader: Install the first stage of your Linux boot loader into the boot sector of some partition. Another boot loader must be installed into the MBR, and configured to pass control to your Linux boot loader

# GRUB and grub.conf

- GRUB – the GRand Unified Bootloader
  - Command-line interface available at boot prompt
  - Boot from ext2/ext3, ReiserFS, JFS, FAT, minix, or FFS filesystems
  - Supports MD5 password protection
- /boot/grub/grub.conf
- Changes to grub.conf take effect immediately
- If MBR on /dev/hda is corrupted, reinstall the first stage bootloader with:
  - /sbin/grub-install /dev/hda

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

7

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2904 or +1 (616) 754 3700.

The RHEL installer provides the GRUB boot loader, GRUB (the GRand Unified Boot-loader).

/boot/grub/grub.conf has a format of global options followed by boot stanzas. Here is a sample grub.conf:

```
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
password --md5 $1$/ix9y$Bk4yt37Ch2fZ5GFN
default=0
title Red Hat Enterprise Linux AS (2.6.9-648_EL)
    root (hd0,1)
    kernel /vmlinuz-2.6.9-648_EL ro root=/dev/VolGroup00/LogVol100 rhgb quiet
    initrd /initrd-2.6.9-648_EL.img
title Windows XP Pro
    rootnoverify (hd0,0)
    chainloader +1
```

- Seconds before booting default image.
- Splash screen to display at boot
- Encrypted password for CLI
- The first stanza (stanza 0) is the default
- Label for stanza 0
- Files listed below are on (hd0,1) device
- Kernel image and root filesystem
- Initial RAM disk to load
- Label for stanza 1
- Root is (hd0,1), don't mount in GRUB
- Boot from first sector of (hd0,0)

Changes to grub.conf take effect immediately GRUB reads the configuration file at boot time, so the grub.conf file must be stored on a filesystem GRUB understands. These include ext2/ext3, reiserfs, FAT, minix, and FFS. If for some reason your MBR becomes corrupted and you need to reinstall GRUB, you can do so with the command /sbin/grub-install <boot-device>.

Occasionally it may prove necessary for the user to set up grub manually. If grub-install fails for some reason try the following:

1. type the command grub and press enter
2. type root (hd0,0)
3. type setup (hd0)
4. type quit

# Starting the Boot Process: GRUB

- Image selection
  - Select with space followed by up/down arrows on the boot splash screen
- Argument passing
  - Change an existing stanza in menu editing mode
  - Issue boot commands interactively on the GRUB command line



Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## The GRUB Boot Screen

When GRUB starts up, a graphical splash screen can be accessed by pressing `<RETURN>` or `<SPACE>`. This screen has a list of menu entries, normally bootable images. You can select between the different images with the up and down arrow keys, and press `<RETURN>` to select a particular entry for booting.

If you want to pass arguments to boot images through menu editing mode or access the GRUB command line, and a GRUB password is set, you'll need to type `p` followed by your GRUB password.

## Menu Editing Mode

If you then select an entry and type `e`, you'll be dropped into menu editing mode. This mode allows you to modify an existing boot stanza to pass options to the kernel or `init`, or select alternate root filesystems or kernel files than you have configured in your existing stanzas. You can use arrow keys to select a line, `e` to edit a line, `d` to delete a line, `o` to add a line, and `b` to boot. For example, to boot into runlevel 2, you could select menu editing mode, select your Linux boot stanza, add a `2` to the end of your `kernel` line, and type `b` to boot the modified menu entry.

## The GRUB Command Line

GRUB provides a command-line interface which can be used to write a temporary boot command from scratch, view the contents of files on the filesystem, perform diagnostic tests, or experiment with GRUB configurations. Most commands supported by the configuration file are available for interactive use. Editing commands are similar to those used by the bash shell, and `<TAB>` completion is available. If GRUB is not able to find a valid `grub.conf` file, it will default to the command line.

To exit menu editing mode or the command line and go back to the main GRUB menu, type `<ESC>`.

For more information about GRUB and `grub.conf`, look at `info grub`.

# Kernel Initialization

- Kernel boot time functions
  - Device detection
  - Device driver initialization
  - Mounts root filesystem read only
  - Loads initial process (`init`)

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## Examining kernel initialization

Although they generate good output, the kernel initialization activities take place so quickly that if you don't watch carefully during boot, you may miss them. A good way to "freeze time" and examine this output is to view `/var/log/dmesg`, which contains a snapshot of these kernel messages taken just after control is passed to `init`. Review of this output will reveal the basic initialization steps of the Linux kernel:

Device drivers compiled into the kernel are called, and will attempt to locate their corresponding devices. If successful in locating the device, the driver will initialize and usually log output to the kernel message buffer.

If essential (needed for boot) drivers have been compiled as modules instead of into the kernel, then they must be included in an `initrd` image, which is then temporarily mounted by the kernel on a RAM disk to make the modules available for the initialization process.

After all the essential drivers are loaded, the kernel will mount the root filesystem read-only.

The first process is then loaded (`init`) and control is passed from the kernel to that process.

# init Initialization

- `init` reads its config: `/etc/inittab`
  - initial run level
  - system initialization scripts
  - run level specific script directories
  - trap certain key sequences
  - define UPS power fail / restore scripts
  - spawn gettys on virtual consoles
  - initialize X in run level 5

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photographed, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 299 2994 or +1 (919) 754 3700.

## init

`init` is the parent of all processes This is easily shown by running the `ps tree` command:

```
$ ps tree
init--apmd
|-atd
|-automount
|-crond---crond
|-deskguide_apple
|-gdm--X
|      ^-gdm---gnome-session
```

Because `init` is the first process, it will always have a PID of number 1.

The file `/etc/inittab` contains the information on how `init` should set up the system in every run level, as well as the run level to use as default

If the `/etc/inittab` file is missing or seriously corrupt, you will not be able to boot to any of the standard run levels (0-6) and will need to use single or emergency mode instead This procedure is discussed in depth in Unit 10 of this course.



## Run levels

- `init` defines run levels 0-6, S, emergency
- The run level is selected by either
  - the default in `/etc/inittab` at boot
  - passing an argument from the boot loader
  - running `init x` after boot (where `x` is the desired run level)
- Show current and previous run levels
  - `/sbin/runlevel`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2884 or +1 (919) 754 3700.

The following chart details the run levels that Linux defines by default:

Run Level	Effect
0	Halt (Do <i>Not</i> set <code>initdefault</code> to this)
1, S, emergency	Single-user modes (Only the root user can be logged on. Used to perform Maintenance)
2	Multi-user, without NFS networking
3	Full multi-user mode. (Includes networking)
4	User definable, but duplicate of runlevel 3 by default
5	X11 (Includes networking)
6	Reboot (Do <i>Not</i> Set <code>initdefault</code> to this)

The `initdefault` line in the file `/etc/inittab` controls the default run level after the system is started. Its format is as follows:

```
id:x:initdefault:
```

where `x` is the run level desired after the system is started.

Run levels (7-9) are also valid, though undefined and not really documented. This is because "traditional" UNIX variants don't use them.

If `inittab` does not have a default run level selected the system will attempt to boot to run level 9 which is undefined.

## /etc/rc.d/rc.local

- Run after the run level specific scripts
- Common place for custom modification
- In most cases it is recommended that you create a System V `init` script in
- `/etc/rc.d/init.d` unless the service you are starting is so trivial it doesn't warrant it. Existing scripts can be used as a starting point.

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2994 or +1 (918) 754 3700.

### `rc.local` - Final System V Initialization

Because the `rc.local` script is run each time the system enters a run level, it is a convenient place to start processes that need to be running

# Virtual Consoles

- Multiple independent VT100-like terminals
- Defined in `/etc/inittab`
- Accessed with `Ctrl-Alt-F_key` from an X session
- `/dev/ttyn`: virtual console `n`
- `/dev/tty0`: the current virtual console
- Default Red Hat Enterprise Linux configuration:
  - 12 consoles defined
  - consoles 1-6 accept logins
  - X server starts on first available console, usually 7

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



17

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2984 or +1 (818) 754 3700.

RHEL provides multiple VT100-like terminals, accessible via the `/dev/ttyn` devices. Users may switch from one console to another using an `Alt-function_key` sequence. The `Alt-LeftArrow` and `Alt-RightArrow` keys can also be used to cycle through consoles. `Shift-PageUp` and `Shift-PageDown` provide scrollback buffer browsing, although this buffer is cleared when changing consoles.

By default, the `init` process respawns `mingetty` processes for the first six virtual consoles, allowing six independent login sessions. This is specified by `/etc/inittab`. When an X server is started, it attaches to the first available virtual console, generally `/dev/tty7`.

Because the `Alt-function_key` sequence is common within X, an additional `Ctrl-key` sequence has been added: `Ctrl-Alt-function_key` is required to switch out of an X session to a text-based console. It is also possible to switch from one virtual console to another by using `chvt`.

The root user can make use of additional consoles by accessing the `/dev/ttyn` device node directly, as in the following example:

```
tail -f /var/log/messages > /dev/tty9 &
```

When `mingetty` displays the contents of `/etc/issue`, it expands certain escape sequences that may appear in that file. See the `mingetty(8)` man page for details.

## Controlling Services

- Utilities to control default service startup
  - **system-config-services**: graphical utility that requires an X interface
  - **ntsysv**: ncurses based utility usable in virtual consoles
  - **chkconfig**: a fast, versatile command line utility that works well and is usable with scripts and Kickstart installations
- Utilities to control services manually
  - **service**: immediately start or stop a standalone service
  - **chkconfig** immediately starts and stops **xinetd**-managed services

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 296 2994 or +1 (919) 754 3700.

The job of the System V initialization scripts is to start services at boot time. Most of these services run as daemons, such as *cups*, *crond* and *sendmail*. Red Hat Enterprise Linux includes several utilities that facilitate the management of System V initialization.

- **system-config-services** is an X client that presents a display of each of the services that are started and stopped at each run level. Services can be added, deleted, or re-ordered in run levels 3 through 5 with this utility.
- **ntsysv** is a console-based interactive utility that allows you to control what services run when entering a given run level. This utility is used during system installation, but can be run from the command line. It configures the current run level by default. By using the **--level** option you can configure other run levels.
- **chkconfig** is a command-line utility. When passed the **--list** switch, it displays a list of all System V scripts and whether each one is turned on or off at each run level. Scripts can be managed at each run level with the **on** and **off** **chkconfig** directives. The **--level** option can be used to specify the runlevels affected if the defaults are unacceptable.
- The **service** command is used to start or stop a standalone service immediately; most services accept the arguments **start**, **stop**, **restart**, **reload**, **condrestart**, and **status** as a minimum.
- The **system-config-services** and **chkconfig** commands will start or stop an **xinetd**-managed service as soon as you configure it on or off. Standalone services won't start or stop until the system is rebooted or you use the **service** command.

# System Shutdown

- Shutting down the system
  - `shutdown -h now`
  - `halt`
  - `poweroff`
  - `init 0`

Rev RH133 RHEL 4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

RHEL is a very reliable operating system, and rarely requires a shutdown or reboot. In fact, usually the only time Linux requires a reboot or shutdown is when you need to add or remove hardware, upgrade to a new version of RHEL, or upgrade your kernel

## **shutdown**

The `shutdown` command supports several options such as:

- a Use `/etc/shutdown.allow`.
- k Don't really shutdown; only send the warning messages.
- r Reboot after shutdown
- h Halt after shutdown.

## **halt**

`halt`, like `shutdown`, also supports several options such as:

- n Do not sync the hard disk before halting.
- i Shutdown all interfaces before halting

## **poweroff**

`poweroff` allows you to shutdown and power the system off. `poweroff` also uses the same switches as `halt`.

## **init 0**

Issuing this command simply tells the system to change to run level 0, which is shutdown.

# System Reboot

- Rebooting rarely fixes problems in Linux
  - If you feel a reboot is necessary try bringing the system down to runlevel 1 and then back up to runlevel 3 or 5. This is much faster than a reboot.
- Rebooting the system
  - `shutdown -r now`
  - `reboot`
  - `init 6`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



20

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2864 or +1 (919) 754 9700.

## Rebooting

In the event that you need to reboot your RHEL system, there are a few options available

`shutdown -r now`

Invoking `shutdown -r now` tells the system to shutdown and restart

`reboot`

`reboot` does exactly what it says

`init 6`

Issuing this command tells the system to switch to run level 6, which is reboot

You can also reboot by pressing `Ctrl+Alt+Del` at a virtual console. The standard Red Hat Linux `/etc/inittab` binds this keystroke combination to `shutdown -t3 -r now`.

## End of Unit 2

- Questions and answers
- Summary
  - What functions does the kernel perform at boot?
  - What are the System V run levels?
  - What commands can you use for shutting down and rebooting?

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



21

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2994 or +1 (919) 754 9700.

### Important files covered in this Unit:

```
/boot/grub/grub.conf
/etc/inittab
/etc/rc.d/rc.sysinit
/etc/rc.d/rc
/etc/rc.d/init.d
/etc/rc.d/rc.local
```

### Important commands covered in this Unit:

```
init
mingetty
shutdown
reboot
halt
poweroff
chkconfig
ntsysv
redhat-config-services
service
```

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100



# Unit 2 Lab

## Managing Startup

---

Estimated Duration: 30 minutes

Goal: To build skills customizing system services

Setup at Start: A Red Hat Enterprise Linux System

Situation: Your company has decided that security is a concern and would therefore like you to disable certain services that might pose a security risk. You also want to setup a login banner to "warn away" potential intruders.

**Sequence 1: Disabling services with chkconfig****Scenario/Story:**

You have decided to disable unneeded services on your machine

**Tasks:**

1. Use `chkconfig` to view the status of the system services:

```
chkconfig --list
```

2. Following the example below use `chkconfig` to turn off `isdn` in all runlevels:

```
chkconfig --del <service name>
```

*chkconfig --del isdn*

3. Using `chkconfig`'s syntax information displayed with `chkconfig --help`, turn off service `kudzu` in runlevels 3 and 5 only

*chkconfig --level 35 kudzu off*

4. Observe the differences between `on` and `--add`, and between `off` and `--del` using the following commands:

```
chkconfig isdn --list
```

```
chkconfig isdn on
```

```
chkconfig isdn --list
```

```
chkconfig isdn off
```

```
chkconfig isdn --list
```

```
chkconfig isdn --del
```

```
chkconfig isdn --list
```

*not found*

```
chkconfig isdn --add
```

```
chkconfig isdn --list
```

5. Use `chkconfig` to view the status of the system services and to verify your changes.

**Deliverable:**

A machine with several default services disabled.

**Sequence 2: Changing the system login banner****Tasks:**

1. We're going to set up the `rc.local` script so that it regenerates the login banner every time the system reboots. Open the file `/etc/rc.local` in a text editor and locate the following line:

```
touch /var/lock/subsys/local
```

2. Insert the following lines immediately before that line:

```
echo "Welcome to \n" > /etc/issue  
echo "All access to this system is monitored" >> /etc/issue  
echo "Unauthorized access is prohibited" >> /etc/issue  
echo >> /etc/issue  
echo "Last reboot complete at $(/bin/date)" >> /etc/issue
```

3. Save the file, then copy `/etc/issue` to `/etc/issue.orig`.
4. Reboot your system.
5. When the system comes up verify the new login banner by changing to a virtual console. (Hint: press `ctrl-alt-F1`.) Look in `/etc/issue`. Note that `mingetty` expands the `\n` in your `/etc/issue` file into your machine's hostname on the screen

**Deliverable:**

A system with a login banner similar to the following:

```
Welcome to station10.example.com  
All access to this system is monitored  
Unauthorized access is prohibited  
  
Last reboot complete at Tue Nov 27 16:03:59 EDI 2001
```

Sequence 3: Changing the default run level

Tasks:

- 1 Edit the /etc/inittab file and change the default run level as shown below from level 5 to level 3.

```
id:3:initdefault:
```

2. Reboot the system. What happens?

*TEXT Mode - multiuser.*

3. Log in and edit /etc/inittab to change the default run level to 1 and reboot.

*single user mode.*

4. Change the default run level back to level 5 and reboot.

Sequence 4: Adding a Message Of The Day (motd)

Tasks:

- 1 Edit the file /etc/motd, which should currently be empty. Add the following lines:

```
#####
#           Welcome to Station xx           #
#####

<date>    The sysadmin is playing today.
          Expect frequent system downtime.
```

Where <date> is today's date which you manually enter and xx is your station number.

- 2 Change to a virtual console and login.

**Sequence 5: GRUB**

**Tasks:** Use GRUB at boot time to bring up Linux in various run levels.

1. Reboot Linux so that GRUB appears on your screen. If you have specified a "timeout=" value in grub.conf you will notice that the timer is counting down.
2. Before the timer counts down to zero, press the space bar to halt the timer.
3. Take note of the help text in the lower part of the GRUB display. Use the up/down arrow keys to navigate to the kernel you wish to boot. Then press the "e" key to edit the contents of grub.conf for this kernel.
4. Once again, take note of the help text in the lower portion of the GRUB display. Use the up/down arrows to navigate to the line starting with the text "kernel" and press the "e" key.
5. You are now in GRUB edit mode with the cursor at the end of the line. Press the spacebar followed by the "s" key, then press the "enter" key. You will note that the GRUB display returns to the prior screen and now has the new text "S" appended to the kernel line. If you wish to undo all changes you have made in GRUB, simply press the "ESC" key to return to the prior screen.
6. Press the "b" key to boot using these grub.conf options. In this example, you will come up in runlevel "S" or single user.
7. Following reboot, review the contents of the grub.conf file. You will note that the change you made at the GRUB screens did not update this file.
8. Repeat steps 1-6 above, trying different runlevels such as "emergency", "1", etc.



THE UNIVERSITY OF CHICAGO  
LIBRARY

# UNIT 3

## Kernel Services and Configuration

RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 296 2964 or +1 (919) 754 3700.

## UNIT 3: Objectives

- Upon completion of this unit you should be able to:
  - Load, list, and unload kernel modules
  - View system configuration information in the `/proc filesystem`
  - Configure running kernel parameters with the `/proc filesystem`

Rev RH133-RHEL4.1

Copyright © 2005 Red Hat, Inc.

redhat



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.



## UNIT 3: Agenda

- Kernel modules
- The `/proc` filesystem

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2864 or +1 (616) 754 3700.

# Kernel Modules

- Modular kernel components
  - components that need not be resident in the kernel for all configurations and hardware
    - peripheral device drivers
    - supplementary filesystems
  - modules configurable at load time
- `/lib/modules`
- Controlling modules
  - `lsmod`, `modprobe`
- Kernel Tainting

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



4

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## Modular kernel components

Many components of the kernel can be compiled as dynamically loadable modules. This allows for increased kernel functionality without increasing the size of the kernel image loaded at boot time. Good candidates for modularization are any supplementary capabilities that are not needed at boot time, including peripheral device drivers, supplementary filesystems

## The `/lib/modules` directory

Kernel modules reside in `/lib/modules/<kernel-version>`. The directory name must match the kernel version as returned by `"uname -r"`

## Controlling Modules

Modules are generally loaded on demand by the kernel. `lsmod` lists the modules currently resident in the kernel. The kernel can be prompted to load a particular module with `modprobe`, and various module parameters can be specified as command line arguments. Additionally modules loaded with `modprobe` may have aliases, options or actions provided by `/etc/modprobe.conf`. Modules may be inserted with `modprobe <module-name>` and removed with `modprobe -r <module-name>`.

*modules.conf*

## Kernel Tainting

If a module with a proprietary license is inserted into the kernel, it becomes tainted. Tainted kernels can cause support issues since developers, with no access to the source, cannot tell what the proprietary code is doing or what problems it may be causing. Unless there is a specific exception, Red Hat will not support tainted kernels. The current taint status of the kernel is displayed among the headers from `lsmod` and the licenses of modules can be viewed with `/sbin/modinfo`.

# Kernel Module Configuration

- Module examination: `/sbin/modinfo`
  - parameters, license
- Module Configuration: `/etc/modprobe.conf`
  - aliases, parameters, actions
- Module Dependencies: `modules.dep`, `depmod`
- Manual Control: `insmod`, `rmmmod`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 299 2994 or +1 (919) 754 3700.

## Module Examination

Available module parameters and the modules license can be viewed with `modinfo`.

## Module Configuration

Aliases provide uniform ways to address various types of hardware. By default, aliases are used for ethernet interfaces, sound cards and usb controllers. Some examples from a `modprobe.conf`:

```
alias eth0 e100
alias eth1 airo
alias snd-card-0 snd-intel8x0
alias usb-controller uhci-hcd
```

Many modules accept parameters that can be specified at load time. The `softdog` module provides a watchdog that can reboot the system if it appears to be locked up. The `soft_margin` defines how many seconds to wait before issuing a reboot. This option can be set in `/etc/modprobe.conf` with the line:

```
options softdog soft_margin=120
```

Often when modules are loaded, an action should be run to configure the new device. Common actions include setting and storing the volume settings on sound cards and forcing duplex settings on network cards.

```
install snd-intel8x0 /sbin/modprobe --ignore-install snd-intel8x0 && \
/usr/sbin/alsactl restore >/dev/null 2>&1 || :
remove snd-intel8x0 { /usr/sbin/alsactl store >/dev/null 2>&1 || : ; }; \
/sbin/modprobe -r --ignore-remove snd-intel8x0
install eth0 /sbin/modprobe --ignore-install eth0 && /sbin/mii-tool -F
100baseTx-FD eth0
```

## Module Dependencies

Some modules depend upon functionality provided by other modules. The `depmod` command can be used to rebuild the dependency database:

```
/lib/modules/$(uname -r)/modules.dep
```

This command may be needed when custom modules are built for the kernel

## Manual Control

The `modprobe` functionality can be bypassed with calls to `insmod` and `rmmod`. These are mentioned from sake of completeness. `modprobe` and `modprobe --remove` are the preferred methods to insert and remove modules from the kernel.

# The /proc filesystem

- /proc is a virtual filesystem containing information about the running kernel
- Contents of "files" under /proc may be viewed using `cat`
- Example:  

```
cat /proc/interrupts
```
- Provides information on system hardware, networking settings and activity, memory usage, and more

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

The /proc filesystem, which is not a disk-based filesystem, is enabled (or disabled) in the kernel itself, and is a map to the running kernel process. The /proc filesystem is mounted during system initialization through an entry in /etc/fstab.

Listing the "files" and "directories" under /proc will reveal that virtually all of them have a size of zero: they are not really files and directories in the typical sense. You can `cd` into the directories as you would directories on a disk-based filesystem, but the appropriate way to view the contents of the files is by using the `cat` command, rather than using an editor or even paging commands like `more` or `less`. Do not use `cat` on /proc/kcore, as this special file is an image of the running kernel's memory at that particular moment -- `cat`'ing this file will leave your terminal unusable.

Some of the key files in the top-level directory include:

- /proc/interrupts -- IRQ settings
- /proc/cpuinfo -- information about the system's CPU(s)
- /proc/dma -- DMA settings
- /proc/ioports -- I/O settings
- /proc/iomem -- memory ranges for PCI devices
- /proc/meminfo -- information on available memory, free memory, swap, cached memory, and buffers
- /proc/loadavg -- system load average
- /proc/uptime -- system uptime and idle time
- /proc/version -- information on Linux kernel version, build host, build date, etc.

## The /proc filesystem, cont'd

- /proc subdirectories
- The /proc/sys subdirectory allows administrators to modify certain parameters of a running kernel



Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

Beneath the top-level /proc are a number of important subdirectories containing files with useful information. These include:

- /proc/scsi -- information about SCSI devices
- /proc/ide -- information about IDE devices
- /proc/net -- information about network activity and configuration
- /proc/sys -- kernel configuration parameters
- /proc/<PID> -- information about process *PID*

/proc/sys is unique in that its parameters may be modified on a running system if CONFIG\_SYSCTL is enabled in the kernel. For a complete description of the available parameters under /proc/sys, read the documentation of the proc filesystem in /usr/share/doc/kernel-doc-\*/Documentation/filesystems/proc.txt and in /Documentation/sysctl/. Below are a few examples of parameter changes one might make using /proc/sys:

- echo "1" > /proc/sys/net/ipv4/ip\_forward # Turn on IP forwarding
- echo "16384" > /proc/sys/fs/file-max # Double the number of file handles

Some /proc entries contain multiple space-delimited values. As an example, semaphores are data structures used by certain programs to control access to shared resources. The values in /proc/sys/kernel/sem indicate: the maximum number of semaphores per semaphore array, the maximum value a semaphore can contain, the maximum number of semaphore operations that can be requested during a single call to a semaphore function and the maximum number of semaphore arrays that can be created.

```
echo "500 32000 64 256" > /proc/sys/kernel/sem
```

## `/proc/sys` configuration with `sysctl`

- `/proc/sys` modifications are temporary and not saved at system shutdown
- The `sysctl` command manages such settings in a static and centralized fashion:
  - `/etc/sysctl.conf`
- `sysctl` is called at boot time by `rc.sysinit` and uses settings in `/etc/sysctl.conf`

Rev RY133/RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 259 2594 or +1 (618) 754 3700.

During system boot, `rc.sysinit` calls `sysctl -e -p /etc/sysctl.conf`. It also sets values for `/proc/sys/kernel/modprobe` and `/proc/sys/kernel/hotplug`. This automatically re-establishes the otherwise temporary `/proc/sys` values. Left-side values are paths within `/proc/sys`.

Typical entries in `/etc/sysctl.conf` might be:

```
# Disables IPv4 packet forwarding
net.ipv4.ip_forward = 0
# Enables source route verification
net.ipv4.conf.all.rp_filter = 1
# Enables the magic-sysrq key
kernel.sysrq = 1
```

For a complete list of valid `/proc/sys` parameters for `/etc/sysctl.conf`, look in `/usr/share/doc/kernel-doc-*/sysctl/` or just use the `sysctl` command directly:

```
sysctl -a
```

# General Hardware Resources

- `dmesg` and `/var/log/dmesg`
- `kudzu`
  - `/etc/sysconfig/hwconf`
  - `/usr/share/hwdata/`
- `/proc` filesystem
- `hwbrowser`

Rev R0133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (519) 754 3700.

## `dmesg` and `/var/log/dmesg`

The `dmesg` command displays the contents of the kernel ring buffer, which contains boot-time messages immediately after boot time. `klogd` writes to the ring which has a default size of 32768 bytes for uniprocessor, or 65536 bytes for SMP kernels. As new messages are written to the buffer, the boot messages are dropped, hence, the need to dump the output of `dmesg` during system startup to `/var/log/dmesg` in order to retain boot messages.

## `kudzu`

*detected hw*

The `kudzu` utility maintains a database of detected and configured hardware, found at `/etc/sysconfig/hwconf`. As part of the boot process, `kudzu` compares the currently detected hardware to the stored database. This comparison can be forced by calling the `kudzu` command directly. If new hardware is detected, or previously existing hardware is removed, `kudzu` will attempt to automatically reconfigure the system, or steer the administrator to the appropriate interactive configuration utility. `kudzu` uses catalogs of known hardware in the `/usr/share/hwdata/` directory.

## `/proc` filesystem

The `/proc` filesystem contains pseudo-files which provide detailed hardware information. The `meminfo`, `cpuinfo`, `interrupts`, `ioports`, and `iomem` pseudo-files, and `bus`, `ide`, and `scsi` directories, are only a few examples of the wealth of information available. Because the filesystem is implemented internally by the kernel, it exists even in minimal environments.

## `hwbrowser`

`hwbrowser` provides a convenient graphical utility that provides a survey of detected hardware.

*kernel dbus* *bell - dev - manager*



# System Bus Support

- PCI Bus
  - `/sbin/lspci`
  - `/proc/bus/pci/`
- ISA Bus
  - `/proc/isapnp/`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2094 or +1 (919) 754 3700.

The PCI bus plays a primary role in most x86 compatible architectures. The PCI protocol supports Plug and Play configuration, and supports a standard identification protocol. The bus can be probed with varying levels of verbosity with `/sbin/lspci`. Examining the output of `lspci` generally reveals controllers that bridge other busses onto the PCI bus, as well as PCI peripheral devices. `/proc/bus/pci` also provides information about detected PCI devices.

Starting with the 2.4 kernel, Plug and Play compatible ISA devices are configured internally by the kernel. Evidence of detected and configured devices can be found in `/proc/isapnp`.

# Hotswappable Bus Support

- USB and IEEE 1394 Buses
  - `/sbin/hotplug`, (`/etc/hotplug/`)
  - Information in `/proc/bus/` subdirectories
  - `/sbin/lusb` and `/sbin/usmodules` utilities
  - USB devices in `/dev/usb/`
- PCMCIA Bus
  - `/sbin/cardmgr`, (`/etc/pcmcia/`)
  - Information in `/proc/bus/pccard`
  - `/sbin/cardctl` utility

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2964 or +1 (919) 754 3700.

The `/sbin/hotplug` program is used by the kernel to notify processes when a device is plugged into a USB or IEEE 1394 (FireWire) peripheral bus. Specific agents in `/etc/hotplug` are executed to load appropriate modules. For USB, the `/sbin/lusb` command can be used to list detected devices. Device access is often provided by the `usbdevfs` virtual filesystem, which is mounted to `/dev/usb`. More information can be found in `/usr/share/doc/hotplug-version`, and in the kernel source `/usr/src/linux-2.6/Documentation/usb` directory.

PCMCIA support is also implemented by kernel modules, with the system-specific PCMCIA controller defined in `/etc/sysconfig/pcmcia`. Hotswap events are monitored by the `/sbin/cardmgr` daemon. Cards can also be identified using `/sbin/cardctl`, which may be invoked directly. The `/etc/pcmcia/` directory provides numerous configuration files which map detected cards to appropriate kernel modules. The PCMCIA infrastructure is initialized using the `/etc/init.d/pcmcia` service script. The `pcmcia(5)` man page and `/usr/share/doc/kernel-pcmcia-version/` directory provide more information.

## System Monitoring and Process Control

- `top`, `gnome-system-monitor` - display snapshot of processes
- `vmstat` - reports virtual memory stats
- `iostat` - lists information on resource usage, including I/O statistics
- `free` - summary of system memory usage
- `renice` - change priority of a process
- `kill` - send system signal to a process

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 336 2004 or +1 (616) 754 3700.

There are several ways to monitor and effect system performance. The `top` and `gnome-system-monitor` programs display a snapshot of running processes that is updated every few seconds.

`vmstat` reports information on virtual memory usage and `free` supplies a summary of system memory usage.

`renice` can be used to change the priority of a process so that it uses more or less system resources.

`kill` is a command that sends signals to running processes. By default, `kill` asks a process to shutdown. If the process will not respond to a "friendly" kill, you can invoke `kill` with `-9` option which will forcibly kill the process. Another signal that `kill` can send is `-HUP`. This option tells a compliant program to reread its configuration files. For a complete list of the signals that `kill` can send review the `kill` man page.

## End of Unit 3

- Questions and Answers
- Lab

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



13

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (918) 754 3700.

### Important files covered in this Unit:

```
/proc/*  
/etc/sysctl.conf  
/lib/modules/*  
/etc/modules.conf
```

### Important commands covered in this Unit:

```
sysctl  
lsmod, insmod, rmmod, modprobe, depmod  
top, kill, free, renice, vmstat, iostat
```

# Lab 3

## Configuring kernel parameters

---

**Goal:** Develop skills tuning the `/proc` filesystem.

### Sequence 1: Turning off ping responses

1. Check the present value of `/proc/sys/net/ipv4/icmp_echo_ignore_all`

```
cat /proc/sys/net/ipv4/icmp_echo_ignore_all
```

It should be currently set to zero which means your system will respond normally to pings.

2. Change the value of `/proc/sys/net/ipv4/icmp_echo_ignore_all` to a 1 which will prevent other hosts from successfully pinging your host while not effecting your ability to ping them. Verify your work.

```
echo "1" > /proc/sys/net/ipv4/icmp_echo_ignore_all  
cat /proc/sys/net/ipv4/icmp_echo_ignore_all
```

3. Now test pinging server1 example.com. Pressing Ctrl-c will stop the ping command and display some statistics for you. You should have been able to ping server1.
4. Next have someone else try pinging your station. They should not receive any responses back from your system.
5. Now reboot your system and try steps 3 and 4 again. What happened? Why?
6. Remember that changes to the `/proc` filesystem are temporary and if you want them to persist across reboots you need to put an entry in `/etc/sysctl.conf`.

- a) edit `/etc/sysctl.conf` and put the following line at the bottom:

```
net.ipv4.icmp_echo_ignore_all=1
```

- b) execute as root:

```
sysctl -p
```

Check the value in `/proc`. If it is not set to a 1 then recheck the previous two steps. Next reboot your system and check the value in `/proc` again.

**MANDATORY CLEANUP:**

- 1) comment out or remove `net.ipv4.icmp_echo_ignore_all=1` from `/etc/sysctl.conf`
- 2) As root, run `sysctl -p`


This is to prevent other things from breaking during the week and help preserve your and your instructors sanity. Please note, it might be too late for your instructor.

# UNIT 4

## Filesystem Management

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (919) 794-3700.

## UNIT 4: Objectives

- Upon completion of this unit you should be able to:
  - Explain how data is accessed and maintained
  - Understand the filesystem hierarchy
  - Manage the filesystem hierarchy
  - Manage virtual memory with swap partitions
  - Add a hard drive

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being inappropriately used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (319) 754 3700.



## UNIT 4: Agenda

- Initial device access
- Partitions and device preparation
- Filesystem basics
- The filesystem hierarchy
- Manage virtual memory
- Adding a New Drive

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 238 2894 or +1 (513) 754-3700.

## System Initialization: Device Recognition

- Master Boot Record ( MBR ) contains:
  - Executable code to load operating system
  - Space for partition table information, including:
    - Partition id or type
    - Starting cylinder for partition
    - Number of cylinders for partition

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



4

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2984 or +1 (919) 754 3700.

When a system boots, a search for code that can start an operating system is made from a list of devices (e.g. CD-ROM, floppy, hard drive) as defined in the BIOS. The first executable code found is used. Most often, a system boots from a hard drive attached to the system mainboard, and from the code found at the first sector, of the first cylinder--the Master Boot Record, or MBR--of that drive. This executable code is called a *boot loader*.

There are many boot loaders: their function is the same, but their interfaces and capabilities differ. The original MS-DOS boot loader has limited capabilities and is only capable of booting Microsoft DOS and older versions of the Windows operating system. Other boot loaders such as the one used by Microsoft WindowsNT/2000, or GNU GRUB (GRand Unified Boot loader) are more flexible, and can accommodate systems configured to run more than one operating system. Note that boot loaders do not load more than one OS at a time, even on "multi-boot" systems.

### Primary Partitions

In addition to the boot loader described above, the MBR contains a structure describing the hard drive partitions. IDE drives on Linux use legacy structures to describe four primary partitions, with provision for an extended partition and its "logical" partitions. Each partition, described by its size -- in sectors, blocks, or cylinders -- and its offset from the "zeroth" cylinder, has a type which is also stored in this MBR structure. Linux-specific partitions would normally be one of the following types:

- 0x5 (or 0xf) - Extended
- 0x82 - Linux swap
- 0x83 - Linux
- 0x8e - Linux LVM
- 0xfd - Linux RAID auto

# Disk Partitioning

hda (3.0)  
hdb (3.1)

- An extended partition points to additional partition descriptors
- Total maximum number of partitions supported by the kernel:
  - 63 for IDE drives
  - 15 for SCSI drives
- Why partition drives?
  - containment, performance, quotas, recovery



Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 288 2864 or +1 (919) 754 3700.

## Extended Partitions and Logical Partitions

If one of the primary partitions is marked as Extended (type '0x5') or Win95 Extended (type '0xf'), then the first sector of the partition described by the entry will contain another block of partition descriptors. These descriptors define partitions known as logical partitions. Use of logical partitions is a work-around for limitations in the legacy, Microsoft DOS-based partition table structure. Logical partitions permit the definition of more than four partitions per drive. While the PC partition specification does not impose a limit on the number of logical partitions, the kernel does

### Partition Limits

63

The Linux kernel is designed with specific device numbers, the numeric "name" of the device driver for a given device. This allocation of device and number supports a maximum of 63 total partitions on each IDE disk, with one partition assigned one device number. On SCSI disks, the maximum number of partitions supported is 15, again due strictly to device number allocation. For more information about the kernel, install the *kernel-doc* RPM and reference `/usr/share/doc/kernel-doc-*/` for, among others, `devices.txt`. In order to operate on higher-numbered partitions, you may need to create the appropriate device files manually. See the man and info pages for `mknod` for details

### Why partition?

Unix best practices suggest that we should partition our disks for many reasons. By creating a separate filesystem we can contain applications and users to that filesystem. If it fills up because of security breach or user demand, the rest of the operating system is more insulated from the issue. Separate partitions improve performance by keeping data together which reduces disk head seek. If you would like to use quotas, they are enabled at the filesystem level. Also partitioning eases backup and recovery. If your application and its data are on separate filesystems, the operating system can be upgraded or reinstalled without having to restore the data from elsewhere.

# Managing Partitions

- Create partitions using:
  - `fdisk`
  - `sfdisk`
  - GNU `parted` - Advanced partition manipulation (create, copy, resize, etc)
- `partprobe` - reinitializes the kernel's in memory version of the partition table

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (919) 754 3700.

## Partitioning Tools

### `fdisk`

According to its own documentation, `fdisk` is "a buggy program that does fuzzy things - usually it happens to produce reasonable results." The man page is too critical - `fdisk` is the most commonly-used partitioning program. It has the advantage that it has some support for BSD disk labels and other non-DOS partition tables

### `sfdisk`

The user interface is somewhat cryptic, but it is more accurate than `fdisk` as well as more flexible. Moreover, it can be used non-interactively (i.e., in a script).

### GNU `parted`

After installation you may need a program for creating, removing, resizing, and copying partitions containing file systems. `parted` manages these tasks for a variety of filesystem types.

### `partprobe`

At system bootup, the kernel makes its own in-memory copy of the partition tables from the disks. Most tools like `fdisk` edit the on-disk copy of the partition tables. To update the in-memory copies, run `partprobe`.

## Managing Data: Filesystem Creation

- mkfs
- mkfs . ext2, mkfs . ext3, mkfs . minix,  
mkfs . msdos, mkfs . Vfat
- Specific filesystem utilities may be called directly  
mke2fs [options] device

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 288 2894 or +1 (619) 754 3700.

The command for creating a filesystem is `mkfs`. This command is a "front end" or "wrapper" for various filesystem creation programs. `mkfs -t fstype` will look for programs that follow the naming convention `mkfs.fstype`, and then run the program to create the desired filesystem. If `mkfs` is run without `-t option`, it assumes the Linux default ext2 (Second Extended) filesystem. The `mkfs.fstype` form of the program may also be called directly. The ext2 filesystem has an additional interface called by:

`mke2fs [options] device`

Some useful options include: *4K block size*

-b specify the size of datablocks in bytes. Without this option, the datablock size is determined by the size of the partition. Filesystems that will contain many small files, for example, should use smaller block sizes because each file uses an entire block, even if its data is less than the size of one block. One file occupies, at least, one datablock. Compare the variable size of a data block with the invariable size of the physical hard drive sector (usually 512 bytes).

-c check the device for bad blocks (a k a sectors) before creating the file system. This may take several hours if the partition is very large!

-i specify the bytes/inode ratio. `mke2fs` creates an inode table based on the total size in bytes of a potential filesystem. Large datasets managed by a filesystem would benefit from a higher ratio as there would be fewer inodes, freeing more data space. Only one inode is allocated per file, while one inode may reference one or more datablocks. Inodes are 128 bytes in size and unused inodes can occupy a lot of available space.

-N overrides the default calculation of the number of inodes that should be reserved for the file system. By default, this is based on the number of blocks and the bytes/inode ratio. This allows the user to specify the number of desired inodes directly. Note: this may be a more effective method to best utilize the partition or dataspace bytes/inode ratio for extremely large datasets as the largest value passed to the `-i` option above is 8192.

-m specify the percentage of reserved blocks for the super-user. This value defaults to 5%. If the file system being created will be used for some specific application, changing the value to zero will allow the application full use of the filesystem.

-L set the volume label for the filesystem. This option will be very useful as we discuss later how filesystems are connected together on a Red Hat Enterprise Linux system.

-j Create an ext3 journal inode and filesystem

## Journaling for ext2 filesystems:

### ext3

- ext3 is essentially an ext2 filesystem that uses a journal for file transaction atomicity
- ext3 filesystems can be created natively or easily converted from ext2
- ext3 has three journaling modes:
  - ordered - the default, journals only meta-data
  - journalled - journals data as well as meta-data
  - writeback - journal updates are not atomic, but gives better performance at possible expense of data integrity

*disabled in Red Hat*



Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2894 or +1 (919) 754 3700.

ext3 provides an easy migration path from ext2. Creating the ext3 journal can be done on a mounted partition, as the ext2-to-ext3 filesystem conversion is not destructive.

An ext2 partition can be converted to ext3 with the following steps:

- Change `/etc/fstab` to specify ext3 for desired filesystems
- Create the ext3 journal on the ext2 filesystem(s):  
`tune2fs -j <partition(s)>`
- If the kernel needs to have access to the ext3 module at boot-time, create a new initial ramdisk:  
`mkinitrd /boot/initrd-<kernel version>.img <kernel version>`

Since `/etc/fstab` references ext3, the ext3 and the related `jbd` module will be included in the ramdisk that is created. The last step is to reboot the machine, and verify by `cat /proc/mounts` that ext2 filesystems are now of type ext3

The `data=ordered` journaling mode is default, as it provides the best balance between journal size, data integrity, and `fsck` recover time. The `data=journal` mode, while requiring much larger journals, can speed some database operations. The `data=writeback` mode does not fully guarantee the data integrity and so it is not recommended, but it does allow for a potential speed increase in some cases. While `data=ordered` is the the mode that is preferred for almost all situations, this can be changed by providing a filesystem mount option of `data=<mode>` in `/etc/fstab`.

## Managing Data: mount

```
mount [options] [device] [mount_point]
```

- device (or file system label) points to the filesystem to mount.
- mount\_point is the directory under which the files on the filesystem will be located

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2984 or +1 (919) 754 3700.

As mentioned earlier, in order to access files of an individual filesystem, it must first be connected to the filesystem tree. This is done with the 'mount' command. While graphical interfaces for mounting devices exist, they merely call the command line versions of the programs.

If the mount command is invoked without any arguments, it reads the file `/etc/mtab`, maintained by the system, to display the currently mounted or available filesystems, their mount points and their modes. The mode indicates what operations may be performed on the mounted file system, such as whether the device is mounted for read access or read/write access. Filesystem modes are passed as options to the mount command and the modes available will vary for different types of filesystems. The full syntax for the mount command is:

```
mount [ -t fstype ] [options] device mount_point
```

The filesystem hierarchy is first configured at system installation by the Disk Druid component of the installer. In addition to the size of each partition, Disk Druid also requires information on the type of partition being created and the 'mount point' for each partition. The mount point is the directory on the root partition where a particular partition will be mounted.

fstab

## Managing Data: **mount** options

- t **vfstype** (vfat, ext2, ext3, iso9660, etc.) *not*
- Not normally needed
- o **options**
  - Default options for the ext2/ext3 filesystem:  
*no dev*  
 rw, suid, dev, exec, auto, nouser and async

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat 10

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email training@redhat.com or phone toll free (USA) +1 888 266 2964 or +1 (619) 754 3700.

The **mount** command takes several arguments to indicate the type of filesystem to be mounted and any special options that should be set (such as read-only access.) A useful reference to available options is the man page for **mount**. The filesystem type must precede the options because a given option may be specific to the filesystem type. One or both of the arguments for device and mount point must be used. If only one is used, then an entry in `/etc/fstab` must exist to supply the remaining information. If, as root, you specify options for a filesystem listed in `/etc/fstab`, your options will override all others. Some of the default options for mounting an ext3 filesystem are:

<u>rw</u>	read and write access
<u>suid</u>	suid or sgid file modes honored
<u>dev</u>	device files permitted
<u>exec</u>	permit execution of binaries
auto	honor mount -a (automatic)
nouser	permit the superuser only to mount the filesystem
async	File changes managed asynchronously

Other options commonly used are:

uid=henry, gid=henry	All files of the mounted filesystem are owned by, in this example, henry. The numeric values may also be used, if not preferred.
loop	mount the filesystem using a loopback device. Helpful when mounting a filesystem when it is a file of another filesystem.
user	by contrast to the nouser definition above, this option permits any user to mount--or unmount--the filesystem.



Owner

similar to the `user` option, but in this case the mount request and the device, or "special file," must be owned by the same EUID. On Red Hat Linux systems, the user logged in to the console is made owner of the CD-ROM and floppy device files. Therefore, while at the console, one has exclusive, "easy" access to CD-ROM and floppy removable devices.

For security reasons, both the `owner` and `user` mount options also imply the mount options `noexec`, `nosuid`, and `nodev`, but this can be overridden with the appropriate mount options

cat /proc/mounts

## Managing Data: Unmounting Filesystems

```
umount [options] device | mnt_point
```

- A filesystem "in use" may not be unmounted
  - use fuser to check and/or kill processes
- Use the remount option to change a mounted filesystem's options "atomically"

```
mount -o remount,ro /data
```

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email training@redhat.com or phone toll free (USA) +1 888 286 2994 or +1 (919) 754 3700.

There are a few reasons filesystems must be unmounted or disconnected from the root hierarchy. When the system is brought off-line, is rebooted, or requires filesystem maintenance, or when using removable media, filesystems are disconnected with `umount`. This command references `/etc/mntab` and may be run as `umount -a` (only as superuser) to disconnect all filesystems, or

```
umount [options] device | mount_point
```

To provide operating system stability and protection, a filesystem that is in use (open files, file handles, or a process's CWD, etc.) may not be unmounted. For removable media devices, its device driver will attempt to lock the device. If successful, the lock is removed when the device is quiescent, when unused. This lock, or an ignored `umount` command can be frustrating. Should you experience this, `fuser` will be helpful.

`fuser` is used to display information about the processes using a filesystem. After determining what is acting on the filesystem, `fuser` also provides a convenient way to send signals to those processes. The command can prompt you interactively for each process before sending a signal, or "kill" all processes acting on the filesystem, including a user's CWD entry. To display what (or who) is acting on filesystem:

```
fuser -v mnt_point
```

To kill all actions on a filesystem:

```
fuser -km mnt_point
```

Sometimes you may need to change the options of a mounted filesystem atomically, that is, without other operations occurring during the change. Consider, for example, that you currently have the root filesystem mounted read-only (quite common during recovery operations). You have located a configuration file on the root filesystem that is causing problems, and you want to edit that file. To edit the file, you must mount the filesystem read-write. If the root filesystem is on `/dev/hda5`, the following `mount` command would reinitialize the root filesystem's mount in a single operation (atomically) using the new options:

```
mount -o remount,rw /dev/hda5 /
```

## Managing Data: Filesystem Labels

- Alternate way to refer to devices
- Device independent

```
e2label <special dev file>  
mount [options] LABEL=fslabel mount_point
```

RH RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



12

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2994 or +1 (619) 764 3700.

A potential problem exists when using special device files to point to file systems in that if the device is somehow relocated on the system, then the special device file that points to it will change. This most commonly happens with SCSI devices. Filesystem labels provide an alternate way to reference file systems for mounting that is not dependent on the special device file.

Two mechanisms exist to support filesystem labels. A filesystem label can be written into the superblock of ext2/ext3 filesystems using the `e2label` command:

```
e2label /dev/hda7 dbdisk1
```

Would create a label of 'dbdisk1' on the filesystem on partition /dev/hda7. The command:

```
e2label /dev/hda7
```

will display the current filesystem label for that device. The filesystem can be mounted using the command:

```
mount LABEL=dbdisk1 /mnt/data
```

don't mount boot to prevent kernel  
collision

```
smbclient -L //server1
```

## Managing Data: mount, by example

- Sample filesystem requirements met using options:
  - Disabling execute access
  - Mounting a filesystem image
  - Mounting a pc-compatible filesystem
  - Disabling access time updates
  - Setting up a mount alias

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 285 2904 or +1 (910) 754 3700.

Listed below are a few `mount` commands and their explanation

```
mount -t ext3 -o noexec /dev/hda7 /home
```

For security, users' home directories should be connected denying permission to execute files managed there

```
mount -t iso9660 -o ro,loop /iso/documents.iso /mnt/cdimage
```

Mount the CD-ROM image file `/iso/documents.iso` read-only using the first available loopback device. This "magic" provides access to a filesystem--in this case, organized as a common CD-ROM file structure--that is itself a file of another filesystem

```
mount -t vfat -o uid=515,gid=520 /dev/hdc2 /mnt/projX
```

Mount the `vfat` filesystem located on the `/dev/hdc2` partition so that each file is owned by a specific UID and GID. This is convenient for the end-user, whose UID and GID are as listed, who wishes to manage this filesystem data. Normally, the filesystem's data would be "owned" by root (the superuser), denying direct file manipulation to others.

```
mount -t ext3 -o noatime /dev/hda2 /data
```

Mount the filesystem using the `noatime` option to increase laptop battery up-time by reducing disk access

```
mount --bind /something /anotherthing
```

This command is available starting with the 2.4 kernel. This mounts a directory already mounted on the filesystem on another mount point.

## Managing Data: Connecting Network Resources

- Mounting NFS resources
  - Requires hostname or address of server
  - Requires name of exported directory
- Mounting SMB resources
  - Requires hostname or address of server
  - Requires share name
  - May require username and password

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (910) 754 3700.

Many filesystem resources are connected through networked systems. Most common are Network File System (NFS) and SMB resources. To access these filesystems, first the network resource must be known, by host and share name. To discover what filesystems are exported by a remote system, use the following commands.

For NFS:

```
showmount -e remote_server
```

For SMB:

```
smbclient -L remote_server -N
```

When the host and share names are known, the following commands are used to connect the network filesystem to the local filesystem tree.

For NFS:

```
mount remote_server:/shared/dir /mnt/remote_nfs
```

For SMB:

```
mount //remote_server/share /mnt/remote_samba
```

Note that the `-t fstype` option is omitted: the syntax for the device is recognized by `mount`, and the specific `mount` command is executed. The NFS method is "built into" the operating system. For SMB, `mount` passes the request to `/usr/sbin/smbmount`.

Options for these filesystem types include:

For NFS:

```
bg, fg, intr, nointr, soft, hard, rsize and wsize
```

For SMB:

```
username=<arg>, password=<arg>, ip=<arg>, uid=<arg>, and gid=<arg>
```

See the man page for `mount` and `smbmount` for NFS and SMB options, respectively.

Zimbra

xpae

## Managing Data: /etc/fstab

- Configuration of the filesystem hierarchy
- Used by **mount**, **fsck**, and other programs
- Maintains the hierarchy between system reboots
- May use filesystem volume labels in the device field

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

15

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2884 or +1 (919) 754 3700.

/etc/fstab is referenced each time the system boots to create the desired filesystem hierarchy. It consists of six fields per line for each filesystem to be connected to the "tree" as follows:

```
# device          mount_point  FS_type  options  dump_freq  fsck_order
LABEL=/mnt/data  /mnt/data   ext3     defaults 0 0
```

*server: /var/ftp/pub /mnt/Server nfs res2c32760,soft,nar,lg 00*

- device: The special device file name, or filesystem label of the device to mount
- mount\_point: The path used to access the filesystem
- FS\_type: The filesystem type
- options: A comma-separated list of options
- dump\_freq: Level 0 dump frequency: 1=daily, 2=every other day, etc; 0=never dump
- fsck\_order: 0 = ignore, 1 = first (the root filesystem should have this value), 2-9 = second, third etc : filesystems that have the same number greater than 1 are checked in parallel. Network filesystems and CD-ROMs should be ignored.

During system initialization (see /etc/rc.d/rc.sysinit), /etc/fstab is used to create the filesystem hierarchy. Entries are parsed and used as arguments to mount if, among their options, noauto is not present. At best, /etc/fstab is a "wish list". If a filesystem is not available after reaching the desired runlevel (discussed later) run mount without options to display what is available. Floppy and CD-ROM entries typically have noauto as an option.

Using /etc/fstab after system initialization, an example: If the partition /dev/hda5 contained a filesystem labeled /mnt/data and the directory /mnt/data exists (and is already available!), then the following mount commands will connect this filesystem to the filesystem tree, using /etc/fstab to provide specifics:

```
mount /dev/hda5
mount -L /mnt/data
mount LABEL=/mnt/data
mount /mnt/data
```

## Managing Data: The Auto-Mounter

- System administrator specifies mount points to be controlled by the `automounter` daemon process
- The `automounter` monitors access to these directories and mounts the filesystem on request
- Filesystems automatically unmounted after a specified interval of inactivity
- Enable `/etc/auto.net` to “browse” all NFS exports on the network

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2994 or +1 (010) 754 3700.

By default, all mounted file systems are owned by the root account and can only be unmounted by the root account. This behavior can be overridden through the use of the `owner`, `user` or `users` options in `/etc/fstab`. This is normally done, for example, for removable media devices so that a non-privileged user can access the contents of a floppy disk or cdrom.

In addition to removable media devices, users may often need access to files located elsewhere on the network. This access can be obtained by mounting a network file share. Whether a network file share or some type of removable media device, one drawback is that casual users must learn the syntax of the `mount` and `umount` commands. The `automounter` can be configured to monitor certain directories and automatically mount the appropriate devices when a reference is made to files in that directory. The `automount` daemon is provided by the `autofs` RPM.

To control filesystems with `automount`, modify the supplied `/etc/auto.master`. `/etc/rc.d/init.d/autofs` parses this file and launches the `automount` daemon based on this configuration. Each line of `auto.master` lists a directory, present in the hierarchy, and a reference to yet another file that further defines specific mount options for those base mount points.

Referred from an `auto.master` entry, an `autofs` mount configuration file lists the filesystems to be mounted under the directory, including options required. Sample syntax can be found in the `autofs(5)` man page.


Once a mount point has been configured for `automount`, and the daemon started, merely requesting the filesystem (i.e., `cd`, or other action) completes the mount. The daemon maintains the connection as long as the filesystem is in use plus an interval of time. The default interval is 60 seconds.

## ext2/ext3 Filesystem Attributes

- ext2 and ext3 support attributes that affect the manipulation of file data
  - `lsattr` displays file attributes
  - `chattr` changes file attributes
    - Some attributes are not currently supported by the Linux kernel

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2994 or +1 (919) 754 3700.

In addition to the more common, UNIX-derived modes of a file (a.k.a. permissions), the ext2 filesystem structure provides for a few data controls, or file attributes. To set file attributes, use:

```
chattr +|-|=attribute[attribute...] file [file...]
```

The following attributes may be set:

- A when file is modified, its atime record is not modified.
- a the file may only be opened in append mode. The file may not be deleted. Only the superuser can set or clear this attribute.
- d the file is skipped for backup by dump.
- i the file is immutable. It cannot be deleted or renamed, nor linked to any other name. No data can be written to the file. Only the superuser can change the immutable attribute.
- j the file's data as well as the meta-data is written to the ext3 journal, even if the ext3 filesystem is mounted with the `data=ordered` or `data=writeback` option.
- S when the file is modified, the changes are written synchronously to the filesystem; this is equivalent to the mount option `sync`.

The attribute operators used with the `chattr` command are not typical to most commands; use `+` to set an attribute and `-` to unset it. The `=` operator will set absolute values. Attribute values not currently supported, as indicated in the man page, may be in the future.



## Virtual Memory

- Swap space is a supplement to system RAM
- Basic setup involves:
  - Create swap partition or file
  - Write special signature using `mkswap`
  - Add appropriate entries to `/etc/fstab`
  - Activate swap space with `swapon -a`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 296 2994 or +1 (619) 754 3700.

### Setting up a swap partition

Use a partitioning program to add a partition. Set the partition id type to `0x82`. Create the signature needed on the partition using `mkswap`:

```
mkswap -v1 /dev/hda6
```

Add an entry for the swap to `/etc/fstab`. It will look similar to the following:

```
/dev/hda6 swap swap defaults 0 0
```

Activate the swap partition using `swapon -a` (which reads `/etc/fstab` and turns on all swap entries it lists).

Check the swap partition's status using `swapon -s`

### Setting up a swap file

Use the following to create a file, where count X defines the file size in kilobyte blocks:

```
dd if=/dev/zero of=swapfile bs=1024 count=X
```

Run `mkswap` to create the signature. The swapfile can also be activated manually with `swapon`, or you can initialize it in a startup script such as `/etc/rc.d/rc.local`.

# Filesystem Maintenance

- Maintaining consistency with fsck
- Filesystems checked at boot up
- `sulogin` session started if errors are severe
- `lost+found`

`e2fsck -f0 /dev/hda2`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat 19

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## System initialization and fsck

Filesystems are susceptible to corruption if the system is shut down improperly. When the system boots, `rc.sysinit` runs `fsck` on any filesystems marked for checking in `/etc/fstab`. If any of these filesystems are marked as "dirty" or have data in the journal, `fsck` will attempt to repair them. If `fsck` succeeds, the filesystems will be mounted and the boot process will continue. If `fsck` fails, `rc.sysinit` will run `sulogin` and will report that `fsck` needs to be run manually. You should only run `fsck` on a filesystem that is unmounted or mounted read-only.

### e2fsck

Like `mkfs`, `fsck` is a front end for a utility specific to the filesystem type. If a type is not specified, an `ext2` filesystem is assumed and `e2fsck` is invoked. Options not recognized by `fsck` are passed on to this utility. See the man page for options to `e2fsck`, such as `-f`, which will force a check even if the filesystem seems clean.

It is important to note that `fsck` can find and repair problems with the filesystem's structure, which may include unlinking datablocks. This does not corrupt data, nor destroy the unlinked datablocks on the filesystem, but it may seem so because the file will not appear in the filesystem hierarchy. Check the directory named `lost+found` in the root of the filesystem; your data might be relinked there, named after its old inode number.

## Filesystem Maintenance (cont.)

- `tune2fs`
- `dumpe2fs`
- `debugfs`
- `parted`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (619) 754 3700.

Use `tune2fs` to modify many filesystem attributes. Examples include: modifying the percentage of reserved blocks, or setting the maximum count before a check by `e2fsck` is forced. The default count value (20), much like the consistency value, is reserved in the filesystem structure. After mounting the filesystem with a count equal to this value, `e2fsck` is run regardless of the filesystem's consistency. Mount count and frequency checking is usually disabled with ext3 filesystems.

`dumpe2fs`, as the name suggests, provides a "dump" of filesystem information to standard out--by default, the console. Saving the output of this command to a file may be useful as a "written record" of the filesystem.

`debugfs` is an interactive utility to examine and debug an ext2 filesystem. `debugfs` can be used to manually verify inode integrity, or as an aid towards recovering data.

`ext2online` can be used to grow the size of an ext2 or ext3 filesystem. If you want to make a file system larger, you must first make the partition larger with `fdisk` or another appropriate tool, then grow the file system with `ext2online`.

Because ext3 is based on ext2, these utilities will work on ext3 filesystems, too.

## Adding a Drive

- Physically connect the new drive
- Create partitions
- If required, reread partition table with `partprobe`
  - verify with `fdisk -l` and `cat /proc/partitions`
- Create filesystems for new partitions, or
  - write signature to new swap partitions
- Optionally create disk label
- Create any needed mount points
- Add new entries to `/etc/fstab`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2894 or +1 (313) 754 3709.

To add a disk drive to a system, first shut down the system, power it off and attach the new drive. For IDE drives, be sure to set the drive as master or slave as appropriate. For SCSI drives, you must select an unused SCSI id for the new device and ensure proper termination of the SCSI bus to which the drive will be attached. When the system is powered up, watch the output from the kernel during its initialization. If you don't see any references to the new drive, check `/var/log/mesg` once the system has booted. If the drive doesn't show up there, try restarting the system and checking the system BIOS: the drive may not be recognized there.

Once recognized, run `fdisk` or one of its variants to create the partitions you need. If the partition is going to be a swap partition, change the partition id type to `0x82`.

Once the partition table on disk is modified, it may be necessary to also update the in-memory copy of the partition table. To do this, run the `partprobe` command.

Use `mkfs` to create filesystems on each of your new, non-swap, partitions. Swap partitions are "marked" with `mkswap`. Take into consideration the intended use of each filesystem and make any appropriate changes to your `mkfs` command for tuning purposes. If you intend to use filesystem labels to mount the filesystem later, specify them using the `-L` option. Another way to label a disk is the `e2label` command.

Create any needed mount points in your current filesystem hierarchy. Keep in mind that directories used as mount points need not be empty, but any files in the directory are temporarily unavailable when a filesystem is mounted on that mount point.

Add entries for the new filesystem to `/etc/fstab`. Check these entries with `mount` manually before you reboot. This will not only make filesystem management simpler, but will call the system initialization scripts to mount, check and provide information to other utilities, like `dump`.

## End of Unit 4

- Questions and answers
- Summary
  - What tools are available for partitioning?
  - What two ways can swap space be implemented?

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2664 or +1 (910) 764 3700.

### Important files covered in this Unit:

```
/etc/fstab
/etc/mtab
/etc/auto.master
/etc/auto.misc
/proc/partitions
```

### Important commands covered in this Unit:

```
fdisk, sfdisk, parted, partprobe
e2label, devlabel
fuser
mkfs
mount
umount
showmount
smbclient
fsck
lsattr
chattr
mkswap
swapon
swapoff
```

5676 232

~~26500~~ 123

793123

1608 139

mke2fs -b 2K -i 4K

2068 4096



# Lab 4

## Filesystem Management

---

**Goal:** To build skills and knowledge related to filesystems

### Sequence 1: Creating and Mounting File Systems

#### Tasks:

1. Use `fdisk -l` to locate information about the partition sizes on `/dev/hda`. Use this information to calculate the amount of unpartitioned space on the hard drive.
2. Now use `fdisk` to add a new logical partition that is 1GB to 2GB in size. (Make sure to write the changes to disk using the "w" command.) What device is the new partition: `/dev/hda__`? Why?
3. Reboot to reread the revised partition table or use `partprobe` to refresh the kernel's view of the partition table.
4. Using `mke2fs`, make a new ext2 file system on the new logical partition you just created. Try creating the ext2 filesystem with 2k blocks and one inode per every 4k (two blocks) of filesystem. You may need to consult the man page for `mke2fs` for these options.
5. Create the directory `/data` which will be the mount point for the new file system.
6. Use the `mount` command to mount the new file system on `/data`. Copy `/etc/passwd` into `/data` and verify that the copy was successful.
7. `umount /data`
8. Add a label to the new partition using `e2label`:

```
e2label /dev/hdax /data (where x is the number of the newly-created partition)
```

You can check the labels of this and any other partitions with `e2label` as well, by specifying the partition as the only argument.

9. Add a line to `/etc/fstab` to mount the new file system on `/data`. Add the following line to use the label you just created:

```
LABEL=/data /data ext2 defaults 1 2
```

You could use the following line instead:

```
/dev/hdax /data ext2 defaults 1 2
```

Both lines produce the same results in this instance. However, if you were to move the IDE drive to another channel or make it the slave instead of the master, using the label in `/etc/fstab` would enable the system to locate the partition and mount it regardless of its device designation

10. Mount the new file system:

```
mount /data
```

11. Copy files into this new file system to make it 75% to 85% full. Use the `df` command to determine how full the file system is. You may use `du` to determine which directories are large enough to be suitable sources.

Examples:

```
df -hT /data
```

```
du -h --max-depth=2 /
```



**Sequence 2: Converting ext2 to ext3**

- 1 Type `sync`. This will flush any information in disk buffers to the disk. This is normally done periodically, but the next step may preempt this automatic syncing.
- 2 **Crash** your computer by running the command `reboot -f`, or by turning the machine off by holding the power button until it powers off then turn it back on (never a good idea). Observe the boot process and the time it takes to fsck the ext2 filesystem.
- 3 If you are presented with a "Repair filesystem", prompt repair the problem filesystems as necessary with `e2fsck /dev/<partition>`. Otherwise, go on to the next step.
- 4 After successful boot, convert the ext2 filesystem that you created to ext3 by creating a journaling inode. Since ext3 has great data integrity, and much greater filesystem integrity guarantees, turn off the automatic per-mount and per-time-period filesystem checking.

```
tune2fs -j -c 0 -i 0 /dev/<partition>
```

5. Examine the filesystem's characteristics:

```
tune2fs -l /dev/<partition>
```

6. Edit the line pertaining to the `/data` filesystem in `/etc/fstab` to use ext3, not ext2, as the filesystem type.
- 7 Unmount and remount the filesystem as type ext3, and verify that it's mounted as ext3:

```
umount /data; mount /data  
df -T /data
```

8. Make sure that your initial ramdisk in `/boot` contains the needed ext3 and jbd kernel modules. If `/data` is the first filesystem on your machine to use ext3, `initrd` almost certainly does not contain these modules. This issue is only important if the root filesystem needs ext3 capability that is not in `initrd`... but let's proceed as if this is the case. Make, or remake, the `/boot/initrd-<version>.img` file:

```
mkinitrd -f -v /boot/initrd-$(uname -r).img $(uname -r)
```

8. Type `sync`, and then **crash** your system again using the `reboot -f` command or power button
9. Observe what happens during the boot process. Which filesystems are checked? Do you see the "recovering journal" message for the `/data` filesystem? How much faster is recovering from an improper shutdown with ext3 than fsck'ing with ext2?

**Sequence 3: Automounting data with autofs**

1. Ensure that `iptables` firewalling is disabled:

```
service iptables stop
chkconfig iptables off
```

2. Edit the `/etc/auto.master` file. Uncomment the line for `/misc`.

3. Add a line to the `/etc/auto.misc` file that will mount the `/var/ftp/pub` export from `server1.example.com` to your own `/misc/server1` target. Use the `ftp.example.org` line as an example of how to accomplish an `nfs` mount using the automounter ( although the hostname says 'ftp' this is an `nfs` example, and presently the linux kernel does not support mounting `ftp` resources without an unsupported kernel patch ).

4. Restart the `autofs` service: `service autofs restart`

5. Try to use the `/misc/server1` directory (eg. `cd /misc/server1`)

# UNIT 5

## Network Configuration

RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (019) 754 3700.

## UNIT 5: Objectives

- Upon completion of this unit you should be able to:
  - Understand network device recognition
  - Know how to configure network interfaces
  - Use network configuration utilities
  - Understand IP aliases
  - Understand IP route configuration
  - Know how to configure client-side DNS

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2894 or +1 (919) 754 3700.

## UNIT 5: Agenda

- Network device recognition
- Network interfaces
- Network configuration utilities
- IP aliases
- IP route configuration
- client-side DNS

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 256 2594 or +1 (919) 754-3700.

# Device Recognition

- All drivers for network interface cards are built as modules
- Networking scripts reference logical interface names, eg:  
eth0
- `/etc/modprobe.conf` maps logical names to specific module name
- Example:  
alias eth0 3c59x

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



4

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

While you can compile specific ethernet card driver code into the kernel statically, Red Hat compiles network card drivers as kernel modules for easy adaptability to any hardware configuration. Network interface modules are loaded at boot time if networking has been enabled. The appropriate module is loaded based on a alias line in `/etc/modprobe.conf`.

Below is an example of the contents of `modprobe.conf`:

```
alias eth0 3c59x
alias parport_lowlevel parport_pc
alias sound-slot-0 es1371
alias usb-controller usb-uhci
```

The logical ethernet interface name (i.e. `eth0`, `eth1`, etc) is used in the configuration file and scripts to associate kernel driver module with a specific interface. If you have an ISA network card, you can specify options for each card by its IRQ and/or I/O address:

```
alias eth0 3c509
alias eth1 tulip
options 3c509 io=0x210
```

Documentation for networking module options is in

```
/usr/src/linux-2.6/Documentation/networking/net-modules.txt
```

# Network Interfaces

- Interface names
  - Ethernet: `eth0, eth1, ethN`
  - Token Ring: `tr0, tr1, trN`
  - FDDI: `fddi0, fddi1, fddiN`
  - PPP: `ppp0, ppp1, pppN`
- Data link layer addresses
  - `ifconfig`



*LIBRE DIGITAL DISTRIBUTED DATA INTERFACE*

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 285 2294 or +1 (810) 724-3700.

## Interface names

The Linux kernel names interfaces with a specific prefix depending on the type of interface. For example, all ethernet interfaces start with `eth`. Notice that regardless of the specific hardware vendor, the interfaces start with a common prefix. Following the prefix, each interface is numbered, starting at zero. For example, `eth0, eth1, eth2` would refer to the first, second and third ethernet interface.

## Layer 2 hardware addresses

The hardware address of network interfaces can be determined by running the `ifconfig` command. Another method is to examine the output from device driver (kernel module) as it loads. Check the output of the `dmesg` command and/or `/var/log/dmesg`.

*arpstar, - procedure for ARP  
ethereal*

## mii-tool

- Views and controls the negotiated media speed (100baseTX, 10baseT) of some ethernet cards.
- Useful for forcing specific ethernet speed and duplex settings
- Changes with `mii-tool` should be made on inactive interfaces

*ethtool*

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2094 or +1 (919) 754 3700.

`mii-tool` allows a system administrator to view, monitor, log, and change the negotiated speed of ethernet network cards. This capability depends upon the card having a chipset compatible with the `mii-tool` and not all cards are compatible.

```
$ mii-tool -v
eth0: negotiated 100baseIx-FD, link ok
  product info: Level One LXI970/971 rev 0
  basic mode:   autonegotiation enabled
  basic status: autonegotiation complete, link ok
  capabilities: 100baseTx-FD 100baseIx-HD 10baseT-FD 10baseI-HD
  advertising: 100baseTx-FD 100baseIx-HD 10baseT-FD 10baseI-HD
  link partner: 100baseTx-FD 100baseTx-HD 10baseT-FD 10baseI-HD
```

To force 100Mbps full duplex operation of eth1:

```
$ ifdown eth1
$ mii-tool -v --force 100baseTx-FD eth1
$ ifup eth1
```

To restart autonegotiation :

```
$ ifdown eth1
$ mii-tool -vr eth1
$ ifup eth1
```



# ifconfig

- Used to configure and set IP addresses on network interfaces
  - Not usually called directly, but by other scripts
- Also used to view properties of active and inactive network interfaces

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2994 or +1 (910) 754 3700.

When `ifconfig` is run with no arguments, it displays information on all currently active interfaces. The information displayed includes the type of interface, IP address information, hardware addresses, various statistics, and hardware resources.

```
[root@localhost /tmp]# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:A0:CC:37:38:88
          inet addr:10.100.0.1  Bcast:10.100.0.255  Mask:255.255.255.0
          UP BROADCAST NOTRAILERS RUNNING MTU:1500  Metric:1
          RX packets:151110 errors:0 dropped:0 overruns:0 frame:0
          TX packets:88699 errors:1 dropped:0 overruns:0 carrier:2
          collisions:33100
          RX bytes:7591398 (7.2 Mb) TX bytes:173492 (169.4 Kb)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING MTU:3924  Metric:1
          RX packets:166 errors:0 dropped:0 overruns:0 frame:0
          TX packets:166 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0
          RX bytes:51066 (499.0 Kb) TX bytes:51066 (499.0 Kb)
```

In order to view inactive interfaces, use the `-a` option.

# ifup / ifdown

- `if(up|down)` interface
- Start and stop network interfaces
- Take care of details specific to interface
  - Changing/adding/deleting routes
  - Obtains addresses as needed
    - BOOTP, DHCP

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



8

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2894 or +1 (919) 754 3700.

## Bringing up and down a network interface

The process of bringing up a network interface involves several possibilities and options. It isn't as simple as just configuring an IP address. For example, what IP address should the interface be configured with? Should DHCP or BOOTP be used? If the interface is a PPP interface associated with a modem, the modem needs to be instructed to dial, and the `pppd` daemon started.

Another important task that is performed with bringing up an interface is the addition, deletion, or changing of routes in the routing table.

The `ifup` and `ifdown` scripts take care of all the extra tasks that need to be performed when activating and deactivating a network interface.

## Interface Configuration Files

- *eth0*  
ifcfg-xxx
- Located in:
  - /etc/sysconfig/network-scripts/
- Configuration methods
  - static
  - dhcp
  - bootp

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 288 2994 or +1 (919) 754 3700.

### Interface Configuration Files

Red Hat Enterprise Linux stores network interface configuration information in files in the directory `/etc/sysconfig/network-scripts`. The file names are prefixed with `ifcfg-` and then the name of the interface. For example, the file for the first ethernet interface would be `ifcfg-eth0`.

For example, to put the `eth0` interface under `dhcp` control and have it activated at boot time, your interface config file would need these contents:

```
DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
```

For an interface with static IP address configuration, the following is needed:

```
DEVICE=eth0
IPADDR=xxx.xxx.xxx.xxx
NETMASK=xxx.xxx.xxx.xxx
BOOTPROTO=static
ONBOOT=yes
```

# Configuration Utilities

- **netconfig**
  - Text-based network configuration tool
  - Only writes config files. Does not activate device or changes. Use `ifup/ifdown` to activate changes
  - Used by `kudzu` when new network card found at boot time
- **system-config-network**
  - GNOME-based network configuration tool
  - Can be launched by a non-privileged user, but requires authentication as root

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



10

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2994 or +1 (319) 764 3700.

`netconfig` is a curses-based tool that is used to configure network interfaces, either as a DHCP client or with a static IP address, nameserver, and gateway. By default it modifies the settings for the first ethernet interface (`eth0`), but the "`--device interface`" argument can be used to set up other network interfaces:

```
netconfig --device eth2
```

The Red Hat Network Administration Tool is a X-based utility that can be used to set up Ethernet, ISDN, PPP, xDSL, token ring, CIPE, or wireless network interfaces. It can be started from the command line with `system-config-network` command, or by selecting the "Network Configuration" tool through the panel menu in GNOME or KDE.

Note that using `system-config-network` creates an alternate file hierarchy under `/etc/sysconfig/networking`. Modifying an interface with this tool will create a hard link between the files in `/etc/sysconfig/network-scripts` and in `/etc/sysconfig/networking/profiles/<profile>`. Both files will have the same name. (e.g. `ifcfg-eth0`)

## Binding multiple IP addresses

- Use multiple IP addresses on a NIC
  - Virtual Interface(s)
- For a small number of IPs, create an ifcfg file for each virtual interface

```
ifcfg-ethX:xxx
```
- For a large number of IPs, create an ifcfg range file

```
ifcfg-ethX-rangeX
```

Rev 1/11/03/RHEL4-1

Copyright © 2003 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 238 2694 or +1 (919) 754 3700.

### Virtual Interfaces

Linux is commonly deployed in the role of web or ftp server. Often a situation calls for many web sites or ftp sites running on the same server. This is called virtual hosting. If you plan on supporting SSL or ftp for different servers, then you are required to use a separate IP address for each server. Note that non-SSL web hosting requires only one IP address.

If you have many IP addresses to bind to a single network card, it is more convenient to use an interface configuration range file than to create many separate interface configuration files. For example the file `ifcfg-eth0-range0` with the following contents would bind the IP addresses 10.100.13.75 through 10.100.13.210 to the `eth0` network card:

```
IPADDR_START=10.100.13.75
IPADDR_END=10.100.13.210
CLONENUM_START=0
```

You can have multiple range files, just be sure that you don't create conflicts. Currently each range file is limited to addresses within the same class C-sized block. A maximum of 256 IP addresses may be bound to a single NIC.

## DHCP / BOOTP

- The `dhclient` daemon manages client-side DHCP and BOOTP
  - For DHCP, `dhclient`:
    - Obtains a lease
    - Performs automatic lease renewal
  - Normally run by `ifup/ifdown`
  - Can be run manually to force renewal or release of a lease

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

12

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 286 2984 or +1 (919) 754 3700.

`dhclient` stores interface configuration in

`/var/lib/dhcp/dhclient-ethX.leases`

Below is a example of a typical file:

```
lease {
  interface "eth1";
  fixed-address 192.168.0.6;
  filename "/kickstart/workstation.cfg";
  option subnet-mask 255.255.255.0;
  option routers 192.168.0.254;
  option dhcp-lease-time 21600;
  option dhcp-message-type 5;
  option domain-name-servers 192.168.0.254;
  option dhcp-server-identifier 192.168.0.254;
  option domain-name "example.com";
  renew 1 2003/10/6 01:16:30;
  rebind 1 2003/10/6 03:49:38;
  expire 1 2003/10/6 04:34:38;
}
```

# Global Network Parameters

- `/etc/sysconfig/network`  
NETWORKING=yes|no  
HOSTNAME=<fqdn by default>  
GATEWAY=<gateway IP>  
NISDOMAIN=<nis domain name>

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## Global network parameters

`/etc/sysconfig/network:`

```
NETWORKING=yes|no  
HOSTNAME=<fqdn by default, but whatever hostname you want>  
GATEWAY=<gateway IP>  
NISDOMAIN=<nis domain name>
```

## Default Route

- Global default defined in:
  - `/etc/sysconfig/network`  
`GATEWAY=xxx . xxx . xxx . xxx`
- Default gateway can also be defined in
  - `/etc/sysconfig/network-scripts/ifcfg-XXX`
  - `ifcfg-xxx` default overrides Global default routes
  - `GATEWAY=xxx . xxx . xxx . xxx`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



14

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (919) 754 3700.

### The Default Route

The default route specifies where IP packets should be sent that are destined for networks that your machine doesn't know how to reach

With RHEL, if you don't set a system global default route in `/etc/sysconfig/network`, you can define default routes associated to specific interfaces. When such an interface is activated, that default route is set in the machine's routing table.



## Static Routes

- Connected networks
  - Linux kernel automatically creates a network route for connected networks
- Static routes defined per interface
  - `/etc/sysconfig/network-scripts/route-eth0`
  - `/etc/sysconfig/networking/devices/eth0.route`
- Display with:
  - `route -n`
  - `netstat -rn`

*GRAPHICAL*

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2664 or +1 (619) 754 3700.

Static routes are set to activate on a per-interface basis. Either of two files may be used to set a static route

Lines in the `/etc/sysconfig/network-scripts/route-(ifname)` file for the interface use the same syntax as the 'ip route add' command:

IP/CIDR via GATEWAY

*class LESS 187ish Domain looking*

For example:

192.168.2.0/24 via 192.168.0.128

*= 255.255.255.0*

*gIP/ipcalc*

adds a static route to the 192.168.2 network through the 192.168.0.128 router.

The `system-config-network` command uses a different file,

`/etc/sysconfig/networking/devices/(ifname).route`, to set static routes. This file uses a different syntax:

```
ADDRESS0=192.168.2.0
NETMASK0=255.255.255.0
GATEWAY0=192.168.0.128
```

The second static route uses ADDRESS1/NETMASK1/GATEWAY1, and so on.

# Name Resolution

*- is /etc/hosts IP add*

- **hostname** - display or set the system's name
  - **is initially set by `rc.sysinit` from `$HOSTNAME` variable**
  - **`/etc/sysconfig/network`**
- **`/etc/hosts`** - local database of hostname to IP address mappings
  - Checked before DNS
  - Useful for small isolated networks

Rev 0123 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

16

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 299 2994 or +1 (919) 254 3700.

## Local name resolution

If you have a small network of computers, you can eliminate DNS lookups for communication between the hosts by creating a standard `/etc/hosts` file and using it on all your machines. Once you grow beyond a handful of computers, DNS provides better scalability and manageability. At a minimum, your `/etc/hosts` should contain `localhost` and the IP address for your ethernet interface

A typical `/etc/hosts` file:

```
[root@station1 /root]# cat /etc/hosts
127.0.0.1          localhost.localdomain localhost
10.100.0.1        station1.example.com station1
```

The file `/etc/host.conf` sets the order of name resolution for the `localhost` (whether to check the `/etc/hosts` file or query the name server first):

```
[root@dawg /root]$ cat /etc/host.conf
order hosts,bind
```

# DNS client configuration

- `/etc/resolv.conf`
  - Defines which name servers to use
  - Servers are checked in order listed

Rev RH133RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2994 or +1 (919) 754 3700.

## Client-side DNS configuration

Domain name service is responsible for associating domain names (i.e. `joe.somewhere.com`) with IP addresses (i.e., `192.168.1.24`). A typical `/etc/resolv.conf` file looks like the following:

```
search somewhere.com dyn.somewhere.org
nameserver 192.168.1.2
nameserver 192.168.1.3
```

Nam servers are checked in order. The first name server listed will be queried and if it is unavailable the second name server will be queried and so on down the list.

The `search` specification allows up to six domains to be searched when the system is attempting to resolve a hostname that is not fully qualified.

See the `resolv.conf(5)` man page for information about other options you can include in this file.

## DNS Utilities

- Useful utilities in `bind-utils` RPM package include:

- `host`: gather host/domain information

```
host ns1.redhat.com
```

```
host -a redhat.com
```

- `dig`: send queries to name server directly

```
dig @ns1.redhat.com mx redhat.com
```

- `nslookup`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



18

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2264 or +1 (310) 754 3709.

*host -t MX Google.com*

# Network Diagnostics

- ping
  - Network packet loss and latency measurement tool
- traceroute, mtr
  - Displays network path to a destination
- netstat *-t qn*
  - Multi-purpose network information tool

*Lab  
CLI  
PLD*



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 285 2994 or +1 (619) 754 3700.

## ping

`ping` attempts to test network connectivity by sending ICMP packets to a specific system on the network. If the remote system is accessible via the network, it will reply.

The default behavior of `ping` is to send a 64 byte ICMP packet to the specified host every second until you cancel the operation with `Ctrl-C`. When the `ping` operation is canceled, `ping` will report summary statistics such as average packet loss, number of packets sent/received, etc.

```
$ ping www.redhat.com
PING www.portal.redhat.com (206.132.41.202) from 10.100.0.1 : 56(84) bytes of
data:
64 bytes from 206.132.41.202: icmp_seq=0 ttl=242 time=100.760 msec
64 bytes from 206.132.41.202: icmp_seq=1 ttl=242 time=90.170 msec
64 bytes from 206.132.41.202: icmp_seq=2 ttl=242 time=90.708 msec
64 bytes from 206.132.41.202: icmp_seq=3 ttl=242 time=94.312 msec
--- www.portal.redhat.com ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/mdev = 90.170/93.987/100.760/4.233 ms
```

See the `ping` man page for command-line options to modify its operation.

## traceroute

As network traffic travels across the Internet, it usually travels through multiple routers. When connectivity between a local system and a remote system is sluggish and inconsistent, it is useful to investigate which router is responsible for the network problem.

The `traceroute` command will attempt to show the path of routers network packets take between the local system and a remote system. This command by default uses UDP not ICMP. See the man page for details.

## End of Unit 5

- Questions and answers
- Summary
  - Where are drivers aliased to specific interfaces?
  - Where is the default route set?
  - What file is used for client-side DNS configuration?

Rev RH133/RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

### Important files covered in this Unit:

```
/etc/modules.conf
/etc/sysconfig/network
/etc/sysconfig/network-scripts/ifcfg-*
/etc/sysconfig/static-routes
/etc/hosts
/etc/host.conf
/etc/resolv.conf
```

### Important commands covered in this Unit:

```
ifconfig
ifup
ifdown
dhclient
system-config-network
netconfig
hostname
ping
traceroute
mtr
route
netstat
host
dig
```

# Unit 5 Lab

## Static Network Settings

---

Estimated Duration: 1/4 hour

**Goal:** To build skills needed to manually configure networking

**Setup at Start:** A Red Hat Enterprise Linux System using DHCP networking

**Situation:** The DHCP server is down! You need to get your workstation up on the network, so you will edit the appropriate configuration files by hand to set up static networking

**Instructor:** TURN OFF THE DHCPD SERVICE ON SERVER1.

**Sequence 1: Setting the IP address****Scenario/Story:**

The DHCP server on your network is down (Your instructor will turn it off.) You need to set up a static IP address so that you can get your workstation back on the network

**Tasks:**

1. Begin by shutting down your ethernet interface with the `ifdown` command:

```
ifdown eth0
```

2. Open `/etc/sysconfig/network-scripts/ifcfg-eth0` in a text editor and change the contents to match the following (*where X is replaced with your station number*):

```
DEVICE=eth0
BOOTPROTO=none
ONBOOT=yes
IPADDR=192.168.0.X
NETMASK=255.255.255.0
GATEWAY=192.168.0.254
```

3. View the contents of `/etc/resolv.conf`. It should still have the valid settings obtained from the DHCP server. If not, make sure it matches the following:

```
search example.com
nameserver 192.168.0.254
```

4. Bring up your newly-configured interface with `ifup`:

```
ifup eth0
```

5. Verify your network settings by pinging `server1`.
6. Reboot the machine and again verify your network settings by pinging `server1`.

**Deliverable:**

A system configured to operate with static network settings.

**Clean up:**

Once your instructor has turned DHCP back on, return your configuration files to their original state and bring `eth0` down and back up again. The `ifcfg-eth0` file should once again read:

```
DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
```



# UNIT 6

## RPM and Kickstart

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (919) 754 3700.

*Handwritten signature*

*Handwritten signature*

## UNIT 6: Objectives

- Upon completion of this unit the student should be able to:
  - Use RPM to install, remove, update, and query packages
  - Configure Kickstart and perform automated installations

Rev/RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (918) 754 3700.

## UNIT 6: Agenda

- Using RPM
- Deploying an installation server
- Installation using Kickstart

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 259 2294 or +1 (813) 754-3700.

# The RPM Way

- Package installation is never interactive
- Applies to all software (core OS and add-ons)
- No such thing as a patch to a package

Rev 701133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



4

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2984 or +1 (919) 754 3700.

**Package installation is never interactive** In contrast to package management on some other platforms, RPM's design does not provide interactive configuration of software as part of the package load process. RPM can perform configuration actions as part of the installation, but these are scripted not interactive. It is common for packages to install with reasonable default configurations applying. On the other hand some software installs in an unconfigured state

**Applies to all software** On some other common platforms, the package management system applies only to part of the installed software. The scope of RPM include core operating systems as well as services and applications. It is common and desirable to run Red Hat Enterprise Linux systems such that **all** software installed falls under the management of RPM. This is a great aid for management and configuration control, since one framework applies for the management of every installed file.

**No such thing as a patch.** It is common on other platforms to have operating system updates released as software objects (eg "service packs") which represent incremental changes to a large number of installed component packages. RPM never does this. If part of any given software package is changed as part of an errata or bug fix, then that entire package will be re-released in its entirety at a new version. The implications are that the installed state of an RPM-managed system can be described as the version number of all the installed components

# RPM Package Manager

- RPM Components

- local database
- rpm and related executables
- package files


- Primary Functions

- install/remove
- query
- verify
- build

*main function*

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



5

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (919) 754 3700.

The RPM Package Manager greatly simplifies the distribution, installation, upgrading, and removal of software on Red Hat Enterprise Linux (RHEL) systems. The RPM system consists of a local database, the rpm executable, rpm package files

The local RPM database is maintained in `/var/lib/rpm`. The database stores information about installed packages such as file attributes and package prerequisites. An administrator rarely, if ever, modifies the database directly, but instead uses the rpm command.

Software to be installed using rpm is distributed through rpm package files, which are essentially compressed archives of files and associated dependency information. Package files are named using the following format:

`name-version-release.architecture.rpm`

The *version* refers to the open source version of the project, while the *release* refers to Red Hat internal patches to the open source code.

## Installing and Removing Software

### • Primary RPM options:

- Install: `rpm -i, --install` } INSTALL
- Upgrade: `rpm -U, --upgrade` }
- Freshen: `rpm -F, --freshen` } upgrade ONLY IF IT IS INSTALLED.
- Erase: `rpm -e, --erase`
- Output Options: `-v, -h` } you verify } Eh hash } ALL (Progress)
- URL support: `ftp://` (with globbing), `http://`
- Many other install-options are available to address special cases.

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2884 or +1 (919) 754 3700.

Installing: `rpm -i`

*must be use to for kernel*

The primary function of RPM is to install, upgrade, and remove software from a system. A package is installed using a command such as `rpm -i zip-2.3-8.i386.rpm`.

When installing an rpm package, the rpm command will consult the local database to ensure that (1) any prerequisites (in the form of files, libraries, rpms, or generally defined "provisions") are installed on the system, and (2) installing the rpm will not clobber any preexisting files. The checks can be omitted by enabling the `--nodeps` or `--replacefiles` command-line switches, respectively, or both using the `--force` switch. rpm will provide "pretty" output if called with the `-v` (print package name) and `-h` (print hash marks) options.

Upgrading: `rpm -U` and `rpm -F`

rpm can be used to upgrade already installed software with the `-U` (`--upgrade`) command-line switch. When upgrading, the original package (with the exception of configuration files) on the system will be removed, and the new package installed. Configuration files from the original installation are saved with a `.rpmsave` extension.

*Freshening* is almost identical to upgrading, except when the package specified on the command line is not already installed on the system. When upgrading with `-U`, the package will be installed whether or not it is already installed; when freshening, the package will be ignored if not already installed. To apply all errata released by Red Hat for all packages installed on your system, ignoring errata for uninstalled packages, execute the following:

```
rpm -Fhv ftp://updates.redhat.com/current/en/os/i386/*.*rpm
```

Uninstalling: `rpm -e`

Software is removed from your system using the `-e` (`--erase`) command-line switch. The package argument must be the installed package's name, not the package file name. For example, these commands first install the `zip` package file, and then remove it:

```
rpm -ihv zip-2.3-8.i386.rpm
rpm -e zip
```

## Updating a Kernel RPM

- Make sure to install kernel updates
- Do not use `rpm -U` or `rpm -F!`
  - `rpm -ivh kernel-version.arch.rpm`
  - Boot new kernel to test
  - Revert to old kernel if a problem arises
  - `rpm -e kernel-oldversion` if no problems

Rev RH133RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (319) 754 3700.

Installing a new kernel is one of the few things you will do on your system that requires a reboot, unlike some other operating systems. It also requires a little more thought and caution, as it is quite simple to render a system temporarily inoperable if one is careless when updating the kernel. Unlike just about any other upgrade you might do, you should NOT upgrade the kernel using `rpm -U` or `-F`.

Recall how `rpm -U` and `-F` functions: it determines if a version already exists on the system, and if so, whether the version to be installed is newer. If it is, it removes the old version's files, excluding those tagged as configuration files when the RPM was built. It then installs the new files, runs any installation scripts that are included in the RPM, then updates the installed database.

Because upgrading removes the previous kernel version, if for some reason your newly-installed kernel proved unstable, you could be left with an unbootable system, and would have to resort to alternate boot media such as a boot floppy or the CDROM. By installing instead of upgrading, the old version of the kernel is still available and can be chosen from the boot loader.

In addition, kernel modules are version specific and an upgrade will remove all modules that your present kernel is using, leaving the system unable to dynamically load device drivers or other modules.

Because all of the kernel RPM's files are version specific, i.e., they either include version information in their names, or else are stored in version-specific paths, it is possible to install multiple versions of the kernel package. If you use `rpm -ivh`, instead of `-U`, then the new kernel will be added to your system, but your old kernel will still remain on it as well.

After installing the updated kernel you may want to edit the "default" line in `/boot/grub/grub.conf` and set it to 0 which will cause the system to load your new kernel by default at boot time.

## rpm Queries

- Syntax:  
`rpm -q what_packages what_information`
- Installed Package Options:
  - `rpm -qa` lists installed packages
  - `rpm -qf filename` shows owning package
  - `rpm -qi package_name` general information
  - `rpm -ql package_name` lists files in package
- Uninstalled Package Options:
  - `rpm -qip package_file.i386.rpm`
  - `rpm -qlp package_file.i686.rpm`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



8

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 265 2994 or +1 (519) 754 3700.

RPM provides robust querying, which is invoked with `rpm -q` or `rpmquery`. Query options fall into one of two categories: those that specify which packages to query, and those that specify what information to retrieve. The first must be specified; the second defaults to the package name.

To get a list of all installed packages, query all packages for the default information (package name):

```
rpm -qa
```

To obtain the name, including version and release, of a specific package, query that package:

```
rpm -q zsh
```

Querying a package for something besides its name is generally more useful:

```
rpm -qi zsh          lists package's information
rpm -ql zsh         lists files contained in package
```

Here's a list of common package specification parameters and what they return:

<code>-qa</code>	all installed packages
<code>-q <i>packagename</i></code>	the named package and version
<code>-qf <i>filename</i></code>	the package that owns the file
<code>-qp <i>package_file_name</i></code>	the (possibly uninstalled) package file
<code>--whatrequires <i>capability</i></code>	all packages that require <i>capability</i>
<code>--whatprovides <i>capability</i></code>	all packages that provide <i>capability</i>



The following is a list of common query information parameters and what they return:

-i	general package information
-l	package files
--requires	package prerequisites
--provides	capabilities provided by package
--scripts	scripts run upon installation and removal
--changelog	package revision history
--queryformat <i>format</i>	custom-formatted information (use rpm --querytags to list available tags for use in <i>format</i> string)

Solve dependences.

`rpm -q -redhatprovides`

`--quiet`

`rpm -q -p -changelog rpm`

[diff]

Kill

## rpm Verification

- Installed RPM File Verification:  

```
rpm -V package_name  
rpm -Vp package_file.i386.rpm  
rpm -Va
```
- Signature verification BEFORE package install:  

```
rpm --import gpg_key  
rpm --checksig package_file.i386.rpm
```

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat  9

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email training@redhat.com or phone toll free (USA) +1 888 265 2694 or +1 (919) 754 3700.

### RPM file verification

Verifying an installed package compares the file sizes, permissions, type, owner, group, MD5 checksum, and modify time against the RPM database. Any inconsistencies will be reported. An installed package can also be verified against a package file as well:

```
rpm -V zip - verifies the installed zip rpm against the RPM database  
rpm -Va - verifies all installed RPMS against the RPM database  
rpm -Vp zip-2.3-8.i386.rpm - verifies the installed zip package against the zip package file
```

### RPM signature verification

Red Hat signs all package files with a GPG private signature. The complementary public signature is shipped with every Red Hat distribution. To verify the integrity of any package file, you must first import the Red Hat public key. The rpm utility will automatically verify the signature of any package you install at install time. You can also check the integrity of package files using the `--checksig` option.

```
rpm --import /mnt/cdrom/RPM-GPG-KEY  
rpm -qa gpg-pubkey  
rpm --checksig zip-2.3-14.i386.rpm
```

## Other RPM Utilities and Features

- **rpm2cpio**: file extraction
- **rpmdb-redhat**: distribution database
  - rpm --redhatprovides filename
  - rpm --redhatprovides capability
- **system-config-packages**

Rev R4133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 298 2994 or +1 (919) 754 3700.

### **rpm2cpio**

The `rpm2cpio` command allows the files contained within a package file to be converted into a `cpio` stream. For example, to extract the executables from the `zip` package file into the local directory, try:

```
rpm2cpio /mnt/test/7.0-pub/i386/RedHat/RPMS/zip-2.3-8.i386.rpm | cpio --extract -  
-make-directories *bin*
```

### **rpmdb-redhat**

The `rpmdb-redhat` package allows the `rpm` command to access a database containing information on all packages in a RHEL release. Suppose that the `xjewel` package has the `libX11.so.6` library as a prerequisite. The `--redhatprovides` switch, along with `rpmdb-redhat-version.i386.rpm`, helps determine the prerequisite package:

```
rpm -ivh rpmdb-redhat-version.i386.rpm  
rpm --redhatprovides libX11.so.6
```

### **system-config-packages**

In addition to installing individual package files, you can install package file component groups using the `system-config-packages` utility. The `system-config-packages` utility uses the `/RedHat/base/comps.xml` file which is located on the first disc of the install media.

# Automatic Dependency Resolution

- Automatic installation of dependent packages
- Invoked with `--aid` option.
- Use in conjunction with `rpmdb-redhat`
- Macro can indicate where package files found

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2864 or +1 (619) 754 3700.

RPM supports a system for automatically determining packages required to resolve dependencies, and to load them automatically. The scope of this automatic dependency resolution is limited to situations where packages are dependent upon packages provided in the standard Red Hat release.

The following actions will allow an operator to take advantage of automatic dependency resolution:

1. `rpmdb-redhat`. Install the `rpmdb-redhat` package so that file dependencies can be determined.
2. `load packages`. Make the RPMs from the distribution available. The contents of the directory `RedHat/RPMS` on each of the install CDs should be available in a single directory on the local system.
3. `macro`. Create two RPM macros. One called `_solve_pkgsdir` and one called `_solve_name_fmt`. The first macro should name the directory holding the packages from step 2 above. One way of doing this is to create a file `/root/rpmmacros` containing the following lines

```
%_solve_pkgsdir /path/to/my/packages
```

```
%_solve_name_fmt %{?_solve_pkgsdir}%%{NAME}-%%{VERSION}-%%{RELEASE} %%\ {ARCH} rpm
```

4. `--aid`. When installing a package, use the `--aid` option to invoke the dependency resolution, eg.  
`# rpm -ivh --aid xsane-0.89-4.i386.rpm`

```
Preparing..... [100%]  
1:libusb..... [33%]  
2:sane-backends..... [67%]  
3:xsane..... [100%]
```

# Red Hat Network (RHN)

- RHN Components
  - RHN account
  - System identity
  - `/usr/sbin/up2date`
  - `rhnsd` daemon and queued actions
- Advantages
  - Errata concurrency
  - Collective and remote administration
  - Bare metal provisioning

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2294 or +1 (818) 754 3700.

The Red Hat Network allows administrators to efficiently manage software installation and upgrades using a combination of your RHN account and the **up2date** utility

In March of 2001 the "Lion" worm, a self-spreading program intended to compromise security on Linux systems, made headlines. According to a report found at [www.sans.org](http://www.sans.org), the worm makes use of a "BIND vulnerability .. that was reported back on January 29th, 2001". Red Hat had posted an updated version of the bind RPM that same day in January. If all Linux administrators had promptly updated their systems, the Lion worm would not have made headlines. Red Hat Network attempts to simplify the task of keeping software updated, reducing a system's vulnerability to exploits.

# RHN in the Enterprise

- Management Entitlements
  - System grouping
  - Multiple administrators
- Proxy Server
  - Updates cached locally conserving bandwidth
  - Private channels
- Satellite Server
  - Client profiles stored locally
  - Custom channel management
  - Provisioning Module

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 205 2994 or +1 (319) 754 3700.

Red Hat Network has several offerings for managing a large number of workstations, including Management Entitlements, a Proxy Server, and a Satellite Server. Each provides an increasing level of collective management and local control

## Management Entitlements

A Management Entitlement service account allows for the grouping of client systems, including collective software management and errata notifications. Multiple administrators may be defined and assigned to the various system groups.

## RHN Proxy Server

Software updates and errata may be locally cached using a RHN Proxy server. Client profiles are still maintained on RHN servers. The base channel (i.e., "Red Hat Enterprise Linux ES i386 2.1") is managed by RHN, but additional private sub-channels that allow the local distribution of custom software can be defined and locally administered. All interactions with RHN are mediated by the RHN Proxy Server, so only the proxy server needs Internet access

## RHN Satellite Server

The RHN Satellite Server allows all aspects to be managed locally, including client profiles and custom channel management. Systems can be provisioned from bare metal with the Provisioning Entitlement. System management is performed using a local web server, and a local database maintains client accounts and profiles. The Satellite Server allows complete channel definition, control, and management. No Internet access is required.

# RHN Registration

- `/usr/sbin/up2date`
  - username, password, system name
- Remote Information
  - Hardware Profile
  - Software Profile (RPM list)
  - Subscribed Channel
- Local Digital Certificate
  - `/etc/sysconfig/rhn/systemid`



Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 293 2994 or +1 (610) 754 3700.

The first time the `up2date` utility is run on a system, you will be prompted for information required to register the system with an existing RHN account. If you do not have a RHN account, you can create one at this time as well. The `up2date` utility can be run graphically from within an X window session or in “text” mode from a console. All communication between the `up2date` utility and the RHN servers is sent securely via the `https` protocol.

You are asked to provide a username/password pair for an existing or new account. Optionally, you may also provide additional contact information. Providing at least an email address is suggested so you can be notified of relevant system updates.

You will assign your system an identifying system name, and may optionally include a hardware profile (including CPU, Memory, disk partitioning, and peripheral devices) and software profile (a list of installed RPMs). While the profiles are optional, they allow RHN to customize update information to your system.

The registration utility creates a local certificate in `/etc/sysconfig/rhn/systemid` that serves to identify the machine to Red Hat Network. If for some reason, you wish to generate a new certificate for the system, you can do so by simply removing this file and re-running the `up2date` utility.

## The `up2date` utility

- Interactive or batch invocations
- Functions
  - Freshen with published errata/updates
  - Install new packages
  - Resolve package dependencies
- `/usr/sbin/up2date-config`
  - install or download only
  - cache dir: `/var/spool/up2date`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 268 2694 or +1 (919) 754 3700.

`up2date` allows you to initiate interactions from your local machine to RHN. From within an X session, or from the console, the utility can be used interactively by just running `up2date`. `up2date` will query RHN for any relevant updates published, download the RPM package files to the `/var/spool/up2date` directory, and optionally upgrade the packages as well.

`up2date` can also be run in "batch mode", by specifying one of the following switches:

```
-u, --update      update according to default configuration
-l, --listlist    relevant updates only
-d, --download    download relevant updates only
-i, --install     download and install relevant updates
-p, --packages    resync RHN profile to currently installed RPMs
```

The following script, when placed into the `/etc/cron.daily` directory, could be used to download daily any relevant updates to the `/var/spool/up2date` directory.

```
#!/bin/bash
/usr/sbin/up2date --download
```

When provided packages as command-line arguments, `up2date` will attempt to download and install the packages, satisfying any dependencies as well. For example, if no `samba`-related packages are installed, the following command would download and install the `samba` package, as well as the prerequisite `samba-common` package

```
up2date samba
```

Note also the following command-line switches, which are helpful in solving package dependencies:

```
--whatprovides=<deps>    list packages that resolve deps
--solvedeps=<deps>       install packages that resolve deps
```



## Remote Administration

- Web based administration
  - `https://rhn.redhat.com`
  - Queuing of actions
- Local polling: `rhnsd`
  - Every 4 hours by default
    - Tuned in `/etc/sysconfig/rhnsd`
  - `/usr/sbin/rhn_check` does the hard work

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 298 2994 or +1 (619) 754 3700.

RHN can be used to perform remote administration of collections of machines. First, actions for the machine (such as specific package installation or upgrades) are queued for the machine using the RHN account.

Client machines use the `rhnsd` daemon to poll RHN periodically for queued actions. By default, `rhnsd` polls every 4 hours, though this can be adjusted in `/etc/sysconfig/rhn/rhnsd`. The `rhnsd` daemon uses the `/usr/sbin/rhn_check` command to actually perform the poll and administer any queued actions. Notably, the `rhnsd` daemon does not open any server networking ports.

## Network Installation Server

- Necessary for network-based installs
- Often faster than CDROM-based installation methods
- Provides an easy distribution platform for the enterprise
- Shares the RedHat directory via NFS, FTP and/or HTTP

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2664 or +1 (919) 754 3700.

A network installation server can be easily created from a host running an NFS, FTP, and/or HTTP server. To create an installation server, select one or more installation protocols:

- If using anonymous FTP, copy each install CD's RedHat directory to the publicly-accessible FTP directory:
  - `cp -a /mnt/cdrom/RedHat /var/ftp/pub`
- If using HTTP, either use the above technique to copy the RedHat directories to a location in your web file tree. If you choose to provide installation services via both FTP and HTTP, create a symbolic link from your web file tree to your existing FTP directory:
  - `ln -s /ftp/pub /var/www/html/pub`
  - `chcon -h -R -t httpd_sys_content_t /var/ftp/pub`
- If using NFS, create an entry in `/etc/exports` to share the `pub` directory, for example:
  - `/var/ftp/pub hosts and or networks(ro)`
- Restart the servers you've chosen to configure, if necessary.

Once a server is set up, an installation can be started on a client machine using a `boot.iso`, a `diskboot.img` USB key, or by PXE, which fetches and executes the second stage installer and utilities from the server. Once the installation configuration is determined (either by interactive prompts or automatically via Kickstart), the server provides the installation RPMs via the selected protocol.

## Using Kickstart to automate installation

- Kickstart is a component of the installer that automates installation
- Kickstart supports all installation methods
- The installer reads information from an ASCII file rather than prompting for it
- Kickstart files can be made available via floppy, cdrom, hard disk, initrd, nfs, ftp and http. They can also be dynamically generated using cgi scripts and specified using dhcp/pxe.

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 298 2994 or +1 (019) 754 3700.

Kickstart allows the installer to read information from a designated file rather than prompting the person doing the installation for it. It especially facilitates the setup of a number of machines that have similar hardware that the installer can autoprobe successfully. If a required item is omitted from the Kickstart file, the installation pauses and the user is prompted for that information. Kickstart "type" refers to the location of the Kickstart file. The first chapter of the Customization Guide describes a number of possible locations, but the most common types are "floppy-based" and "network-based." These approaches are described below.

For a floppy-based Kickstart :

- Create a Kickstart configuration file named `ks.cfg`
- Create a `boot.img` or `bootnet.img` diskette and copy `ks.cfg` to its top-level directory
- Boot the target machine with the boot diskette, and type `linux ks=floppy` at the boot prompt.

For a network-based Kickstart:

- A DHCP server is necessary for network-based Kickstarts, even if the system to be built will use a static IP address. It may also instruct the Kickstarted client about the location of the Kickstart file and the server and NFS share on which it may be found.
- If the location given is a file, then the client will try to mount the file's parent directory to retrieve the file. If the location given is a directory, then the client will try to mount that directory and will look for a file whose name is `<DHCP-IP-Address>-Kickstart`.
- If the DHCP server does not provide a `next-server` name, then the client will assume the DHCP is the server on which the Kickstart file resides. If the DHCP server does not provide a `filename`, then the client will assume the directory is `/Kickstart` and will look for a Kickstart file with a name of the `<DHCP-IP-Address>-Kickstart` form.
- Boot the target machine with a `bootnet.img` diskette (or using a PXE image)

RHEL supports additional methods for accessing the Kickstart file. Further information about Kickstart may be found in the [Official Red Hat Enterprise Linux Customization Guide, Chapter 1](#)

## Kickstart: Commands Section

- Constructs arguments that are passed to configuration utilities ("commands")
- The absence of required specifications (e.g., keyboard) will raise the appropriate utility
- Commands section must come first

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyright. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2804 or +1 (910) 754 3700.

The commands section specifies configuration actions taken by the installer. This section must come first in the Kickstart file, but the order of lines within the section is not especially significant.

### Utility-based Directives

Many of the command section directives are simply switches and arguments passed to utilities used by the installer. These directives include the following:

<u>Directive</u>	<u>Utility</u>
keyboard	system-config-keyboard
timezone	system-config-time
xconfig	system-config-display
auth	system-config-authentication
mouse	system-config-mouse
rootpw	system-config-rootpassword
firewall	system-config-securitylevel

### Partitioning Directives

Partitioning directives provide information used by Disk Druid to partition a Kickstart-installed system and specify mount points. As with interactive installation, Disk Druid may act as both partitioning tool and utility for setting mount points (i.e., creating `/etc/fstab`), or it may simply set mount points if partitions already exist. Partitioning directives can also create software RAID arrays, logical volumes, and volume groups.

```
part <filesystem> --size <size> [--grow --maxsize <size>]
part swap --size
part <raid component>
raid
lvmlogvol
volgroup
```

## Other


The remaining Kickstart directives handle other configuration issues and provide various mechanisms for fine-tuning the installation process. These include `bootloader`, `lilocheck`, `skipx`, `lang`, `langsupport`, `zerombr`, `clearpart`, `reboot`, `install|update`, and `nfs|url|cdrom|harddrive`.

## Kickstart: %packages

- %packages specifies components groups and RPMs to install
- Component groups in the comps.xml file are specified with @ component-group
- Third-party RPMs cannot be specified without modifying hdlist
- Package names only (not version)

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 265 2994 or +1 (315) 754-3700.

The %packages section must come after the command section, although it need not be second. The packages section is a list with one item per line. The item may be an individual RPM package name, a component group specified in the comps.xml file, or @ **Everything** to install all packages.

It is possible to specify third-party RPMs in the %packages section, but it requires modifications to the database of information on the distributions RPMs, **hdlist**. See the discussion on the Network Installation Server earlier in this unit for more information.

```
# The choices below would install three component groups plus the mutt and vlock RPMs.
@ Workstation Common
@ GNOME
@ Kernel Development
mutt
vlock
```

## Kickstart: %pre, %post

- %pre gives you the first word
  - executes as a bash shell script
  - executes after Kickstart file is parsed
- %post gives you the final word
  - Can specify interpreter (bash is default)
  - chroot'ed by default, but may be run without chroot

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (919) 754 3700.

The %pre and %post sections make just about anything possible. The %pre section executes as a bash script after the Kickstart file is read, but before partitioning, formatting, and copying of packages commences. The %pre script runs in a somewhat limited environment, as the only executables available are the ones provided by the installer. Unlike %pre, all of the packages specified in the %packages section are available in the %post section, which means many more utilities and capabilities are available.

The default behavior of %post is to execute the contents of the section as a bash shell script in an environment that chroots to /mnt/sysimage -- the newly installed system. In other words, all paths and commands are as they will be on the installed system. The default use of bash as the interpreter may be overridden using the --interpreter switch on the %post line, e.g., %post --interpreter /usr/bin/perl. A non-chrooted environment is also possible through the --nochroot switch.

Examples:

```
%pre
# Create partitions by copying an MBR image
mknod /tmp/hda
dd if=/mnt/source/pub/mbr.img of=/tmp/hda
%post
# Download a customized XF86Config file via ftp -- the "echo" allows a
# DHCP client system to use hostnames in the %post section
echo "nameserver 192.168.0.254" >> /etc/resolv.conf
lynx -source ftp://server1/pub/XF86Config > /etc/X11/XF86Config
# Suppress the rewriting of /etc/resolv.conf on a DHCP client
cat >> /etc/sysconfig/network-scripts/ifcfg-eth0 <<EOF
PEERDNS=no
EOF
```

## End of Unit 6

- Questions and answers
- Summary
  - What are the primary functions of RPM?
  - What rpm options should be used to install a kernel RPM?
  - Where can the Kickstart configuration file be stored?

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

22

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 268 2994 or +1 (919) 754 3700.

### Important files covered in this Unit:

```
/etc/sysconfig/rhn/systemid  
ks.cfg
```

### Important commands covered in this Unit:

```
rpm  
rpm2cpio  
rhn_register  
up2date  
rhnsd
```



# Lab 6

## RPM and Kickstart

**Goal:** Install Red Hat Enterprise Linux using Kickstart

**Sequence 1:** Kickstart Installation

**Tasks:**

Before you begin, read the troubleshooting suggestions at the end of the sequence.


1. Copy the `/root/anaconda-ks.cfg` to `/root/ks.cfg`

2. Insert the following line at the top of the `/root/ks.cfg` file:

```
nfs --server=server1.example.com --dir=/var/ftp/pub
```

3. Rather than using the commented partition information, enter the following into your `ks.cfg` file:

```
clearpart --all
1 part / --fstype ext3 --size=400
2 part /boot --fstype ext3 --size=100
3 part /tmp --fstype ext3 --size=128
4 part /usr --fstype ext3 --size=2800
5 part /var --fstype ext3 --size=400
6 part /home --fstype ext3 --size=128
part swap --size=512
```



4. Add the following below the `%post` directive:

```
rpm -i ftp://server1.example.com/pub/RedHat/RPMS/vim-enhanced*.rpm
rpm -i ftp://server1.example.com/pub/gls/RPMS/rhce-ts*.rpm
useradd <insert your username here>
echo <insert password here> |passwd --stdin <username from previous line>
```

5. Format and then mount a floppy. Then copy your `ks.cfg` to the floppy:

```
fdformat /dev/fd0H1440          # Low level floppy format
mkfs -t ext2 /dev/fd0          # puts an ext2 filesystem on the floppy
mount /media/floppy
cp /root/ks.cfg /media/floppy
umount /media/floppy          # VERY IMPORTANT! Floppies must be unmounted
```

6. Reboot your system using `cd #1` or from media provided by the instructor. The kickstart floppy is not bootable so if your system's bios is set to boot from floppy first you will need to remove the floppy and reinsert it after the system boots from the cdrom and you see the Red Hat Enterprise Linux installer `boot :` prompt.

- 7 After first making sure that the floppy is in the drive, when the system comes to the `boot :` prompt type:

```
linux ks=floppy
```

If anything is missing on the kickstart floppy the installer will raise a dialogue allowing you to add in the required information.

You will use this installation for the remainder of this course

### Troubleshooting suggestions:

If you have a typo in your kickstart file, boot the system to runlevel 1 to fix it. This is much faster than booting to runlevels 3 or 5.

If an installer screen appears, such as the ones for configuring the language or keyboard, then you may be missing a line from your `ks.cfg`

If Disk Druid appears, then you probably misspecified your partitions. Make sure there is sufficient space for your partitioning scheme and that you included a swap partition

The Python interpreter will spew ugliness everywhere if there is a fatal error. Examine this mess carefully -- you can use the `<Shift><Page Up>` and `<Shift><Page Down>` keystroke combinations to scroll the screen up and down. Careful examination of the Python traceback will usually reveal where the error is, even if you are not fluent in Python.

If there is a problem somewhere other than the `%post` section, it will probably appear before your system is overwritten. Consequently, you can reboot your system to examine and fix your `ks.cfg` file. Booting into single user mode should speed progress.

## Sequence 2: Installing Errata and Dealing with RPM

## Tasks:

1. Use rpm queries to answer the following questions. In the blank spaces, write in the command used to find the answers:

- a. What files are in the *initscripts* package?

`rpm -qf initscripts`

- b. On what host was the *bash* RPM built, and what is its installed size?

`rpm -qf bash`  
 5113068  
 perl4.build@redhat.com

- c. Has the *pam* package changed since it was installed?

`rpm -Va pam`  
 Yes

- d. Which installed packages have "gnome" in their names?

`rpm -qa *gnome*`

- e. Which RPM provides `/etc/inittab`?

`rpm -qf inittab`  
 initscripts

- f. Which RPM provides `/etc/fstab`? Why?

- g. What was the last changelog entry for your kernel?

- h. What are the differences between the following commands?

`rpm -ivh <package file>` INSTALL

`rpm -Uvh <package file>` UPGRADE

`rpm -Fvh <package file>` FRESHEN

2. Practice checking the signature and integrity of an RPM package file of your choosing from your CDROM or from server1

Use `rpm --import` to add Red Hat's GPG key to a system-wide keyring:

```
rpm --import /usr/share/rhn/RPM-GPG-KEY
```

The following `rpm` invocation will test whether the package was signed by the private key associated with the public key you added to your keyring, and whether the MD5 checksum of the package is unchanged from the time the package was created. The `-K` option is the short option for `--checksig`.

```
rpm -K <RPM package file>
```

3. Apply errata. On server1 is a directory containing some errata that is applicable to the release of Red Hat Enterprise Linux that you are running. Freshen your system by applying these errata packages. Bear in mind the following ( read all before starting ):
  - a. You can copy the files to a local file system using ftp or http, or access them from server1 using an NFS mount ( preferred ). If you copy them to the local system, make sure you have sufficient space in the target file system.
  - b. Do not apply errata to the kernel using the normal package upgrade method. **New kernel packages should be applied with an (I)nstall operation, not an (U)pgrade operation or (F)reshen.**
  - c. You probably want to use (-F) to freshen the existing RPMS on your system. Remember that (-U) will upgrade existing RPM(s) but also install RPM(s) that are not currently installed on your system.

**Extra:**

Create a corrupted RPM then verify it. Start by copying a RPM file to /tmp, then use the cat command to append some extraneous data to the end of the file:

```
cp /mnt/server1/errata/foobar*.rpm /tmp
cat /bin/date >> /tmp/foobar-2.9-14.i386.rpm
```

Then use RPM to see if the file is a genuine Red Hat RPM:

```
rpm -K /tmp/foobar-2.9-14.i386.rpm
```

This should fail horribly.

**Sequence 3: Automatic dependency resolution.**

Initial situation. Before commencing this lab ensure that none of the following packages are loaded. If you need to remove a package and run into a dependency problem, use the `--nodeps` option to circumvent it and force the removal of the package. Please note, normally this is a bad idea.

```
rpmdb-redhat
xsane
sane-backends
```

`--aid`

Packages to be loaded are available as usual by NFS from `server1.example.com:/var/ftp/pub` which may be mounted to the local system. Packages to be used are in this share under `RedHat/RPMS`.

1. Observe no resolution case. To help appreciate the benefits of automatic dependency resolution, first attempt a package load without it. Attempt to install the `xsane` package from `{mount-point}/RedHat/RPMS`

This should fail and give an indication like:

```
error: Failed dependencies:
  libsane.so.1 is needed by xsane-0.89-3
  libusb-0.1.so.4 is needed by xsane-0.89-3
```

Do not attempt to complete the installation by this method

2. Using `rpmdb-redhat`. Install the package `rpmdb-redhat` then re-attempt the installation of the `xsane` package. This should fail again, but give more useful additional information like:

Suggested resolutions:

```
libusb-0.1.6-3.i386.rpm
sane-backends-1.0.9-5.i386.rpm
```

3. `Aid` Make sure your current directory is `{mount-point}/RedHat/RPMS`. Re-attempt the installation of the `xsane` package using the `--aid` option, ie

```
rpm -ivh --aid xsane-0.*.rpm
```

You should see that the package `sane-backends` is loaded automatically to satisfy the dependency.

Note, in this case the `rpm` macro method mentioned in the notes was not required because the package and its dependency were together in the current directory.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100

100

100

100

100

# UNIT 7

## User Administration

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (910) 754 3700.

## UNIT 7: Objectives

- Upon completion of this unit you should be able to:
  - Create, modify, and delete user accounts
  - Create, modify, delete group accounts
  - Modify file ownership and permissions
  - Limit access to files with "special" permissions
  - Set group access to files and directories with `umask` and the UPG scheme
  - Configure a user's shell environment

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2864 or +1 (919) 754 3700.



## UNIT 7: Agenda

- User accounts
- Group accounts
- File ownership and permissions
- "Special" permissions SUID / SGID / Sticky
- Switching accounts with `su`
- `umask` and the UPG scheme
- Shell environment

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3


For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2094 or +1 (619) 754 3700.

ACTIVILUS

## User Policy Considerations

- Amount of system access outside of user's account
  - Determine "need to know"
- Expiration of passwords and accounts
- Disk usage and CPU limits

6 or 1295 and CPU

redhat  4

Rev RH133-RHEL4-1 Copyright © 2005 Red Hat, Inc.

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 296 2994 or +1 (519) 754 3700.

### User Policy

When setting up or administering user accounts, there are a number of things to be considered that have no hard and fast rules. In some cases, it may be fine to grant users a large amount of latitude and in other situations, unthinkable. For example, on a home computer or a test system granting users lots of latitude is usually fine. On a corporate server, the user policy will likely dictate strict practices in order to maintain the security and integrity of the server.

### Key Considerations

There are several issues to consider in defining user policy. These include:

- The amount of access to system files and resources
- Whether to force periodic password changes
- Whether to limit logins to certain times and places
- Whether to enforce CPU and memory limits
- Whether to enable disk quotas

Some of these policies may apply to the entire site while others may be applied to individual users. Often these policies are set by corporate policy and may vary from system to system depending on the sensitivity and value of data and resources on the system.


As a rule of thumb it is probably best to have strict policies, at least in the beginning. If a situation comes up that requires the relaxing of a policy, the change can be made. Having a rather open system may lead to problems with security that may be difficult to deal with after a problem occurs.

## User Account Database: `/etc/passwd`

- Contains account information used at login and by other programs
  - One account per line with seven colon-delimited fields
  - Should have permissions `rw-r--r--`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (919) 754 3702.

The `/etc/passwd` file is the standard Linux database for user accounts. Every user on the system will have a single one-line entry in the `/etc/passwd` file. The `/etc/passwd` file contains information about user and system accounts that is required at login and by other programs. The permissions of `/etc/passwd` should always be 644. This allows utilities that read the file to operate, and still prevents non-root users from adding or modifying entries. Each entry in the file follows a precise syntax. There are seven fields separated by colons:

```
uname uid uid gid password home dir shell
```

```
bcroft:x:502:504:Bryan Croft:/home/bcroft:/bin/bash
```

644

The first field is the user ID name of the account

The second field is the password field. This field will contain an 'x' if shadow passwords are implemented, or else will contain the user's encrypted password.

Next comes the UID field. The UID number should be unique for each account on the system. A common convention is to number them sequentially starting from 500. Accounts with a UID under 100 are generally system accounts such as `lp`. These accounts rarely, if ever, have valid passwords and are not meant to be used as a login account. Many services and daemons, however, need to have an account in order to behave properly. Note: A UID number of 0 indicates that the user is privileged, i.e., they have superuser access.

After the UID field comes the GID field, which specifies the initial group that the user is placed into when they logon.

The fifth field is the GECOS, or comment field. Items such as the user's real name go here. Some utilities, such as `finger`, use this field.

Next is the user's home directory and finally, the command executed by login, which is usually a shell (although it could be an application or a shell script.)

## Adding a New User Account

- Most common method is `useradd`:  
`useradd username`
- Running `useradd` is equivalent to:
  - `edit /etc/passwd, /etc/shadow, /etc/group`
  - create and populate home directory
  - set permissions and ownership
- Set account password using `passwd`
- Accounts may be added in a batch with `newusers`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll-free (USA) +1 866 286 2894 or +1 (519) 754 3700.

The command-line utility `useradd` provides a simple method for adding new users to the system:

```
# useradd joshua
```

The above command will add a new account to the machine called `joshua` as well as set up that user's home directory, and create a private group for the user, also called `joshua`. The next step would be to assign `joshua` a password which you can do by simply typing the following command:

```
# passwd joshua
```

When you need to add several users, you can use the `newusers` command. Create a file, formatted like `/etc/passwd`, that contains the usernames and plain text passwords of the users. The command: `newusers <yourfile>` will create those users. One apparent drawback of using this method is that the users home directories do not get populated with the files from `/etc/skel`.

template for password

## User Private Groups

- When user accounts are created, a private group is also created with the same name
  - Users are assigned to this private group
  - User's new files affiliated with this group
- Advantage: Prevents new files from belonging to a "public" group
- Disadvantage: May encourage making files "world-accessible"

*Handwritten notes: A circle around "world-accessible" and the initials "htc".*

Rev Rn133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2694 or +1 (918) 754 3700.

Each new user added to a Red Hat Enterprise Linux system also triggers the creation of a new User Private Group, with the same name as the new user. This scheme allows the user's `umask` to be set to 002, but still restricts write access to files and directories created by the user by other users. Since every user has a corresponding private group, the traditional Unix `umask` of 022 (disallowing write permission to everyone but the user) is not used.

The idea is that if the user has a `umask` setting of 002, new files and directories will be protected since they will belong to a group that has only one member, namely the user. Of course, the user may belong to other groups and switch to one with the `newgrp` command and/or change the file's group ownership with `chgrp`:

```
$ chgrp <groupname> <filename>
```

The drawback is that often users will not do any of these and just make the file world-accessible with `chmod` so that anyone *can* access it when only a few may *need* access.

When a system administrator decides to enable multiple users write access to files in a directory, those users are added to a common group. The group ownership on the directory and its files is changed to the common group, and those directories and files are given group-write permission.

When the group set bit is set on the common directory with `chmod g+s`, a file or subdirectory created there by a member of the common group will be owned by that user, but will be assigned group ownership of the common group.

# Group Administration

- Entries added to `/etc/group`  
`groupadd`  
`groupmod`  
`groupdel`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

8

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2884 or +1 (910) 754 3700.

New groups may be created by hand-editing the file `/etc/group` or by using `groupadd`. The basic syntax for `groupadd` is very simple:

```
# groupadd groupname
```

`groupdel` is used in a similar fashion to remove groups:

```
# groupdel groupname
```

`groupmod` can be used, among other things, to change the name of a group:

```
# groupmod -n newname oldname
```

For example, if several users are members of the `employee` group and a root user issues the following command, the group will be changed to `staff` and all the same members will remain:

```
# groupmod -n staff employee
```

`/etc/group` contains one group per line:

```
gurus:x:501:joshua,dax,bryan,chrish,heather,jon
```

The first field is the name of the group. The second field is the group password, or an "x" when using shadow passwords. The third field is the unique group ID. The fourth field is a comma-separated list of group members.

In order to avoid using a GID within the range typically assigned to users and their private groups, use `-r`:

```
# groupadd -r groupname
```

`gpasswd` can be used to define group members in `/etc/group`, group administrators, and to create or change group passwords in `/etc/gshadow`, if desired.

VIPW - user

VIPW

VI PASSWORD.

EDIT - SKELIFE

## Modifying / Deleting Accounts

- To change fields in a user's `/etc/passwd` entry you can:
  - Edit the file by hand
  - Use `usermod [options] username`
- To remove a user either:
  - Manually remove the user from `/etc/passwd`, `/etc/shadow`, `/etc/group`, `/var/spool/mail`
- Use `userdel [-r] username`

home dir

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 200 2984 or +1 (615) 764 3700.

In addition to hand-editing `/etc/passwd`, you may use `usermod` to change account information. The following options are available:

<code>-c &lt;comment&gt;</code>	Change the comment field. This is often the user's full name.
<code>-d &lt;home dir&gt;</code>	Change the home directory.
<code>-e &lt;expire date&gt;</code>	Set date on which the account will expire and be disabled.
<code>-g &lt;group&gt;</code>	Change the initial login group.
<code>-G &lt;group, [...]&gt;</code>	A comma separated list of supplementary groups for the user.
<code>-l &lt;login name&gt;</code>	Change the login name.
<code>-s &lt;shell&gt;</code>	Change the login shell.
<code>-u &lt;uid&gt;</code>	Change the login ID.
<code>-p &lt;password&gt;</code>	Change the string in the password field.
<code>-L</code>	Lock the password. This renders the account unusable.
<code>-U</code>	Unlock the password.

`userdel` deletes a user from the system. It may be prudent to lock the user's account first with `usermod -L`, and to delay deletion of the user's account until you are sure that none of the files in the user's home directory are still needed.

When deleting an account with `userdel`, you should consider using the `-r` option, which recursively deletes the user's home directory. Deleting accounts without deleting the associated home directories may cause issues with file ownership when future users are added to the system.

When adding user access to a group with the `-G` option, you must list all groups to which a user belongs.

# Password Aging Policies

- By default, passwords do not expire
- Forcing passwords to expire is part of a strong security policy
- Modify default expiration settings in `/etc/login.defs`
- To modify password aging for existing users, use the `chage` command  
`chage [options] username`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2934 or +1 (319) 754 3730.

By default passwords do not expire. This means that it is possible for a user to have the same password indefinitely. This situation is not very secure, because if a password has leaked, or been compromised, it will remain so forever. This can be adjusted in the `/etc/login.defs` file

The `chage` command is used to set up password aging. You may set the maximum amount of time that a password is considered valid before the system will force the user to change his password. The security policy of an organization will generally define the amount of time between password changes. It is also important to set the minimum amount of time that a password must be used before it can be changed. This prevents users from changing their password when required to by the system, and then changing it right back to the old value.

```
# chage [options] username
```

Common options used with the `chage` command:

- m minimum days between password changes
- M maximum days between password changes
- I number of days inactive since password expired before locking account
- E date expire the account on this date (YYYY-MM-DD format)
- W number of days before a required change to start warnings



# Login Shell Scripts

- `/etc/profile`
  - `/etc/profile.d/*.sh`
- `~/.bash_profile`
  - `~/.bashrc`
    - `/etc/bashrc`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2864 or +1 (915) 754 3700.

## Shell configuration scripts

`/etc/profile` is executed every time a user logs into the system. It will set environment variables for the user such as `HISTSIZE` and `MAIL`. This is the first script executed at login. The user's `~/.bash_profile` script runs next, which typically calls `~/.bashrc` and `/etc/bashrc`. The `/etc/profile` contains system-wide environment settings -- whatever is added or removed here affects all users.

`/etc/profile.d` contains initialization scripts specific to software packages installed by RPM. These scripts are called by `/etc/profile` on login, or by `/etc/bashrc` if called by a non-login interactive shell.

`~/.bash_profile` is executed once at login time. It is usually used to set environment variables and to start programs at login, as opposed to every time you open a terminal window.

## Non Login Shell Scripts

- `~/ .bashrc`
  - `/etc/bashrc`
    - `/etc/profile.d/*.sh`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



12

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2094 or +1 (619) 754 3700.

`~/ .bashrc` allows users to customize their own aliases and functions without the intervention of the administrator. It runs whenever a user starts up a non-login interactive shell, and the default user `~/ .bash_profile` also calls it whenever the user logs in.

For example, RHEL aliases `ll` to `ls -l --color=tty`, but a user might want `ll` aliased to `ls -laF --color`. To do this, one would add this alias to the end of `~/ .bashrc`, then `source .bashrc`. Every time a user starts up an interactive shell, the personalized alias would be set, overriding the system alias.

The `/etc/bashrc` script is used for system-wide functions and aliases. It allows a system administrator to set aliases for every user, like `c` for `clear` or `h` for `history`. This script is usually called by a user's `~/ .bashrc` file.

## Switching Accounts

- Syntax

```
su [-] [user]
```

```
su [-] [user] -c command
```

- Allows the user to temporarily become another user

- Default user is root

- The "-" option makes the new shell a login shell

Rev 09133 RHSL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

The `su` command is used to switch to another account from the command line. This command is most often used by system administrators to temporarily become the root user without logging out of their non-privileged account. It is very important to only use the root account when absolutely necessary, because of the awesome power of the root account. Day to day use of the system should never be conducted with the root account.

The password of the account being switched to must be supplied unless the superuser issued the `su` command.

Without the `-` option, the original user's environment is maintained. Using the `-` option causes the new shell to be a login shell which, among other things, unsets the original user's environment variables in the new shell. This is usually preferred so that actions performed as the "new" user will not have any effect on the "old" user, and so the "new" user's environment is available, rather than the old.

There are a number of options to the `su` command. See the online documentation for more information.

Most systems log the use (or attempted use) of `su` to change to the root account. System administrators will often log on to the system as an ordinary user and then use `su` to gain access to the root account rather than log in directly as root so that the switch is logged.

Most systems administration tasks are best performed using `sudo`. `sudo` is safer than a `su` to root.

# sudo

- Users listed in `/etc/sudoers` execute commands with:
  - an effective user id of 0
  - group id of root's group
- An administrator will be contacted if a user not listed in `/etc/sudoers` attempts to use `sudo`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2664 or +1 (919) 754 3700.

`sudo` access is controlled by `/etc/sudoers`. This file should be edited with `visudo`, an editor and syntax checker. `/etc/sudoers` contains mappings of variables to reference groups of users, hosts, or commands. To give a specific group of users limited root privileges, edit the file with `visudo` as follows:

- In the user alias specification section, list users and groups allowed to use the `sudo` command:  
`User_Alias LIMITEDTRUST=student1,student`
- In the command alias specification section, list the commands specifically allowed or denied execution as root:  
`Cmnd_Alias MINIMUM=/etc/rc.d/init.d/httpd`  
`Cmnd_Alias SHELLS=/bin/sh,/bin/bash`  
(NOTE: the above list should match the entries in `/etc/shells` for more secure results!)
- In the user privilege specification section, list the users and groups allowed to use `sudo` and the commands that they may use:

```
LIMITEDTRUST ALL=MINIMUM
```

Users `student1` and `student2` can use `sudo` only with the commands listed with `MINIMUM`.

```
student3 ALL=ALL,!SHELLS
```

Means that user `student3` may use `sudo` with every command except `/bin/sh` and `/bin/bash`.  
`%development station1=ALL,!SHELLS`

This declaration means that every member of the `development` group can use `sudo` with every command when they are logged into `station1`. The only commands they may not issue are `/bin/sh` and `/bin/bash`.

## Network Users

- Information about users may be centrally stored and managed on a remote server
- Two types of information must always be provided for each user account
  - Account information: UID number, default shell, home directory, group memberships, and so on
  - Authentication: a way to tell that the password provided on login for an account is correct

Rev RH133/RHEL4-1

Copyright © 2004 Red Hat, Inc.

redhat  15

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (018) 754 3700.

We can store information about user and group accounts in local files like `/etc/passwd` on each workstation or server. However, it may be easier to keep account information synchronized among many computers by storing it centrally in a remote network server

Two basic types of information must be provided for each user account. First, name service information must be provided to the standard C library, *glibc*. This information maps the account's username to a UID number, primary group (by GID number), GECOS field information (such as the user's real name), home directory, and default shell. This is controlled by a system called **Name Service Switch**, or **NSS**. Second, a mechanism must be provided which can determine if a password provided to authenticate login or other access to a particular account is the correct password for that account. This is configured for programs through the **Pluggable Authentication Modules** system, or **PAM**

As shipped, NSS is configured to get information about users and groups only from local files. PAM is configured by default to authenticate users by encrypting the password provided to the login program and comparing it to the encrypted password information provided by NSS (and normally stored in `/etc/shadow`) for that username.

# Authentication Configuration

- **system-config-authentication**
  - GUI tool to configure authentication
  - For text-based tool, use `--nox` option
- **Supported account information services:**
  - (local files), NIS, LDAP, Hesiod, Winbind
- **Supported authentication mechanisms:**
  - (NSS), Kerberos, LDAP, SMB, Winbind

Auth-config

Rev RH133-RHEL4-1

Copyright © 2004 Red Hat, Inc.

redhat 16

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2994 or +1 (919) 754 3700.

The default mechanisms used to manage user information can be changed with `system-config-authentication`. By default, this provides a graphical tool to configure network authentication. The interface provides two tabs: "User Information" (which changes NSS settings) and "Authentication" (which changes PAM settings). Check boxes are used to select which user information sources or authentication mechanisms are desired, and then each service must be configured through button selections. In addition, use of shadow passwords or the MD5 password encryption algorithm can be toggled on or off. The `--nox` option will open this tool as a text-based menu interface instead, usable without an X display. In that case, the first screen will allow selection of the user information sources or authentication mechanisms, and subsequent screens will be used to configure those interfaces.

For "User Information", five data sources are supported:

If nothing is selected, only **local files** will be used as a source of NSS information. **NIS** gets information from database maps stored on a NIS server. **LDAP** allows account information to be stored as entries in a LDAP directory server. **Hesiod** stores user information as special resource records in a DNS name server, and its use is relatively uncommon. **Winbind** uses `winbindd` to automatically map accounts stored in a Microsoft Windows domain controller to Linux users by storing SID to UID/GID mappings in a database and automatically generating any other NSS information that is required.

For "Authentication", five data sources are also supported:

If nothing is selected, it is assumed that NSS will provide an encrypted password with the other NSS user information that can be compared to the entered password normally. **Kerberos** authenticates users by requesting a "ticket" for the user from the Kerberos server, and if the user's password decrypts the ticket the authentication passes. **LDAP** authentication maps the username provided to a LDAP directory entry and tries to bind to the directory using that entry and the provided password; if this succeeds, the authentication passes. **SMB** and **Winbind** use different approaches to authenticate using a Microsoft Windows domain controller.

The older version of this tool in Red Hat Enterprise Linux 3 and earlier was `redhat-config-authentication`, and it had a slightly different look and feel. The `authconfig` command may start up the `--nox` version of the GUI tools, and also supports a `--kickstart` option which can make these settings through command-line flags.

## Example: NIS Configuration

- Must install `ybind` and `portmap` RPMs
- Run `system-config-authentication`
  - Enable NIS to provide User Information
  - Specify NIS server and NIS domain name
  - Keep default authentication (through NSS)
- What does this actually do?
  - Four text-based configuration files are changed

Rev RH133-RHEL4-1

Copyright © 2004 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

One popular service used to centrally manage system and account information is NIS. NIS uses one or more NIS servers, each running `ypserv`, to share information with NIS client systems, running `ybind`. The master server may also run `rpc.yppasswdd`, which allows users on NIS clients to update the passwords stored in NIS. Both NIS clients and NIS servers also must run a local service called `portmap` which helps remote systems contact the local `ypserv` or `ybind` program. Clients and servers which communicate with each other are normally members of the same *NIS domain*, identified by an arbitrary name.

NIS servers typically are used to synchronize account information. They can share the contents of the `/etc/passwd`, `/etc/shadow`, and `/etc/group` files by converting them into *NIS maps*. Each NIS map consists of a set of key/value pairs. For instance, one typical NIS map used is `passwd.byname`, where the key is a username and the value is the matching line of user information for that account in `/etc/passwd` format.

The easiest way to set up a client to use an existing NIS server is to install the `portmap` and `ybind` packages, and run `system-config-authentication`. Under "User Information", enable NIS, and then a NIS domain and a NIS server for that domain must be specified. No changes are necessary under "Authentication" if NIS will be used for authentication, since it provides password hash information through NSS.

What does this change? The variable `NISDOMAIN` is set in `/etc/sysconfig/network` to the NIS domain's name, and the `nisdomainname` command is run to set it. The `/etc/yp.conf` file has a line added to it which specifies which server to use for that NIS domain. The `/etc/nsswitch.conf` file is modified to specify that NIS should be used as a source of information for password, shadow, and group lookups. Finally, `/etc/pam.d/system-auth` is modified so that password change requests for NIS accounts will be sent to the `rpc.yppasswdd` service running on the NIS master server.

NIS is a relatively insecure service. Kerberos authentication can be used in conjunction with NIS to improve password security. A better solution might use LDAP protected with TLS (SSL) encryption to store name service information in place of NIS. However, these solutions can be more complex to set up and manage.

## Example: LDAP Configuration

- Must install `nss-ldap` and `openldap` RPMs
- Run `system-config-authentication`
  - Enable LDAP to provide User Information
  - Specify server, the search base DN, and TLS
  - Enable LDAP to provide Authentication
- What does this actually do?
  - Four text-based configuration files are changed

Rev RH133-RHEL4-1

Copyright © 2004 Red Hat, Inc.

redhat 18

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 299 2994 or +1 (919) 754 3700.

LDAP is a protocol used to talk to a distributed directory service based on X.500, which can be used to store system and account information. Client programs can use the LDAP protocol to get information from directory servers running `slapd`, the standalone LDAP service. The OpenLDAP packages in the distribution provide a `slapd` server implementation as well as client tools and development libraries which can be used to work with LDAP services. Information is stored in the directory in individual *entries* organized into a hierarchical tree. Each entry and its location in the tree is uniquely identified by its *distinguished name*, or DN. A particular LDAP server is normally responsible for only the part of the tree under a particular entry, and the DN of this entry is normally used as a *base DN for searches* when looking up information stored in that server. (It may help to think of this base DN like a DNS domain name, and the entire LDAP directory tree like all of DNS.)

One use of a LDAP directory service is to synchronize account information between multiple networked systems. An individual entry may represent a single user (by using information from the relevant lines in `/etc/passwd` and `/etc/shadow`), or a single group (using information from the relevant line in `/etc/group`), or may store other information.

The easiest way to set up a client to use an existing LDAP server is to install the `nss_ldap` and `openldap` packages, and run `system-config-authentication`. To get NSS information from LDAP, under "User Information", enable LDAP, and then specify an LDAP server, and a base DN to use for searches. If the LDAP service provides standard password hashes to NSS, this may be sufficient. Alternatively, under "Authentication" you can also enable LDAP, which will let PAM test and change passwords by accessing the directory service using the DN of the account's entry and the password entered. If you select LDAP on the "Authentication" tab, it is **crucial** to also check "Use TLS to encrypt connections", or your unencrypted password will be transmitted over the network to the LDAP server as clear-text on every authentication!

What do these settings change? The `/etc/ldap.conf` file is modified to specify the location of your LDAP server, your search base DN, and whether or not TLS is enabled. The `/etc/openldap/ldap.conf` file is modified with the same information so that command-line OpenLDAP tools and the `automounter` can use the same server and base DN. The `/etc/nsswitch.conf` file is modified to specify that LDAP should be used as a source of information for password, `shadow`, and group lookups. Finally, `/etc/pam.d/system-auth` is modified so that PAM will try accessing the directory service as your account entry to test and change your password.



# File Ownership

- Every file has both user and group "ownership"
- A newly created file will be owned by:
  - the user who creates it
  - the current primary group of that user
    - SGID directories may change this behavior
- The chown command can be used by root to change ownership

S U G O  
7 7 7 7

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyright. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email training@redhat.com or phone toll free (USA) +1 866 295 2004 or +1 (919) 754-3700.

File ownership and permissions can be listed by using `ls` with the `-l` option (for "long" listing):

```
-rw-rw-r-- 1 joshua joshua 43557376 Sep 29 11:51 database
-rw-rw-r-- 1 joshua joshua 29958144 Sep 29 11:51 project
```

The files in the above listing, and any files created by `joshua`, are owned by user `joshua` and group `joshua`.

If `joshua` creates files in a directory with the SGID bit set, then those files will inherit the group ownership of that directory. The SGID permission will be explained in greater detail in the following pages.

```
-rw-rw-r-- 1 joshua webstaff 29958144 Sep 29 11:51 project
```

`chown` can be used by root to change the ownership of a file:

```
# chown bob /tmp/somefile
```

It can also be used to change the owner and group simultaneously:

```
# chown bob:webteam /tmp/somefile
```

Note that you may use a period "." in place of the colon

*used should joshua*

*-rwxrwxrwx  
# F S S t A  
C E T U S P E T A  
E R J O O P P S T T  
I T I T  
I*

*U G O*

*U+S  
G+S  
O+T*

*U CASE=NOX  
L CASE=X  
U digit  
= 237*

# Linux File Permissions

- Access levels
- Access modes
- Flags indicate access mode for each access level
- File mode is a concise collective expression of flags' values

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 263 2994 or +1 (619) 754 3700.

**Access levels.** There are three access levels; "user", "group" and "other" - usually referred to as "u", "g" and "o".

**Access modes.** There are three access modes; "read", "write" and "execute" - usually referred to as "r", "w", and "x"

- Read access on a file gives permission to read the contents. Read access on a directory gives permission to view the contents of the directory.
- Write access on a file gives permission to write the contents. Write access on a directory gives permission to remove or add files to the directory. Note that write permission on a directory is sufficient for a user to remove files, i.e. a user can remove files they do not own.
- Execute access on a file allows it to be treated as an executable program. To be successfully executed, the file contents must also be executable. Execute access on a directory gives permission to `cd` to a directory.

**Flags.** For each file there is a flag (a setting which is on or off) for each of the three basic access modes in each of the three access modes. For example, the following flags might be set: user read, user write, group read, group execute, others read

**File mode.** Collectively, the state of the nine flags is called the file mode, and can be concisely represented with a field of nine characters. From left to right, the first three are for user, the second three for group, and the last three for others. In these fields, a letter "r", "w" or "x" indicates read, write or execute respectively. If a flag is not set, the field contains a "-" (dash) character. With this scheme, the example file mode in the example above would be represented as `rw-r-xr--`.

Read access for a user is determined according to the following steps:

- a If the user (the person accessing the file) is also the owner then user-level access applies. If user-level access applies and the user read flag is set, then read access is available, otherwise read access is not available.
- b If user-level access does not apply, and the user's group is the same as the group with which the file is associated, then group-level access applies. If group-level access applies, and the group read flag is set, then read access is available, otherwise read access is not available.
- c If neither user-level access or group-level access apply, then others-level access applies. If others-level access applies, and the others read flag is set, then read access is available, otherwise read access is not available.

The same logic applies for both write access and execute access

An implication of this scheme that may not be obvious is that if user-level access applies without the user read flag being set, then read access is not available, regardless of the flags available at group level and others level. Likewise, if group-level access applies without group read being set, then read access is denied regardless of the others-level flags

## SUID / SGID Executables

- Normally processes started by a user run under the user and group security context of that user
- SUID and/or SGID bits set on an executable file cause it to run under the user and/or group security context of the file's owner and/or group

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2904 or +1 (319) 754 3700.

When a user starts a process, it runs with the permissions of that user. For example, if you run `vi`, and try to edit a file which you do not have permission to read or write, the operation will fail. However, if the SUID or SGID bit is set on an executable, it runs with the permissions of its owner (or group owner). For example, consider the file `/etc/shadow` that stores users' encrypted passwords:

```
-r----- 1 root root 805 Sep 29 11:19 /etc/shadow
```

The file is owned by `root`, who has exclusive read access. Users may still change their passwords with the `passwd` command, because the `passwd` command has its SUID bit set, and is owned by `root`:

```
-r-s--x--x 1 root root 13536 Jul 12 05:56 /usr/bin/passwd
```

SUID and SGID bits are set using the `chmod` command:

```
# chmod u+s <filename> (SUID)
# chmod g+s <filename> (SGID)
```

Since the SUID and SGID permissions are displayed overlying the execute permission for either user or group, respectively, the case of the indicates whether the execute permission is turned on or off. For example, if a capital is in the execute field, the SUID or SGID permission is on and the execute permission is off. If a lowercase is in the execute field, both the SUID or SGID and the execute permissions are on.

## The Sticky Bit

- Normally users with write permissions to a directory can delete any file in that directory regardless of that file's permissions or ownership
- With the sticky bit set on a directory, only the owner of a file can delete the file
- Example: /tmp

```
drwxrwxrwt 12 root root 4096 Nov 2 15:44 tmp
^
```

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat



22

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 260 2694 or +1 (619) 754 3700.

It may seem counterintuitive that write permissions on a directory would allow one user to delete another user's file within that directory. Consider, however, that a directory is really just a file itself whose contents are references to other files. Deleting a file is therefore an edit to a directory file's list of other files

Many users need to be able to create and delete files in /tmp. Even if users do not actively create files in /tmp, many applications they run will use /tmp as a location for temporary files. Setting the sticky bit prevents users from deleting each others' files, even though they have full access to the directory.

Note that the sticky bit on /tmp is set by default, and can be seen as a "t" in the file permissions:

```
drwxrwxrwt 13 root root 4096 Sep 29 12:42 tmp
```

To set the sticky bit on a directory, use `chmod`:

```
# chmod o+t /home/share
```

The sticky bit will appear as a T if the directory's execute permission for "others" is off.

## The Setgid Access Mode

- Normally, files created in a directory belong to the default group of the user
- When a file is created in a directory with the setgid bit set, it belongs to the same group as the directory

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2894 or +1 (610) 754 3700.

When a file is created in a directory, it belongs to the primary group of the user that created the file. However, if the setgid bit is set for the directory, new files that are created in the directory have their group ownership set to the same group as the owner of the directory. This provides a mechanism to allow one level of access to users, who are members of the same group that owns the directory while allowing a different level of access to non-member users of files in the directory.

Recall that the setgid bit can be set with `chmod`:

```
# chmod g+s <directory>
```

This sets the setgid bit without affecting current permissions while

```
# chmod 2770 <directory>
```

sets the setgid bit and gives read, write, and execute permissions to the owner of the directory and members of the group whose ownership is on that directory.

## Default File Permissions

- Read and write for all is the default for files
- Read, write and execute is the default for directories
- `umask` can be used to withhold permissions on file creation
- Non-system users' `umask` is 002
  - Files will have permissions of 664
  - Directories will have permissions of 775
  - Supports user private groups
- System users' `umask` is 022

Rev RH133/RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 285 2954 or +1 (919) 754 3700.

Without a `umask` in effect, any file created will have 666 permissions. This means that anyone on the system will have read and write access to any newly-created file. In order to withhold permissions we use a `umask`. The `umask` lists the permissions to withhold. A `umask` of 002 will result in files created with 664 permissions.

The default `umask` on a Red Hat Enterprise Linux system is 002. To change your `umask` to 022, use the `umask` command:

```
$ umask 022
```

Note that when setting your `umask` from the command line, the next time you login your `umask` will be set back to your default, so this is typically set by one of the `bash` initialization scripts.

## Access Control Lists (ACLs)

- Grant RWX access to files to multiple users or groups

```
mount -o acl
getfacl file|directory
setfacl -m u:gandolf:rwX
setfacl -m g:nazgul:rw
setfacl -m d:u:frodo:rw
setfacl -x u:samwise
```

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2954 or +1 (818) 754 3700.

The ext3 filesystem includes support for access control lists which allow finer grained control of file system permissions than are possible with the standard three access categories that are normally provided. Many filesystem commands, such as cp and mv, have been modified to copy the associated ACL's for a file.

In order to enable ACLs on a file system, the file system must be mounted with the -o acl mount option.

To view the ACL's for a file, use the getfacl command:

```
getfacl /tmp/schedule.txt
```

ACL's can be set using the setfacl command:

```
setfacl -m u:visitor:rx /tmp/schedule.txt
```

The above command would grant the user visitor read and execute access to the file schedule.txt in the /tmp directory.

To remove an ACL:

```
setfacl -x u:visitor /tmp/schedule.txt
```



# SELinux

- Each process or object (file, directory, network socket) also has a SELinux context.
  - identity:role:domain/type
- The SELinux policy controls
  - what identities can use which roles
  - what roles can enter which domains
  - what domains can access which types

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (616) 754-3700.

SELinux adds another layer of access control permissions on top of standard file permissions and ACLs, which are defined by the system's security policy. Each process or object (such as a file, directory, or network socket) on the system also has a SELinux security context. This context consists of a SELinux user identity, a role, and a domain (for processes) or a type (for objects).

The policy controls what SELinux identity a process is assigned when it starts. The identity determines which roles are accessible to the process, and the roles are used to determine which domains the process can switch to. Once running in a particular domain, the policy also determines what access a process will have to objects of particular SELinux types. The policy also determines the default type for a new object when it is created.

In general, the access a process will be granted to an object is determined by the domain of the process and the type of the object being accessed.

The default policy is set by the contents of the `selinux-policy-targeted` RPM.

# Controlling SELinux

- `system-config-securitylevel`
- `setenforce` and `setsebool`
- `/etc/sysconfig/selinux`
- `enforcing=0`
- `/selinux` virtual file system

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 256 2894 or +1 (919) 754 3700.

The SELinux policy may be adjusted or disabled through a number of utilities. The easiest to use is the graphical `system-config-securitylevel` tool, which can turn on or off SELinux or place it into a warn-only mode. It also allows the adjustment of "booleans" which can fine-tune the rules enforced by the policy.

SELinux enforcement may also be changed from on to warn-only and vice versa with the `setenforce` command-line tool. The file `/etc/sysconfig/selinux` can be edited to make enforcement changes persist across reboot. The `setsebool` command-line tool can be used to adjust and save booleans, and `sestatus` will print out the current state of SELinux to standard output.

The kernel option `enforcing=0` can be passed through GRUB at boot time to put the kernel in warn-only mode; `enforcing=1` puts it in enforcing mode.


The `/selinux` virtual filesystem is similar to `/proc` and `/sys`. It presents information about the state of SELinux in the kernel to user programs like the ones above.

## SELinux Contexts

- List process contexts: `ps -Z`
- List file contexts: `ls -Z`
- Change file contexts: `chcon`  
`chcon -t httpd_sys_content_t index.html`  
`chcon --reference=/var/www/html index.html`

Rev RH133RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

Process contexts can be listed by adding the `-Z` option to the `ps` command. Most processes run in the `unconfined_t` domain, which means that they are not restricted by the default SELinux policy. Processes running in other domains are most likely restricted by the default policy. Services affected include `dhcpcd`, `httpd`, `mysqld`, `named`, `nscd`, `ntpd`, `portmap`, `postgres`, `snmpd`, `squid`, `syslogd`, and `winbind`.

The security contexts of files can be displayed through the `ls -Z` command. A newly created file is assigned the context of its parent directory unless the policy specifies otherwise. The `chcon` command changes the context of a file, and works much like `chown` and `chmod`. Generally, the type of a file is the critical part of its context to set. To recursively set the type of all files in `/var/www/html` to `httpd_sys_content_t`, without chasing symlinks, run the command

```
chcon -R -t httpd_sys_content_t -h /var/www/html
```

The `--reference` option can be used with `chcon` to apply the current SELinux context of one file or directory to another file or directory.

# Troubleshooting SELinux

- What is the error?
  - Check `/var/log/messages` for AVC denials
- Is the process doing something it shouldn't?
- Does the target have the right context?
- Does a boolean setting need adjustment?

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2894 or +1 (919) 754 3700.

SELinux policy violations are logged to `/var/log/messages`. An error might read like the following:

```
Feb 25 01:38:23 station15 kernel: audit(1109313503 808:0): avc: denied { read } for pid=4346 exe=/usr/sbin/httpd
name=joe dev=hda2 ino=311297 scontext=root:system_r:httpd_t tcontext=system_u:object_r:user_home_dir_t tclass=dir
```

This translates as:

PID 4346, a `/usr/sbin/httpd` process, with the context `root:system_r:httpd_t` was denied read access to a directory, named `joe`, which is inode 311297 on `/dev/hda2`, which has the context `system_u:object_r:user_home_dir_t`

At this point several questions must be asked. Is the process being blocked for legitimate reasons -- is it doing something inappropriate? If not, then is the target's context wrong? If so, the correct context needs to be determined and set with `chcon`. If the policy is being too strict, perhaps a "boolean" setting can be adjusted with `system-config-securitylevel` or `setsebool`. In the worst case, perhaps SELinux can be disabled for just the affected service, or entirely.

Resources that can help troubleshoot SELinux problems include the Red Hat Enterprise Linux 4: Red Hat SELinux Guide on [www.redhat.com](http://www.redhat.com), and [Understanding and Customizing the Apache SELinux Policy for Fedora Core 3](http://www.redhat.com) at [fedora.redhat.com](http://fedora.redhat.com). The source used to build the SELinux policy is included in the `selinux-targeted-policy-sources` RPM.

## End of Unit 7

- Questions and answers
- Summary
  - User & Group accounts
  - File ownership and permissions
  - Extended file modes: SUID / SGID / Sticky
  - Switching accounts with `su`
  - `umask` and the UPG scheme
  - Shell environment

Rev RH133-RHEL-4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2884 or +1 (619) 794-3700.

### Important files covered in this Unit:

```
/etc/passwd
/etc/group
/etc/skel
/etc/profile
/etc/bashrc
-/.bashrc
-/.bash_profile
/etc/sysconfig/selinux
```

### Important commands covered in this Unit:

```
useradd, usermod, userdel, newusers
groupadd, groupmod, groupdel
chage
system-config-authentication
chown, chgrp, chmod
umask, su
chcon, setenforce, setsebool, sestatus, system-config-securitylevel
```

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100

100

100

100

# Unit 7 Lab

## User and Group Administration

---

Estimated Duration: 45 Minutes

Goal: To build skills at user and group administration

Setup at Start: A Red Hat Enterprise Linux System

Situation: You need to setup groups and accounts for several users in your company. Each department also requires a shared directory for their files.

**Sequence 1: Creating the groups and users****Scenario/Story:**

You need to set up groups for different departments in your company. You also need to set up user accounts for the employees in those departments.

**Tasks:**

1. Using the `useradd` command, add accounts for the following users to your system: joshua, alex, dax, bryan, zak, ed, and manager. Remember to give each user a password.
2. Using the `groupadd` command, add the following groups to your system. Use the `-g` option to set the correct GID

group	gid
====	===
sales	10000
hr	10001
web	10002

Why should you set the GID in this manner instead of allowing the system to set the GID by default?

3. Using the `usermod` command add joshua and alex to the sales auxiliary group, dax and bryan to the hr auxiliary group. Add zak and ed to the web auxiliary group. Add manager to all auxiliary groups. The `-G` option can be used to add users to supplemental groups.
4. Login as each user and use the `id` command to verify that they are in the appropriate groups. How else might you verify this information?

**Deliverable:**

A system with users joshua and alex in the sales group, dax and bryan in the hr group, zak and ed in the web group, and manager in the sales, hr, and web groups.



**Sequence 2: Setting up shared directories****Scenario/Story:**

Each department for which you created a group also needs a shared directory. This will allow users in each department to share files, but will prevent users in other departments from altering, or even seeing those files.

**Tasks:**

1. Create a directory called `/depts` with a `sales`, `hr`, and `web` directory within the `/depts` directory.

```
mkdir -p /depts/{sales,hr,web}
```

2. Using the `chgrp` command, set the group ownership of each directory to the group with the matching name as in the example:

```
chgrp sales /depts/sales
```

3. Set the permissions on the `/depts` directory to `755`, and each subdirectory to `770`

4. Set the `sgid` bit on each departmental directory so that files created within those directories will be owned by the appropriate group.

```
chmod g+s /depts/*
```

5. Experiment by logging in as each user and creating or altering files in each of the directories. Only manager should be able to enter all the directories. You can also use `su -`, but make sure to include the dash and to exit one `su` session before starting another.

**Deliverable:**

A shared directory for each department that allows only users in that department to enter it or create, view, and alter files within.

**Sequence 3: Client-side NIS account management****Task: Configure your system as an NIS client system****Scenario/Story:**

Your site is centrally managing user account information in a NIS directory service on server1.example.com, using the NIS domain name notexample. Set up your client to get user information from NIS, authenticating users with the NSS information provided by the NIS server. In addition, home directories for these users are exported through NFS to all workstations as well. You should set up the automounter to automatically mount and unmount these directories as needed.

**Tasks:**

1. Install the following RPMs on your system if they are not already installed: portmap, ypbind, yp-tools, authconfig, authconfig-gtk
2. Use system-config-authentication to configure your client to use NIS for authentication.

If you are using the graphical interface, check the option "Enable NIS Support" on the "User Information" tab. You can leave the "Authentication" tab alone. Click the "Configure NIS..." button, and in the "NIS Settings" window specify notexample for "NIS Domain" and server1.example.com for "NIS Server". Click on the "OK" button for the "NIS Settings" window, and again on the main window.

If you are using a (text-based) virtual console or the --nox option, then under "User Information" check "Use NIS". Under "Authentication" leave "Use MD5 Passwords" and "Use Shadow Passwords" checked. Then select the "Next" button, and on the next screen specify 'notexample' for "Domain" and 'server1.example.com' for "Server". Select "Ok".

Either way, ypbind should now start up successfully. If it does not, check /var/log/messages for errors

3. Restart the sshd service to make sure that it registers the changes to authentication: service sshd restart.
4. Run 'getent passwd'. You should see all the /etc/passwd lines for local users on your system, followed by user information from the NIS server's passwd maps.
5. Next, try to log in as one of the NIS users. Switch to virtual console 5 by typing Ctrl-Alt-F5. (You can switch back to X11 by typing Ctrl-Alt-F7.) Assuming that you have not used this console already, it should be at a text login prompt. Log in as guestN, where N is your station number plus 2000. (So, if you are on station15, you should use guest2015.) The account's password is 'password'. You should be able to log in, but there will be an error because your account's home directory does not exist on your local system. Your shell should start in / instead. Log out as the NIS user

- 6 Use the automounter to mount the home directories for your NIS users from server1 example.com. Begin by editing `/etc/auto.master` to add the following line:

```
/home/guests /etc/auto.guests --timeout=60
```

This line specifies that `/etc/auto.guests` defines mount points in `/home/guests` managed by the automounter. When not in use for more than 60 seconds, filesystems mounted on those mount points are automatically unmounted.

- 7 Create and edit `/etc/auto.guests` so it contains the line

```
* -rw,soft,intr 192.168.0.254:/home/guests/& *
```

This line specifies that accesses to any immediate subdirectory of `/home/guests` should make `autofs` mount a NFS export from `192.168.0.254` where the `&` is the same as the name of the local subdirectory. (So the automounter would mount `192.168.0.254:/home/guests/foo` on `/home/guests/foo`.) The middle column specifies the mount options that will be used; read-write, timeout eventually if the NFS server is not available, and timeout immediately if an interrupt is sent.

- 8 Configure `autofs` to start in run levels 3, 4, and 5, then start it manually:

```
chkconfig autofs on  
service autofs start
```

- 9 Now try logging in again and see whether the home directory gets mounted automatically. It should. Try logging into to your neighbors system once it is also configured. You should be able to access your home environment from any system in the `notexample` domain.

**Sequence 4: Client-side LDAP account management****Scenario/Story:**

Your site is migrating new user account information into a LDAP directory service, also served by server1.example.com under the search base dc=example,dc=com. The administrator of the LDAP server requires clients to use TLS (SSL) encryption, and the LDAP server's TLS certificate is digitally signed by a locally-run TLS Certificate Authority so that clients may verify that they are communicating with the real LDAP server.

Using TLS encryption, set up your client system to get user information and authenticate users using the LDAP server.

Home directories for these LDAP-managed users use the same NFS export as your old NIS-managed users, so you will not have to make any changes to your automounter configuration.

**Tasks:**

1. Install the following RPMs on your system if they are not already installed: `openldap`, `openldap-clients`, `nss_ldap`.
2. Use `system-config-authentication` to configure your client to use LDAP for authentication. In the graphical interface, on the "User Information" tab check "Enable LDAP Support". Switch to the "Authentication" tab and also check "Enable LDAP Support".
3. Select the "Configure LDAP..." button on either tab. On the window that opens, set your "LDAP Search Base DN" to `dc=example,dc=com`. Set your "LDAP Server" to `server1.example.com`. Finally, select "Use TLS to encrypt connections." Using TLS is critical for security if you enable LDAP on the Authentication tab, because the PAM module used for LDAP will send passwords as cleartext over the network to the directory server on each authentication if this is not selected.
4. Your site uses a local certificate authority to allow clients to verify that they are getting their TLS certificate from the real server. You need to download and install the local CA's public certificate to make this work. This certificate is at <http://server1.example.com/pub/EXAMPLE-CA-CERT>. Download and copy this file to `/usr/share/ssl/certs` on your client system. Then

```
chown root:root /usr/share/ssl/certs/EXAMPLE-CA-CERT
chmod 644 /usr/share/ssl/certs/EXAMPLE-CA-CERT
```

5. Next, you need to make sure that the client can find and use the CA certificate. Open `/etc/ldap.conf` in a text editor and make sure the following two directives appear in that file as listed below:

```
tls_checkpeer yes
tls_cacertfile /usr/share/ssl/certs/EXAMPLE-CA-CERT
```

Save the file and exit.

6. Restart the `sshd` service to make sure that it registers the changes to authentication: `service sshd restart`.

7. Run 'getent passwd'. You should see all the /etc/passwd lines for local users, followed by user information from the NIS maps set up in the previous sequence, followed by user information from the LDAP directory
8. Finally, try to log in as one of the LDAP users. Switch to an unused virtual console's login prompt. Login as the user ldapuserX, where X is your station number (so if you are on station15, use ldapuser15) The password for these accounts is 'password'

If you see any errors, look at the logs in /var/log/messages and /var/log/secure. You can also try running the command 'ldapsearch -x -Z', which should dump user information in LDIF format from your LDAP server to standard output if the server is reachable.


1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100

# UNIT 8

## Printing and Administration Tools

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 296 2094 or +1 (919) 754-3700.

## UNIT 8: Objectives

- Upon completion of this unit you should be able to:
  - Configure printing
  - Perform task automation with cron
  - Configure system logging
  - Perform backup and restore

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (519) 754 3700.



## UNIT 8: Agenda

- Printing
- Task Automation
- System Logging
- Backup and Restore

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

## Controlling Access to cron

- Restrict / allow user access to cron
  - /etc/cron.allow
  - /etc/cron.deny
- Contains usernames to allow / deny access

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



8

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (919) 754 3700.

### **cron access control**

If the file `cron.allow` exists and your username appears in it, you may use the `crontab` command. If the `cron.allow` file does not exist and the file `cron.deny` does, then you must not be listed in `cron.deny` to use `crontab`. If neither file exists, the default behavior is to allow all users to schedule jobs with `cron`.

Note that denying a user through the use of the above files does not disable their installed `crontab`.

## System `crontab` Files

- Different format than user `crontab` files
- Master `crontab` file `/etc/crontab` runs executables in
  - `/etc/cron.hourly`
  - `/etc/cron.daily`
  - `/etc/cron.weekly`
  - `/etc/cron.monthly`
- `/etc/cron.d/` directory contains additional system `crontab` files

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat  9

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2994 or +1 (919) 754 3700.

The format of `/etc/crontab` and the files in `/etc/cron.d` are different from user `crontab`s. The sixth field is a username which will be used to execute the command in the seventh field.

A common command in these files is the `run-parts` shell script. This script takes one argument, a directory name, and invokes all of the programs in that directory (The `run-parts` script is located in `/usr/bin` and has no online documentation.)

The following is an example `/etc/crontab` file:

```
SHELL=/bin/bash
PATH=/sbin:/bin:/usr/bin:/usr/sbin
MAILTO=root
HOME=/

# run-parts
01 * * * * root run-parts /etc/cron.hourly
02 4 * * * root run-parts /etc/cron.daily
22 4 * * 0 root run-parts /etc/cron.weekly
42 4 1 * * root run-parts /etc/cron.monthly
```

Thus, at 4:02 every morning, all of the executables in the `/etc/cron.daily` directory will be run as `root`. A default installation's `cron.daily` directory will contain scripts to update the `slocate` and `whatis` databases, to clean up temporary directories, and to perform other housekeeping tasks.

## System Cron Job: tmpwatch

- Cleans old files out of specified directories
- Useful for keeping /tmp directory from filling up
- tmpwatch is run daily in /etc/cron.daily

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2884 or +1 (919) 754 3700.

The contents of the /etc/cron.daily/tmpwatch file installed by RHEL:

```
/usr/sbin/tmpwatch 240 /tmp
/usr/sbin/tmpwatch 720 /var/tmp

for d in /var/{cache/man,catman}/{cat?,X11R6/cat?,local/cat?}
do
    if [ -d "$d" ]; then
        /usr/sbin/tmpwatch -f 720 $d
    fi
done
```

As shown above, tmpwatch can be used for cleaning a variety of directories in the Linux filesystem.

Usage:

```
tmpwatch [-u|-m|-c] [-adfqtv] [--verbose] [--force] [--all] \
[--nodirs] [--test] [--quiet] [--atime|--mtime|--ctime] \
<hours-untouched> <dirs>
```

In the cron script above, tmpwatch is instructed to clean up files which are more than 240 hours (10 days) old. This means that any file in the above directories that has not been accessed for more than 10 days will be deleted.

## System Cron Job: logrotate

- Maintain with **logrotate**
  - Keeps log files from getting too large
  - Keeps filesystem from filling up
- **logrotate** is run daily in `/etc/cron.daily`
- Highly configurable
  - Configure all logs in `/etc/logrotate.conf`
  - Configure individual log files in files within `/etc/logrotate.d`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 290 2994 or +1 (919) 754 3700.

Left unchecked, system logs will grow until you run out of disk space. Red Hat Enterprise Linux ships with **logrotate**, a powerful tool for log file maintenance:

- Log files from different subsystems are rotated at predefined intervals, or when they reach predefined sizes, and old logs are optionally compressed.
- When you install a new system service by RPM, it should preconfigure itself automatically for log file rotation.
- You should monitor the system logs in `/var/log` and increase your rotational frequency if logs are becoming too large.
- For example, `/var/log/messages` is rotated weekly to `/var/log/messages.1`, with older log files rotated to `/var/log/messages.2`, etc, optionally compressed.
- Rotational configuration is stored in `/etc/logrotate.conf`, for general settings, and
- `/etc/logrotate.d/subsystem`, for subsystem-specific settings.

For more information see the **logrotate** man pages.

## System Cron Job: **logwatch**

- Monitor with **logwatch**
  - Helps catch problem issues
  - Detects suspicious behavior
- **logwatch** is run daily in `/etc/cron.daily`
- Configuration file:  
`/etc/log.d/conf/logwatch.conf`
- Sends nightly email report
- Other tools

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2694 or +1 (919) 754 3700.

Monitoring system logs is an onerous but important task. If you do not properly monitor your logs, you may miss security problems, hardware problems, or software problems. For example, a system may have a runaway maintenance problem every Monday at 5 AM that filled up the filesystem and them promptly cleaned it up again. Only the system log would reveal that there was a problem

**logwatch**, installed by default on most Red Hat Enterprise Linux systems, monitors log files, reporting nightly on activity, and, potentially, on any anomalies located. **logwatch** is highly configurable. It can be programmed to detect most any type of activity. See `/usr/share/doc/logwatch-version` for information on writing log filters

Other tools can display log files in real time, such as `tail -f`. Several log files can be monitored simultaneously with a command similar to the following:

```
cd /var/log; tail -f messages -f maillog -f secure
```

This will display log messages from any of these files as the files grow, preceding each change of log file with an indicator of the log file reporting the message

222

# System Logging

- Centralized logging daemons: `syslogd`, `klogd`
- Log file examples:
  - `/var/log/dmesg`      Kernel boot messages
  - `/var/log/messages`   Standard system error messages
  - `/var/log/maillog`    Mail system messages
  - `/var/log/secure`     Security, authentication, and `xinetd` messages
- Application log files and directories also reside in `/var/log`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2994 or +1 (919) 754 3700.

A variety of log files are maintained by the system, an understanding of which is often vital for troubleshooting system problems:

*hardware*

`/var/log/dmesg` - This logfile is written upon system boot. It contains messages from the kernel that were raised during the boot process.

*services*

`/var/log/messages` - This is the standard system logfile, which contains messages from all your system software, non-kernel boot issues, and messages that go to `dmesg`. Readable only by root.

*mail server*

`/var/log/maillog` - This logfile contains messages and errors from your sendmail. Readable only by root.

*security*

`/var/log/secure` - This logfile contains messages and errors from security-related systems such as `login`, `tcp_wrappers`, and `xinetd`. Readable only by root. Very useful in detecting and investigating network abuse.

There are also various other system logfiles that store information from other applications (i.e. Apache, Squid, etc) that also may be found under `/var/log`

*honeypot.rules.org*

*how to build a network of honeypots under 1000 dollars*

## syslog Configuration

- `syslog` System V initialization script in `/etc/rc.d/init.d` controls both the `syslogd` and the `klogd` daemons
- `/etc/syslog.conf`
  - Configures system logging
- `/etc/sysconfig/syslog`
  - Sets switches used when starting `syslogd` and `klogd` from the System V initialization script

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2094 or +1 (619) 754 3700.

RHEL provides a central system logging facility that allows all applications to store debugging information, errors and messages in a central, manageable place. System logging is provided by `syslogd`. `klogd` intercepts kernel messages and provides them to `syslogd`. `syslogd` is configured in `/etc/syslog.conf`.

Messages can be logged to files, broadcast to connected users, written to the console, or even transmitted to remote logging daemons across the network. Messages for system logging all have an associated severity. You may choose not to log less-severe messages, or to log severe messages in a separate file. On most systems however, the default settings are adequate

By default, messages of emergency or higher (more severe) are broadcast to all users, and most other messages are written to `/var/log/messages`, which is where you should look for non-kernel boot errors, error messages from most application-level services, such as automount, login services, etc. After system boot, kernel messages are also written to this file.

### `syslog.conf` Format:

Each entry in the system log file has four main entries:

The date and time of the message

The hostname from whence the message came. This field is important when logging across a network to a centralized log host

The name of the application or subsystem from whence the message came; for example, `kernel`, `ftpd`, etc. This may include the process identifier.

The remainder of the line (following the colon) is the actual message itself.

Repeated log messages will be marked as such.



# Tape Drives

- SCSI tape devices (i.e., DDS, DLT)
  - `/dev/[n]st0`, `/dev/[n]st1`, etc.
  - devices with 'n' do not automatically rewind
- Use the `mt` utility to control tape drive
  - `mt -f /dev/st0 rewind` (Rewind)
  - `mt -f /dev/st0 fsf 50` (Position)
  - `mt -f /dev/st0 offline` (Eject)
  - `mt -f /dev/st0 erase` (Erase)
  - `mt -f /dev/st0 rewoff` (Rewind, Eject)

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 288 2994 or +1 (619) 754 3700.

SCSI tape drives will be in use in most commercial settings. These drives will often be DAT (Digital Audio Tape) tape drives, sometimes called DDS (Digital Data Storage)

Tape devices whose names beginning with *n* refer to no rewind devices. When a tape device is referenced with a no rewind name it will not automatically rewind when the device is closed, e.g., when a utility such as `dump` is finished writing.

`mt`'s general syntax looks like the following:

```
mt -f /dev/st0 rewind
```

This syntax consists of a tape device and an `mt` command (such as `rewind`)

The device given as the argument with the `-f` option is determined by the target tape device:

The standard SCSI tape devices are named `st0`, `st1`, ..., `nst0`, `nst1`, ...

The standard IDE tape devices are named `ht0`, `ht1`, ..., `nht0`, `nht1`, ...

The standard floppy tape devices are named `ftape` (`rft0`) and `ntape` (`nrft0`)

If an environment variable `TAPE` is set, `mt` will use its value for the device if no `-f device` argument is supplied.

`mt` is used in backup scripts in many cases, but may be used from the command line as well. RHEL also includes remote version of `mt` named `rmt`. As with most of the so-called "r" commands, the use of `rmt` is not recommended in security-conscious environments

## Using tar/star

- Archives to tapes or other media or files
- **star** backs up SELinux contexts and ACL attributes
- Parameters:
  - c create x extract
  - t list v verbose
  - z gzip compression j bzip2 compression
- Examples:

```
cd /tmp && tar xvf ~/archive.tar
tar cvf /dev/st0 /data /foo /bar
```

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754-3700.

**tar** can be used to create archives on tape, floppy, or Zip (or other removable device) in the same way it is used to create file archives.

To back up a file or directory to SCSI tape:

```
tar cf /dev/st0 file_or_directory
```

**tar** can be used with compression. To use gzip compression, use the following syntax:

```
tar zcf /dev/st0 file_or_directory
```

To extract a compressed archive, use:

```
tar zxf /dev/st0
```

To use **tar** with a removable device, such as a Zip or Jaz drive, simply point **tar** at the device:

```
tar cf /dev/<device> file_or_directory
```

## Using dump/restore

- Back up and restore ext2/3 filesystems
  - Does not work with other filesystems
  - `dump` should only be used on unmounted filesystems or filesystems that are read-only.
- Can do full or incremental backups
- Examples:

```
dump -0u -f /dev/nst0 /dev/hda2
restore -rf /dev/nst0
```

*of level backup?*

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 295 2994 or +1 (919) 754 3700.

### Using `dump`

The `dump` command can be directed to read `/etc/fstab` and do backups based on information it keeps as to which filesystems need to be backed up.

For example,

```
dump -0u -f /dev/nst1 /home
```

*Don't rewind after*

will do a full backup of the `/home` filesystem onto the tape device `nst1`. The `-u` option will update the file `/etc/dumpdates`, which will record dump information for future use by `dump`. After a level 0 backup, `dump` will perform an incremental backup every day on active filesystems listed in `/etc/fstab`.

The command:

```
dump -4u -f /dev/nst1 /home
```

will perform an incremental update of all files that have changed since the last backup of level 4 or lower, as recorded in the `/etc/dumpdates` file

### Using `restore`

To restore data backed up with `dump`, make a clean filesystem (using `mkfs`), mount the filesystem, and `cd` to the directory where the filesystem is mounted. Then run the `restore` command:

```
restore -rf /dev/nst1
```

## Using `cpio`

- Similar to `tar`
  - Does not recurse directories by itself
  - Can archive special files
  - Piping output from `find` into `cpio` is common
- Examples:

```
find /data | cpio -ocv > /dev/nst0
cpio -icdvm < /dev/nst0
cpio -tvf < mybackup.cpio
```

*↓  
Test*

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

18

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 266 2994 or +1 (519) 754 3700.

`cpio` copies files into or out of a `cpio` or `tar` archive, which is a file that contains other files plus information about them, such as their file name, owner, timestamps, and access permissions. The archive can be another file on the disk, a magnetic tape, or a pipe. `cpio` has three operating modes:

In copy-out mode, `cpio` copies files into an archive. It reads a list of filenames, one per line, on standard input and writes the archive onto standard output.

```
find /tmp | cpio -ocv > /dev/nst0
```

In copy-in mode, `cpio` copies files out of an archive or lists the archive contents. It reads the archive from standard input.

```
cpio -icdvm < /dev/nst0
```

*copy in*

`cpio` also has a copy-pass mode, which copies files from one directory tree to another. It combines the copy-out and copy-in steps, without using an archive file.

## Remote Backups

- `dump` and `tar` can use `rmt` (remote tape mgr)  
`dump -0uf joe@svr:/dev/nst0 /home`
- Use `user@host:path` format to specify the remote user, host, and device.
- `dump` can use `ssh` for secure backups when `RSH` environment variable is set to `ssh`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2094 or +1 (610) 754 3700.

Remote backups come in very handy when only one computer on the network has a tape backup device. In this situation, all the other machines can use `dump` or `tar` and `rmt` to send their backups to the network backup server. Using the `ssh` option, backups can even be made securely over the Internet.

## Other Backup Software

- Higher-level applications for tape backup include:
  - Amanda
    - Highly-scalable command-line client-server archiver included with RHEL
  - Commercial applications
    - Arkeia, Bru, Tivoli, Veritas (client), UNiBACK, ArcServe

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 255 2904 or +1 (515) 754 3700.

Tools such as `tar` and `dump` are considered an essential part of any Linux or UNIX system, and they can fulfill their role as backup and restore utilities well in most cases. In some cases, organizations may elect to use a more comprehensive tool. The tool may have more extensive network backup capabilities, may provide a GUI, or may have more extensive data verification routines.

A number of backup applications are available for RHEL. Amanda, an open source product, is part of the distribution. Commercial applications which are available include Arkeia, Bru, Tivoli, Veritas, UNiBACK, and ArcServe.

## End of Unit 8

- Questions and answers
- Summary
  - Configuring Printing
  - Task Automation
  - Configuring System Logging
  - Backup and Restore

Rev 09/133/RHEL4-1


Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 206 2084 or +1 (918) 754 3700.

### Important files covered in this Unit:

```
/etc/cups/cupsd.conf
/etc/cups/printers.conf
/var/spool/cron
/etc/crontab
cron.{allow,deny}
/etc/cron.daily/tmpwatch
/etc/syslog.conf
/var/log/messages
/var/log/dmesg
/var/log/secure
/var/log/maillog
/etc/logrotate.conf
```



Handwritten text, possibly a date or page number, located at the top of the page.

Handwritten text, possibly a name or title, located below the top line of text.

Handwritten text, possibly a date or page number, located in the middle of the page.

Handwritten text, possibly a name or title, located below the middle line of text.

Handwritten text, possibly a date or page number, located at the bottom of the page.



# Unit 8 Lab

## Printing and Administration Tools

---

Goal: Develop skills using system administration tools and setting up and administering CUPS.

Setup at Start: Running Red Hat Enterprise Linux System, logged on as the root user

Situation: Not Applicable

**Sequence 1: Using cron****Scenario/Story:**

You need to schedule another job to run every ten minutes today between the hours of 0800 (8:00 AM) and 1700 (5:00 PM).

**Tasks:**

- 1 You want to know some information about the status of the system every ten minutes today to help investigate some performance issues you have been having. You suspect it might be memory or I/O related and want to keep an eye on those resources. As root, use the command `crontab -e` to edit your cron file (If you are not comfortable with `vi`, export an `EDITOR` environment variable to set another editor.)
- 2 Enter the following line in your crontab file:  

```
*/10 8-17 * * * /usr/bin/free; /bin/ps
```
- 3 How could you send the output from these cron jobs to another e-mail address?  
*edit crontab + change mail to*
- 4 Use `mail`, or `mutt` as root to check for email from the `at` and cron jobs you have scheduled
- 5 Be sure to delete your cron job when you have received several emails from it

**Sequence 2: Logging to a centralized loghost****Scenario/Story:**

Your boss thinks it is a great idea to have one central logging host.

**Tasks:**

Work together with your neighbor.

- 1 First, set up `syslogd` to accept remote messages. Edit `/etc/sysconfig/syslog`:

```
SYSLOGD_OPTIONS="-r -m 0"
```

2. Restart `syslogd`:

```
service syslog restart
```

Now your machine will accept logging messages from other machines.

- 3 Set up `syslogd` to send some messages to another machine. Append in `/etc/syslog.conf` the following line:

```
user.* @stationX
```

Where `stationX` is your neighbor's machine.

- 4 Restart `syslogd` again:

```
service syslog restart
```

Now your machine sends messages from user programs to your neighbor's machine.

5. Test the new setup by using `logger` to generate a `syslog` message:

```
logger -i -t yourname "This is a test"
```

Does the message appear in your neighbor's `/var/log/messages` ?

**Sequence 3: Restoring individual files with `dump/restore`**

1 Prepare to use `dump` to back up the files in `/boot`. Use `df /boot` to find out which device `/boot` is on. (In the examples below, we will assume this device is `/dev/hda1`.)

2. First, estimate how much space you will need on the backup media. Estimate the size in bytes of a level-zero dump of `/boot` by using the `-S` option:

```
# dump -0S /dev/hda1
```

3. Rather than backing up to a tape, back up the data to a dump file. Check to make sure there is enough room in `/var/tmp` to hold the file, then execute

```
# dump -0u -f /var/tmp/dumpfile /dev/hda1
```

4. Look in `/etc/dumpdates` and see how `dump` recorded the timestamp of the full backup.

5. Now use `restore` to view the contents of the dump file:

```
# restore -tf /var/tmp/dumpfile
```

6. `restore` has an interactive mode we can use to restore only selected files. Restore to a temporary directory:

```
# mkdir /tmp/restored; cd /tmp/restored  
# restore -if /var/tmp/dumpfile
```

7. You should have a `restore >` prompt. Type `help` to see the list of available commands. Use `ls` and `cd` to view and navigate the list of backed-up files.

8. Use `add` to include `/grub/menu.lst` and `/grub/grub.conf` in the list of files to extract. List the directory with those files again. Files marked for extraction should be marked with an asterisk.

9. Type `extract` to restore the selected files. Your next volume number is "1", the first tape (or dump file). Do not set owner/mode for '.' (the directory the files are being extracted into). Quit `restore`.

10. There should now be a `grub` directory (like there is under `/boot`) inside the directory where you ran `restore` that contains your restored `grub.conf` and `menu.lst` files.

**Sequence 4: Setting up a printer and administering a printer with CUPS.**

1. As root run `system-config-printer` in a virtual console (not an X terminal).
2. Select *New* and press *Enter*.
3. Enter the text: `lp0` in the Queue Name field.
4. Select Local Printer Device for the Queue Type.
5. Select *Next* and press *Enter*.
6. Select `/dev/lp0` and select *Next*.
7. Select Postscript Printer and select *Next*.
8. When presented with the screen entitled "Create a New Queue: Name and Type" select *Finish* and press *Enter*.
9. Select *Exit* and press *Enter*. You will be asked to save your changes. Choose *Yes* and press *Enter*.
10. Type the command: `cd` and then type: `lp install log`
11. Type the command: `lpstat`  
(Note: you should see one print job active for root with Job number 1)
12. Type the command: `cancel 1` to remove the job.
13. Type the command: `lpstat` (Note: the job should now be removed)



... ..

... ..



# UNIT 9

## The X Window System

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2094 or +1 (919) 754 3700.

## UNIT 9: Objectives

- Describe the XOrg X11 implementation
- Describe how XOrg manages and displays data
- Configure the XOrg environment

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 236 2004 or +1 (319) 754 3700.



## UNIT 9: Agenda

- XOrg concepts and architecture
- Configuring the XOrg environment
- Manage and provide network access to the XOrg environment

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



3

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1-866-266-2084 or +1 (919) 754-3700.

# XOrg: The X11 Server

- Foundation for the Red Hat Enterprise Linux graphical user interface(GUI)
- Open Source implementation of X11
- Client / Server Architecture
  - Relies on networking
    - IP or local UNIX domain sockets
  - Designed as one server to many clients
  - Highly flexible protocol

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 286 2864 or +1 (519) 754 3700.

The X Window System (also referred to as "X" or "X11") is the foundation for the graphical user interface(GUI) on Red Hat Enterprise Linux. The X Window System is maintained by the X Consortium at <http://www.X.org> and creates the reference implementation of X under an open source license. The XOrg project(<http://xorg.freedesktop.org>) adds hardware drivers for a variety of video cards and input devices, along with several software extensions to manage the visual representation of data.

It is important to understand X's relationship to what you see on the screen. X does not define how anything should look or behave. Instead, X focuses on providing a standard way in which applications, called X clients, may display, or "write" on the screen. The X server is the program that speaks through your video hardware. Any application that wants to communicate through the display is an X client, including the "OK" button changing color when "clicked" with the left mouse button. The visual effects of a mouse cursor(an arrow, or pointing hand) selecting a link on a web page are X *client* activity that spawns an X *server* event informing the web browser to send an HTTP request to the link's target(or "anchor") You do not really see the X server, but X clients. X provides the data I/O infrastructure for X clients, like a the human nervous system, it sends messages when touched by client activity.

Originally designed as a client and server application suite, X11 uses UNIX-domain or TCP/IP networking for its operation, where one server provides many clients--both hardware(hosts and displays) and software(applications and widgets)--a *protocol* through which to pass data. Expressed in this design are two layers: a device dependent, hardware layer, and a device independent software layer. The hardware layer manages the coordination of mouse and keyboard(input) and video card and display(output). The software layer provides an API as the basis of uniform visual characteristics and rendering across varied platforms. The combination of both layers provides X client applications greater hardware and operating system independence. Also, an X client running on one system can display on any X server running on any operating system, *if* sufficient access is granted.

In a single workstation, the X clients and X server still communicate via the X protocol, but instead of using TCP/IP, they use a high-speed Unix domain socket. For each managed display, this socket is `/tmp/.X11-unix/X#` where # is 0 to the greatest number of permitted connections.

# XOrg Server Design

- System video hardware I/O Management
  - Display, video and input device coordination
  - Core server: `/usr/X11R6/bin/Xorg`
  - Enhanced by dynamically loaded modules
    - drivers: ati, nv, mouse, keyboard, etc
    - extensions: dri, glx, and extmod
- Font Rendering
  - Native server: `xf86`
  - Fontconfig/Xft libraries

Rev 08133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



5

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2984 or +1 (919) 754 3700.

XOrg consists of one core server and several dynamically loaded modules. The server core is hardware independent and is extended through the configuration and loading of hardware and X11 extension specific modules stored in `/usr/X11R6/lib/modules`, and the kernel's module directory. Like the OS kernel, through the combination of unique, hardware specific modules, the X server design remains flexible, and therefore capable, regardless of the video and input hardware found on the system. In fact, when the server is invoked, it scans and rectifies its configuration against the hardware, seeking the optimum display and input capabilities. To investigate this process, look through the `/var/log/Xorg.0.log` file.

Fonts, or the character sets used to render numbers, letters and other *glyphs* (like dollar or euro symbols) are managed either through the native X server rendering engine, or the libraries of fontconfig/Xft. The native server is actually a separate service, named `xf86` and must be available to invoke the X server. The Xft rendering engine is implemented within the XOrg core server through dynamically loaded libraries. It is planned to succeed `xf86` for its flexibility, enhanced rendering capabilities, programming API, and utilities to adjust font rendering in the KDE and GNOME object frameworks. We will discuss a few of these utilities later.

The XOrg X server has native support for several types of fonts, including TrueType and Postscript Type 1. To render all fonts accurately, be sure you have the `freetype` module defined in the "Module" section of the server configuration (discussed later).

To provide the local X server fonts not shipped with the RHEL distribution, copy them to a directory under `/usr/share/fonts`, or to the end-user directory `$HOME/.fonts`. The next time an X session starts, the Xft system spawns `fc-cache`, automatically configuring fonts under these directories for most client applications (one notable exception is the OpenOffice Suite which has its own font management system).

By default, `xf86` listens on a Unix domain socket and does not accept remote network connections. To change this behavior, comment out the bottom line of the `xf86` configuration file, `"no-listen = tcp"`. As with any network service, take appropriate security measures at your firewall, or on the font server machine itself with `iptables` or `system-config-security` to limit which clients can connect to it. Network font servers listen on TCP port 7100.

# XOrg Server Configuration

- Typically configured after installation
- Post-install configuration:
  - Best results while in runlevel 3!
  - `system-config-display`
    - options:
      - noui
      - reconfig
  - stored in `/etc/X11/xorg.conf`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (519) 754 3700.

If the X environment was not configured during installation, the utility `system-config-display` is used with the recommended `--noui` and `--reconfig` options. With these options, probed values for the monitor and video card will be used, with a deference toward the greater capabilities of system hardware. For example, if the probed hardware can support a color depth of 24 bits/pixel(TrueColor), then the server will start with this setting. If the probed attempt fails, then typical and conservative values will be set. The resulting video card and display, keyboard, and mouse configuration is stored in `/etc/X11/xorg.conf` and specifies the hardware components' resources in several sections.

Among them are:

- "Server Layout"  
defining individual combinations of "InputDevice" and "Screen." Only one per session is used.
- "Module"  
defining which hardware and extension specific modules are called when X is started.
- "InputDevice"  
defining keyboard, mouse, touchscreen, or other form of supported input device
- "Monitor"  
defining characteristics and capabilities of the physically attached output hardware.
- "Device"  
defining which hardware specific driver is used to communicate with the local video card.
- "Screen"  
defining individual combinations of "Monitor"and "Device" with display properties.

Please Note: X server configuration should be performed in runlevel 3 for best results. Also note that to run an X client to be displayed on a remote system, no local server configuration is necessary.

# XOrg Modularity

- The X server and its clients may be individually configured and combined
  - Server extensions provide enhanced rendering capabilities
    - To view server capabilities: `xdpinfo`
  - Display Managers
    - `gdm`, `kdm` and `xdm`
  - Window Managers
    - `metacity`, `kwin` and `twm`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2994 or +1 (919) 754 3700.

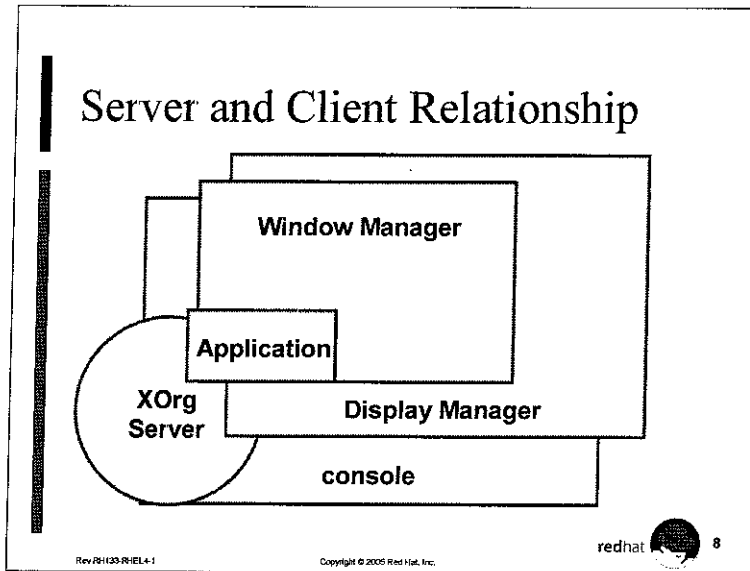
The X environment is a combination of varied hardware and software components and thus extremely flexible. You can view all current extensions, and other interesting information about the capabilities of the current X environment, by running the `xdpinfo` command in an X terminal application. Presented in its output are included the screen resolution, pixel dimensions, color depth and available rendering capabilities, like 3D graphics (GLX) or video frame playback (XVideo).

The display manager component is the X equivalent of the text-based login prompt. RHEL ships with `gdm`, `kdm`, and `xdm` display managers. All three use the PAM authentication subsystem, and their means of authentication may be independently configurable. Display managers are usually started in runlevel 5 from the `/etc/X11/prefdm` script, discussed later.

The `xdm` display manager is shipped with most X11 implementations. Its configuration information is stored under the `/etc/X11/xdm` directory. `kdm` and `gdm` are extremely configurable alternatives that support features such as session types, locale definition, and system commands such as `halt` and `reboot`. Although the `gdm` display manager inherits some of its configuration from `/etc/X11/xdm` (notably the `Xsetup_0` script), it reads its configuration from the `/etc/X11/gdm` directory and the `gdm.conf` file. `xdm` is normally, if only used when neither `gdm` nor `kdm` are available. Once authenticated, the display manager spawns its default window manager. Both display and window manager may be configured in `/etc/sysconfig/desktop`.

Window managers are a special type of X client. They encapsulate other clients, allowing them to be moved, resized, or iconified. They also provide the desktop theme, configurable menus, panel utilities, and session management. Of the many window managers available for X, RHEL ships with three: `metacity` the GNOME window manager, `kwin` for KDE, and `twm`, a basic window manager shipped with XDM and used if neither GNOME, nor KDE are installed. Window managers provide the core of the graphical user interface (GUI).

Widgets are components of a GUI application and include dialog boxes, scroll bars, menus, etc. There are many widget libraries available for applications to use: Athena, Motif, Qt, GTK+, and others. The Qt library is used by the KDE environment, and GTK+ is used by the GNOME environment.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2904 or +1 (619) 754 3700.

Here we explore the relationships between the system hardware, X server and X clients.

**A.** The system console comprises the user input/output interface between keyboard, mouse, video subsystem and graphics display and is owned by the user who successfully authenticates through this interface. This ownership is granted whether the X environment is running or not. The console hosts the X server environment, and the X server, in turn, supports its clients. X clients must connect to an X server and will not run from merely a text-based console session

**B.** If the X environment is not running (the system may be in runlevel 3), the console owner may invoke the X server, which attaches itself to the console, with the environment variable `DISPLAY=<hostname>:0.0`. To meet the most basic hardware and software configuration requirements, the server may be started alone by running `/usr/X11R6/bin/X`. If an X appears in the center of the display, and its movement is reflected with mouse movement, then the X server configuration is operational. Press `Ctrl-Alt-Backspace` to exit the server.

**C.** If the X server configuration is operational, then the program `/usr/X11R6/bin/xinit` should present a similar X in the center of the display, along with a single application(`xterm`) in the upper left corner(a location also known as the origin). This is a default behavior of `xinit` that may be customized as we discuss it later. Note that here we have only the X server and one application: no display, nor window manager is currently running. From the `xterm` prompt, additional X clients may be spawned, including a window manager(`metacity`), the `gnome-panel` application, and the Nautilus file system browser. Each of these, and all applications run from them, are a child process of the original `xterm` and a client of the X server. If the original `xterm` is terminated, then all X client applications are terminated and the X server as well.

**D.** If the system is in runlevel 5, then the X environment is likely running and a display manager(`gdm`) provides the authentication application. On successful authentication, the owner of the display is typically presented an X session with a window manager. Again, the console owner also owns the X display and, as in C above, all applications run from the window manager(parent process) will exit when the user exits the window manager. In runlevel 3(without an X display manager running), a similar session managed by the window manager is achieved by running `usr/X11R6/bin/startx`

## XOrg in runlevel 3

- Two methods to establish the environment
  - `/usr/X11R6/bin/xinit`
  - `/usr/X11R6/bin/startx`
- Environment configuration
  - `/etc/X11/xinit/xinitrc` and `~/.xinitrc`
  - `/etc/X11/xinit/Xclients` and `~/.Xclients`
  - `/etc/sysconfig/desktop`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



9

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2994 or +1 (618) 754 3700.

To start an X session in run level 3, use `xinit` or `startx` from a virtual console shell prompt.

`init` (or `startx`) will pass control of the X session to `/etc/X11/xinit/xinitrc`, unless `~/.xinitrc`, exists

`xinitrc` seeks to read additional system and user configuration files, including:

Resource files:

`/etc/X11/Xresources` and `$HOME/.Xresources`

Input device configuration files:

`/etc/X11/Xkbmap` and `$HOME/.Xkbmap`,  
`/etc/X11/Xmodmap` and `$HOME/.Xmodmap`

`xinitrc` then runs all shell scripts in `/etc/X11/xinit/xinitrc.d`

`xinitrc` then turns over control of the X session to `~/.Xclients`, if it exists. If not, it turns over control to `/etc/X11/xinit/Xclients`

`Xclients` reads `/etc/sysconfig/desktop` to determine whether Gnome or KDE is the preferred desktop environment. If unset, or the defined is not installed, `Xclients` will attempt to run a number of other window managers in the following order:

Gnome

KDE

`twm` (plus `xclock`, `xterm`, and `mozilla`, if installed)

In the unlikely event that `Xclients` does not exist, `xinitrc` will go into failsafe mode, starting an `xclock`, `xterm`, perhaps `mozilla`, and finally the `twm` window manager.

## XOrg in runlevel 5

- Environment established by `/sbin/init`
- Environment configuration
  - `/etc/inittab`
  - `/etc/X11/prefdm`
  - `/etc/sysconfig/desktop`
    - DESKTOP defines the window manager
    - DISPLAYMANAGER defines the display manager
  - `/etc/X11/xdm/Xsession`
    - `/etc/X11/xinit/xinitrc.d/*`
      - `~/Xsession` or `~/Xclients`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyright. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2894 or +1 (910) 764 3700.

If the default runlevel is set to 5 in `/etc/inittab`, `/sbin/init` will run `/etc/X11/prefdm`. This script invokes the X server and a display manager, set in the file `/etc/sysconfig/desktop`. If no system default is set, it tries `gdm`, then `kdm`, then `xdm`.

By default, the three display managers all use the same startup scripts as part of their operation. When a display manager is first started, the `/etc/X11/xdm/Xsetup_0` file is run (as root) before the display manager presents a login widget. Then, once the user authenticates, `/etc/X11/xdm/Xsession` is run. This does many of the same things `/etc/X11/xinit/xinitrc` does in `startx`, including running the executables in `/etc/X11/xinit/xinitrc.d`.

`/etc/X11/xdm/Xsession` determines which desktop environment to run. If the user specified one at login through the display manager, it is run. Otherwise, the script checks for either `~/Xsession` or `~/Xclients`. If all else fails, `/etc/X11/xinit/Xclients` is run as it is for `startx`.

When the user logs out at the end of their session, the X server is restarted by the display manager with a new login window.



# Configuration Utilities

- Server
  - `system-config-display`, `mouseconfig`
- Fonts and Typefaces
  - `xfs`, `chkfontpath`, `fc-cache`
- Display and Window Managers
  - `switchdesk`, `/etc/sysconfig/desktop`, `gconftool-2`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2634 or +1 (919) 754 3700.

As described earlier, XOrg server configuration is typically set during RHEL installation. Similarly, many of its client applications, including window managers, office suites and terminals are pre-configured before their initial execution; or upon first use, derive a site or user specific configuration from the environment. Applications are independent X clients and typically make a Preferences menu item available.

As each end-user(including root) maintains a user-specific set of window manager and X client application configurations, there are opportunities where site-wide configuration is preferred. Note that most methods to prepare a site-wide configuration may be superseded by the end-user invoked runtime options or client configuration file. For example, an end-user may run the X server with a preferred color depth, or pixel dimension, but only while in runlevel 3. In runlevel 5, these options are only available to the superuser, root. While a non-root user may alter most X client attributes, preparation of site-wide X server and client configurations should be considered a default, or fundamental set of options. Recall that the first end-user to authenticate at the system console owns that console and is provided exclusive access to its environment, including the local or 0 (zero) X display.

The superuser may add font directories to `xfs`, using `chkfontpath`, or add the directory path to the existing list in the `xfs` configuration file `/etc/X11/fs/config`. Non-root users may also add fonts to a directory under `$HOME/.fonts`. After copying fonts to this directory, the non-root user may run `fc-cache`, which updates the Xft font configurations in `$HOME/.fonts.cache-1`. This occurs automatically at the next execution of an X server session. Additional font and desktop configuration is available through the desktop panel Preferences menu, or from the command-line in an X client, like `gnome-font-properties`, used to adjust anti-aliasing and other font rendering characteristics.

Although the authenticating display manager will typically spawn its own window manager, an individual user may define his own default desktop environment using `switchdesk`. Both the display and window manager application choices may be set in `/etc/sysconfig/desktop` as a system default. System default window manager options for the Gnome Desktop, may be set using `gconftool-2`. The Gnome GConf utility is extremely extensible, including configuration stored in XML which may be further extended to LDAP or SQL management. For more information on this powerful utility, visit <http://www.gnome.org/learn/admin-guide/2.0/>

# Remote X Sessions

- X protocol communication is unencrypted
- Host-based sessions implemented through the `xhost` command
- User-based sessions implemented through the Xauthority mechanism
- `sshd` may automatically install `xauth` keys on remote machine
  - Tunnels X protocol over secure encrypted `ssh` connection

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

The X11 environment was originally designed to run on one host and serve remote client connections. While this capability is still available, because data is transmitted unencrypted it is disabled by default. All three display managers may be configured to host remote X sessions and act as a system authentication agent. For example, a user on one host `client` may request a display managed session from another host `server` which in turn provides access to `server`'s system resources. The syntax for making the request is `X -query server`, if an X session is not already running on `client`, and `X -query server :1` if one instance of X is already running.

To locally display an X client application from a remote host, you must configure your local X server and remote hosts. This technique is useful when you wish to run a GUI client to manage resources or services on a remote system with no console attached (a "headless server," as found in most data-centers). A few methods are available, including `xhost`, `xauth`, or forwarded through an SSH connection. In all cases, display access is permitted by the local console and X server session owner.

As its name suggests, `xhost` permits connections on a host to host basis. When access to your display is permitted by this method, *all* redirected X client connections are accepted from the permitted host, regardless of the user who spawns them! If you require host-based X access, issue the command `xhost + trustedhost`, which instructs the local X server that clients from `trustedhost` may have access to the display. Similarly, `xhost - friendlyhost` removes this access permission. The command `xhost +` without a hostname argument disables access restrictions, and is not recommended.

For greater accountability, a user-based mechanism may be preferable. The `$HOME/.Xauthority` file, used in conjunction with the `xauth` command, provides the capability of authenticating remote users who are permitted to use the local display. While certainly more secure than `xhost`, data is similarly passed unencrypted between systems. An even more secure method utilizes "Secure Shell" (SSH).

SSH provides an excellent user-based access mechanism that is encrypted, easy to use and compatible with the `xauth` facility. Using the `xauth` program for display access, SSH will tunnel X authentication information between systems. To achieve this, use `ssh -Y remote-host`.

## End of Unit 9

- Questions and answers
- Summary
  - What function does a display manager serve?
  - What tools are available for configuring XOrg?
  - What purpose does `xfs` serve?

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (619) 754 3700.



.....

.....

.....

# Lab 9

## The X Window System

---

### Sequence 1: Explore the X Start-up Sequence

- a) Create and edit `/etc/X11/xinit/xinitrc.d/xeterm.sh` place the following lines in it, and make it executable:

```
#!/bin/bash
xterm &
```

- b) Switch to runlevel 5 if you are in another runlevel (init 5).
- c) Log into your system through your display manager, i.e., `gdm`, `kdm`, or `xdm`. What happens?
- d) Switch to runlevel 3 (init 3), then run the command

```
startx
```

What happens?

- e) Create a user named `xuser`. Create a file in `xuser`'s home directory called `.xsession`:

```
#!/bin/bash
gnome-terminal &
firefox &
exec metacity
```

Make the `.xsession` file executable.

- f) Switch to runlevel 5 (init 5) and log in through your display manager as that user; what happens?
- g) Switch to runlevel 3 (init 3) and log in as that user in a virtual console then use the `startx` command; what happens?

### Other Results or Questions:

1. Outline all the steps you would take to upgrade a video card on a Linux machine, including how you would select the video card to use.
2. Describe what steps you would take to fix a system that was coming up in runlevel 5, but for which the card was misconfigured in X (effectively resulting in a locked system on boot.)



# UNIT 10

## Advanced Filesystem Management

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 285 2994 or +1 (919) 754-3700.

## UNIT 10: Objectives

- Upon completion of this unit you should be able to:
  - Configure software RAID
  - Manage software RAID
  - Configure Logical Volumes
  - Manage Logical Volumes
  - Implement Quotas

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



2

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 255 2654 or +1 (610) 754-3700.




## UNIT 10: Agenda

- RAID concepts and software RAID configuration
- Configuration of Logical Volumes
- Filesystem Quotas

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2994 or +1 (910) 754-3700.

## Software RAID Configuration

- Create and define RAID device using `mdadm`  
`mdadm -C /dev/md0 -l 0 -n 2 /dev/hda5 /dev/hdb7`
- Format each RAID device with a filesystem  
`mke2fs -j /dev/md0`
- Test the RAID devices
- `mdadm` allows you to check the status of your RAID devices  
`mdadm --detail /dev/md0`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 253 2534 or +1 (619) 754 3700.

In order to install and configure software RAID after installation, the *raidtools* RPM is required. It should already be installed by default, since it is included in the minimal Core installer component. Then:

- Use `fdisk` or another tool to create disk partitions of type 0xfd, "Linux RAID"
- If needed, run `partprobe` to have the kernel reload the partition table.
- Initialize and activate your RAID array  
`mdadm -C /dev/md0 --chunk=64 --level=5 --raid-devices=3 /dev/sd{b,c,d}1`
- Create a filesystem on the new software RAID device. With `mke2fs`, there is a special `-R stride=n` option that can improve performance. The stride is the software RAID device's chunk-size in filesystem blocks. For example, with an ext3 filesystem that will have a 4K block size on a RAID device with a chunk-size of 64K, the stride should be set to 16.  
`mke2fs -j -b 4096 -R stride=16 /dev/md0`
- Add the software RAID device and its mount point to `/etc/fstab` as appropriate


man mkraid

# Software RAID Recovery

- Simulating disk failures  
`mdadm /dev/md0 -f /dev/sda1`
- Recovering from a software RAID disk failure
  - replace the failed hard drive and power on
  - reconstruct partitions on the replacement drive  
`mdadm /dev/md0 -a /dev/sda1`
- `mdadm`, `/proc/mdstat`, and `syslog` messages

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 208 2994 or +1 (919) 754 3700.

The `mdadm -f` command can be used to simulate a drive failure. This is useful when testing RAID 1 or RAID 5 arrays, but will destroy a RAID 0 array. Spare disks can be set aside in a RAID 1 or RAID 5 array, that can automatically start rebuilding as a replacement disk (`-x` option). The array can be used while it is rebuilding, although performance will be degraded.

Recovery is fairly simple. Power off the system, replace the failed disk, power on the system, and reconstruct an appropriate partition table on the disk. The new disk can be inserted into the array with the `mdadm -a` command. You may have to remove the failed disk from the array with `mdadm -r`.

Information about disk failures is logged through `syslog` to `/var/log/messages` by default. Information on the current state of software RAID devices is also available in `/proc/mdstat`. The output below shows a RAID 5 device, `/dev/md0`, made up of `/dev/sdb1`, `/dev/sdc1`, and `/dev/sdd1`; `/dev/sdd1` has failed.

```
[root@station ~]# cat /proc/mdstat
Personalities : [raid5]
md0 : active raid5 sdd1[3](F) sdc1[1] sdb1[0]
      272896 blocks level 5, 64k chunk, algorithm 2 [3/2] [UU_]
unused devices: <none>
```

Information on estimated time to reconstruction will appear in this file if the array is rebuilding. You can also use the `mdadm --detail` command to view the status of the RAID array.

## Converting LVM1 to LVM2

- RHEL4 Uses the LVM2 format for metadata
  - more compact
  - supports transactional changes & replication
  - human readable and editable in an emergency
- Existing LVM1 volumes can be converted to LVM2 with the `vgconvert` command.
  - `vgconvert -M2 vg0`
  - converts the volume group `vg0` from LVM1 to LVM2

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 266 2984 or +1 (919) 754 3700.

Red Hat Enterprise Linux 4 and later releases use LVM2, removing many of the limitations present in LVM1. Before you can take advantage of these enhancements, you will need to convert your LVM1 volume groups to LVM2. An example:

```
vgconvert -M2 vg1
```

would convert the volume group `vg1` from LVM1 to LVM2.

You can convert the root filesystem from rescue mode


The volume group information is now stored in human readable format under `/etc/lvm`

## Creating Logical Volumes

- Create physical volumes  
`pvcreate /dev/hda3`
- Assign physical volumes to volume groups  
`vgcreate vg0 /dev/hda3`
- Create logical volumes from volume groups  
`lvcreate -L 256M -n data vg0`  
`mke2fs -j /dev/vg0/data`

Rev 781133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (810) 754 3700.

Any disk partitions to be used as physical volumes need to have their partition types set to 0x8e, Linux LVM. Then the partitions need to be added as physical volumes with `pvcreate device-name`. Software RAID devices may also be set up as physical volumes.

New volume groups need to be created and one or more physical volumes assigned to them with the `vgcreate` command. The new volume group can have almost any name. At this time the size of an extent is determined; by default it is 4 MB. This affects the minimum size of changes which can be made to a logical volume in the volume group, and the maximum size of logical and physical volumes in the volume group. A logical volume can contain at most 65534 extents, so the default extent size limits the volume to about 256 GB; a size of 1 TB would require extents of at least 16 MB:

*extend size*  
`vgcreate -s 16M vg0 (/dev/hda3 /dev/sda2)`

Logical volumes are created from these volume groups, and may have arbitrary names. The size of the new volume may be requested in either extents or in KB, MB, GB, or TB (rounding up to whole extents):

```
lvcreate -L 512M -n data vg0 (asks for /dev/vg0/data of 512 MB)
lvcreate -l 32 -n test vg0 (asks for /dev/vg0/test of 32 extents)
```

Logical volumes may also be striped like a RAID 0 device between multiple physical volumes. A striped logical volume may be extended later, but only with extents from the original physical volumes. Because LVM can not tell if two particular physical volumes in the volume group are on the same drive, it is best not to do this if a volume group will contain striped logical volumes. The following command creates a logical volume, `/dev/vg1/striped`, striped between two physical volumes in the `vg1` volume group, using the default stripe size.

```
lvcreate -i 2 -L 1G -n striped vg1
```

# Resizing Logical Volumes

- `lvextend` and `ext2online` can extend mounted ext2/3 filesystems
  - `lvextend` first grows the logical volume
  - you can not shrink mounted filesystems
- Physical volumes may be added to or removed

```
vgextend vg0 /dev/sdb1
pvmove /dev/hda3
vgreduce vg0 /dev/hda3
```

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.

redhat

8

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2094 or +1 (819) 764 3700.

Logical volumes may be resized dynamically while preserving the data on the volume, if the volume's filesystem supports resizing. The `lvextend` command allows resizing of an ext2 or ext3 based logical volume. `ext2online` can be used to grow mounted ext2/3 filesystems. `lvextend` must be called first to grow the logical volume.

The following commands will grow the mounted `/dev/vg0/data` filesystem.

```
lvextend -L +500M /dev/vg0/data show the size
ext2online /dev/vg0/data
```

For other filesystems, the `lvextend` utility can be used to add unallocated extents in the volume group to a logical volume. Then native utilities for the filesystem can be used to expand it to fill the volume. To reduce a filesystem, first the native utilities should be used to shrink the filesystem, then `lvreduce` should be used to shrink the logical volume.

Additional physical volumes may be added to a volume group to provide more unallocated extents to assign to logical volumes. The physical volumes need to be setup with `pvcreate`, then are added to the volume group with the `vgextend` command.

Physical volumes can also be removed from a volume group. This is useful to remove an old disk from the volume group. The `pvmove` command can redistribute extents from the physical volume being decommissioned to the other physical volumes in the volume group. In its simplest mode, `pvmove` takes the name of the physical volume to be removed as its argument: `pvmove /dev/hda3`.

Then once there are no extents in use on the old physical volume, it can be removed from the volume group with the `vgreduce` command: `vgreduce vg0 /dev/hda3`.

Some commands are available to help you gather information about the state of your physical volumes, volume groups, and logical volumes. Three useful commands are `pvdisplay`, `vgdisplay`, and `lvdisplay`.

# The Linux Quota System

- Overview
  - Implemented within kernel
  - Enabled on a per-filesystem basis
  - Individual policies for groups or users *only one*
    - Limit by number of blocks or inodes
    - Implement both soft and hard limits
- Initialization
  - Partition mount options: `usrquota`, `grpquota`
  - Initialize database: `quotacheck`

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2864 or +1 (619) 754 3700.

The Linux quota system allows an administrator to establish limits on the amount of disk resources users can consume. Because resource accounting must occur with every file creation, quotas are implemented within the kernel. Various required quota administration utilities are found within the *quota* RPM.

## Quota Initialization

During this discussion, examples will be given for implementing user quotas on the `/home` partition. Group quotas are implemented in a nearly identical manner.

**Define partition options:** In order for a partition to implement quotas, it must be mounted with the `usrquota` or `grpquota` options. These options can be added to the appropriate entries in `/etc/fstab`. After editing, the options can be made to immediately take effect by remounting the filesystem

edit `/etc/fstab`, adding the `usrquota` option to the `/home` partition

`mount -o remount /home`

**Create/Update the database:** The disk usage database is stored in specially named binary files within a partition's top-level directory, `aquota.user` and `aquota.group`. These files may have to be created manually (`touch /home/aquota.user`). During initialization, or any time the database file is out of sync with the actual state of a partition (for example, if quotas were turned off for a period of time, or if a system was brought down rudely without unmounting partitions), the database can be brought up to date by running the `quotacheck` command:

`quotacheck -c /home`

## The Linux Quota System (cont.)

- Implementation
  - Start or stop quotas: `quotaon`, `quotaoff`
  - Edit quotas directly: `edquota username`
  - From a shell:  
`setquota username 4096 5120 40 50 /foo`
  - Define prototypical users:  
`edquota -p user1 user2`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2894 or +1 (919) 754 3700.

**Starting and stopping quotas:** Quotas are turned on and off with the `quotaon` and `quotaoff` commands. These commands either take partitions as their arguments, or are invoked with the `-a` command line switch, in which case quotas are turned on for all appropriate partitions defined in `/etc/fstab`. These commands rarely need to be run in practice, because they are included within the default RHEL initialization script

```
/etc/rc.d/rc.sysinit.
```

```
quotaon /home (or quotaon -a)
```

**Editing user policies:** User policies are implemented with the `edquota` command. This command invokes an editor and loads a template, which can then be edited to establish the appropriate values. These values are committed to the database upon exiting the editor. In order to ease the propagation of quotas, user policies can be prototyped from established policies for another user. Often, when first establishing quotas, it is helpful to first define a “prototypical” user, and then edit the user’s properties. (If you do not first prototype the user, no template will be provided by the `edquota` command.) Grace periods are established by using the `-t` command line switch with `edquota`.

```
edquota elvis (implement a policy for this user)
edquota -p elvis prince (mimic user prince’s policy from elvis’s)
edquota -t (establish a grace period)
```



## The Linux Quota System (cont.)

- Reporting
  - User inspection: `quota`
  - Quota overviews: `repquota`
  - Miscellaneous utilities: `warnquota`

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 298 2984 or +1 (910) 754 3702.

**Generating quota reports:** Users can inspect their disk usage and quotas by issuing the `quota` command. An administrator can generate a report of disk usages by all users with the `repquota` command. Users over their quotas can be warned with a `warnquota` cron job.

## End of Unit 10

- Questions and Answers
- Lab

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 268 2994 or +1 (919) 754 3700.

### Important files covered in this Unit:

`/proc/mdstat`

### Important commands covered in this Unit:

`mdadm`  
`ext2online`  
`vgconvert`  
`pvcreeate, lvcreate, vgcreate, vgdisplay, lvextend, lvremove`  
`quotacheck`  
`quotaon`  
`quotaoff`  
`edquota`  
`setquota`  
`repquota`  
`quota`

# Lab 10

## Advanced Filesystem Management

---

**Goal:** Develop skills with installation of LVM, RAID and quotas.

### Sequence 1: Creating Logical Volumes with LVM

1. Use `fdisk` to create four new partitions of type Linux LVM (`0x8e`) in unpartitioned space on your hard disk. These partitions should have the same size; to speed up the lab, do not make them larger than 1 GB in size. Make sure to write the changes to disk when you quit by using the `w` command to exit `fdisk`. Reboot the system.

2. Once the system has booted up, log in as root.

3. Initialize your Linux LVM partitions as physical volumes with `pvcreate`. Assuming the LVM partitions are `/dev/hda9`, `/dev/hda10`, `/dev/hda11`, and `/dev/hda12`, the command would be:

```
pvcreate /dev/hda9 /dev/hda10 /dev/hda11 /dev/hda12
```

You can use the `pvdisplay` command (for example, `pvdisplay /dev/hda9`) to verify that the partitions have been registered as physical volumes.

4. Next, create a volume group called `test0`, using the default 4MB extent size, which initially contains only one of your physical volumes:

```
vgcreate test0 /dev/hda9
```

You can use the `vgdisplay` command to list information on all the volume groups active on the system.

5. Create a small logical volume that does not use up all the space in the volume group. Look for `VG Size` and `Free PE/Size` in the output of the `vgdisplay` command to assist you with this. For example, to create a logical volume about 40 MB in size:

```
lvcreate -L 40M -n data test0
```

6. The command `lvdisplay /dev/test0/data` can be used to verify that this command worked.

7. Now create an `ext3` file system on your new logical volume: `mke2fs -j /dev/test0/data`

8. Make a new directory `/data` and then `mount /dev/test0/data /data`

Copy some files into `/data`. Try creating a large file:

```
dd if=/dev/zero of=/data/bigfile bs=1024 count=20000
```

Run `df` and check disk usage and space free on `/data`. Verify that everything works like a normal `ext3` file system. Set up `/etc/fstab` to automatically mount `/data` at boot, and reboot the system to verify that the logical volume is available after a reboot.

**Sequence 2: Working with Logical Volumes**

1. First, enlarge the logical volume `/dev/test0/data`. Then, use `ext2online` to enlarge the file system. For example, to enlarge it by about 50 MB:

```
lvextend -L +50M /dev/test0/data
ext2online /dev/test0/data
```

2. Verify that your files are still intact. Run `df` and check to verify that more free disk space is now available on `/data`.
3. Use the remaining extents in `test0` to create a second logical volume. Run the `vgdisplay` command, and look at the line `Free PE / Size`. If it says something like `166 / 664 MB`, this means the volume group has 166 extents (or 664 MB of space) free. So to make a second logical volume, `/dev/test0/scratch`, which is exactly 166 extents in size, the command would be:

```
lvcreate -l 166 -n scratch test0
```

4. Rerun `vgdisplay`. There should be no free extents left.
5. Format the new logical volume: `mke2fs -j /dev/test0/scratch`
6. Add one of the unused physical volumes to the volume group:

```
vgextend test0 /dev/hda10
```

7. If you run `vgdisplay` again, there should be free extents (provided by the new physical volume). Extend the new logical volume by 20 MB:

```
lvextend -L +20M /dev/test0/scratch
mkdir /scratch
mount /dev/test0/scratch /scratch
ext2online /dev/test0/scratch
```

Use `lvdisplay` and `vgdisplay` to verify this command worked.

Before moving on to the RAID sequence, disassemble your LVM-managed volumes:

Remove any `/etc/fstab` entries you might have set up

```
umount /dev/test0/data
lvremove /dev/test0/data
umount /dev/test0/scratch
lvremove /dev/test0/scratch
vgchange -an test0
vgremove test0
```

(this deactivates the volume group)  
(this deletes the volume group)

**Sequence 3: Software RAID**

1. You probably do not have multiple disks in your computer in the classroom. In this sequence we are going to construct a practice software RAID device using multiple partitions which are all on the same disk. The system will let us do this (with some warnings), but in real life the RAID device should consist of individual partitions which are each on a separate disk. Run `fdisk` on your hard disk again, converting your Linux LVM (type 0x8e) partitions to type Linux raid auto (type 0xfd). Save your changes and reboot. Remember that the example is assuming your partitions are `/dev/hda9`, `/dev/hda10`, `/dev/hda11`, and `/dev/hda12`. Your actual partitions may vary.

2. Initialize your RAID array:

```
mdadm --create /dev/md0 --level=5 --raid-devices=4 /dev/hda{9,10,11,12}
```

This will also start the array. The array will immediately begin the construction process, but it can be used immediately. In a spare virtual console or terminal window, run `watch cat /proc/mdstat` to keep an eye on the progress of array construction.

3. While waiting for this process to complete, we can start using the RAID device. Format the disk with an ext3 file system. We will use 4K blocks. The stride option should be set to the chunk-size in file system blocks; this allows `mke2fs` to lay out the file system more efficiently for this RAID device.

```
mke2fs -j -b 4096 -R stride=16 /dev/md0
```

4. See if you can mount `/dev/md0 /data`. This command should work, even if the array is still building. Use `df` to check the size of the file system. If all four of your RAID partitions were the same size, then the file system should be the size of three of them added together. (The missing space is used by the redundant parity information striped across the array.)
5. The `mdadm --detail` command can be used to display useful information on the state of your RAID devices. Try out this command, `mdadm --detail /dev/md0`.
6. Try creating some files in `/data`. Add an entry to `/etc/fstab` and test mounting and unmounting the file system.
7. Check `/proc/mdstat` and make sure that the array has finished building. You should see output like:

```
Personalities : [raid5]
md0 : active raid5 hda12[3] hda11[2] hda10[1] hda9[0]
      770688 blocks level 5, 64k chunk, algorithm 2 [4/4] [UUUU]
```

It is now time to break things. Simulate a single disk failure by running the following command:

```
mdadm --manage /dev/md0 --fail /dev/hda11
```

Look for errors in `/var/log/messages`, and note that the output of `/proc/mdstat` has changed:

```
md0 : active raid5 hda12[3] hda11[2] (F) hda10[1] hda9[0]
      2328064 blocks level 5, 64k chunk, algorithm 2 [4/3] [UU_U]
```

You will now need to replace the hard drive that failed and add in a new drive. First, remove the existing drive from the RAID set:

```
mdadm --manage /dev/md0 --remove /dev/hda11
```

Then add in the new good drive:

```
mdadm --manage /dev/md0 --add /dev/hda11
```

8. If everything worked properly, `/proc/mdstat` should show the array rebuilding again.

**Sequence 4: Implementing Quotas**

Set up your system to use quotas to restrict disk usage in /home on a user-by-user basis. Try to meet the following requirements on your own. A solution is provided on the last page of the lab if you are not sure how to proceed.

- 1 Create a new user named `filehog`. For /home, set `filehog`'s quota to have a soft limit of 60 inodes (files) and a hard limit of 100 inodes.
- 2 To test these restrictions, run the following commands:

```
su - filehog
quota

for i in $(seq 1 100); do
echo -n "file${i} "; touch file${i} 2>&1; done | less

quota
```

3. The `quota` commands should report the restrictions in place and current usage for `filehog`. The for-loop is intended to create 100 files. Since `filehog` already owns some files in /home/`filehog`, this loop should not be able to create all 100 files.

What you should see is a series of filenames echoed until the soft limit is reached. When it is reached, a warning will be printed out but the `touch` commands should continue to succeed. Once the hard inode limit is reached, error messages should begin to appear and the `touch` commands will fail to create any more files. Run `ls` in /home/`filehog` to confirm that quotas worked.

4. Create a new user named `diskhog`. For /home, set that user's quota to have a soft limit of 4 MB of disk space and a hard limit of 5 MB of disk space.

To test these restrictions, run the following commands:

```
su - diskhog
quota
dd if=/dev/zero of=bigfile bs=1M count=3
quota
```

(`dd` should work fine.)

```
dd if=/dev/zero of=bigfile bs=1M count=4
quota
```

(Should get a warning.)

```
dd if=/dev/zero of=bigfile bs=1M count=5
quota
```

(Should fail to write the whole file.)

**Solution to Sequence 4 (Implementing Quotas):**

1. Edit `/etc/fstab`. On the line for the `/home` file system, replace the `defaults` mount option with `usrquota`.
2. Remount `/home`: `mount -o remount /home`
3. Create the initial `/home/aquota.user` database: `quotacheck -cM /home`
4. Turn on quotas: `quotaon /home`
5. Run `edquota filehog`. This will open up a text file in your default editor containing `filehog`'s quota settings. (The default is `vi`, but you can override this by setting the `EDITOR` environment variable.) Set the inode soft limit to 60 and the inode hard limit to 100
6. Run `edquota diskhog`. Set the soft block limit to 4096 and the hard block limit to 5120.

If you do not like the `edquota` program, there is an alternative program that we did not discuss in the unit, `setquota`, which is better suited for use in shell scripts. The `setquota` equivalents of steps 5 and 6 are:

5. `setquota -u filehog 0 0 60 100 /home`
6. `setquota -u diskhog 4096 5120 0 0 /home`



# UNIT 11

## Troubleshooting

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (916) 754 3700.

## Unit 11: Objectives

- Upon completion of this unit you should be able to:
  - Describe troubleshooting strategies
  - Identify the things to check when diagnosing X, networking, and boot problems
  - Fix filesystem corruption
  - Boot a system into a recovery run level
  - Perform system recovery in a rescue environment

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2854 or +1 (610) 754 3700.

# Unit 11: Agenda

- Troubleshooting strategies
- Things to check
- Boot procedures
- Rescue environment

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.




For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 296 2994 or +1 (919) 754 3700.

# Troubleshooting

- Treat the problem as a symptom
- Gather data by identifying other problems
- Identify what still works
- Form a hypothesis about what is wrong
- Check log files for supporting evidence
- Backup config files before editing them

Rev RHEL4-1

Copyright © 2005 Red Hat, Inc.



4

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

The process of troubleshooting any system, including those running Red Hat Enterprise Linux (RHEL), is both science and art. The science comes from the concepts of hypothesis testing, experimentation, comparison, and reproducing results. The art of troubleshooting comes from the realization that operating systems, services, and applications do not always work as we hope or anticipate, or even as their creators hope or anticipate. Science allows us to focus on likely causes, while art permits us consider the off-the-wall and unlikely as possibilities.

Regardless of whether the problem is something wrong in a single user's environment or a system-wide crisis that has rendered a system unusable, sensible troubleshooting begins with the easy fixes. This may mean running a configuration tool, or it may mean looking at the system or service-specific log. Logs often provide explicit information on problems, sometimes even identifying the exact line of a configuration file that is causing problems, so it is often far easier to look in a log for an answer than to parse a configuration file for what might be a trivial syntax error. Many services provide switches that enable higher levels of debugging output, and some may be run as foreground applications for debugging purposes. Some services, such as Samba, provide syntax checkers that may also prove useful. Begin looking in configuration files after you have exhausted these possibilities.

Sometimes the problem you are having is a bug, not misconfiguration. If you have followed instructions and documentation to the letter, only to find that these do not work as you assume they should, then perhaps you should visit Red Hat's bug tracking facility, Bugzilla. You can query the database of known and reported bugs by RHEL version, package, and numerous other criteria. You might find that the bug has been fixed in an update, or perhaps there is a workaround that will tide you over until an update is available.

You may sometimes find a package's source RPM worth investigation. Before delving into the source, check the binary RPM's change log (`rpm -q --changelog package`) to see if any changes have been made or patches applied that would explain why are not working as you would expect. Sometimes investigating the patches applied to the original sources will reveal that certain capabilities have been turned off for a good reason. The spec file may also be helpful.

## Things to Check: X

- Never debug X while in runlevel 5!
- Try `system-config-display` first
- `X -probeonly` - *Log verbose 5*
- Is `/home` or `/tmp` full, or has the user reached a hard quota?
- Is `xfs` running?

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 299 2984 or +1 (910) 754 3700.

X can sometimes be problematic. The first thing to try when confronting X problems is `system-config-display`. If the X problem is occurring in runlevel 5, it could make logging in impossible in X or in virtual consoles. Reboot or change the system runlevel to runlevel 3.

Sometimes viewing the output of the X command itself is revealing. The `-probeonly` switch will perform all tasks necessary to start the X server without actually starting it. Thus, the start up messages are displayed. If the text scrolls off the screen, the material off screen can be viewed by holding down the Shift key and pressing the Page Up key (Shift/Page Down to scroll down again). The output may be viewed by holding down Shift and using Page Up and Page Down. `/usr/share/hwdata/Cards` may provide useful information about optional configuration features for your hardware.

A user's inability to create files in the user's home directory or in `/tmp` either because of a full filesystem or because of a hard quota limit typically inhibits the ability of the user to run X. Although the exact symptoms differ between runlevels, messages will appear stating (in runlevel 3) or suggesting (in runlevel 5) that a filesystem is full.

The `xfs` font server is the only font path entry in the `/etc/X11/xorg.conf` file, so if it is not running, X will not run. Make sure that `xfs` is configured to run in the appropriate runlevels. Occasionally it may be necessary to delete stale lock, pid, or socket files. Once in a while, the font indexes in a font directory may be corrupt, and it will be necessary to run `mkfontdir` to recreate them. When this happens, `xfs` may seem to start correctly, but then dies. Try commenting out font paths in `/etc/X11/fs/config`, then run `xfs` from a terminal to determine which directory has problems.

Remember that X is a network service, even when you have not enabled access to your display. Unsuccessful hostname resolution can produce various behaviors, including the inability to launch applications from the panel. Changing hostnames can cause similar problems; if you need to change your computer's hostname, switch out of runlevel 5 and make sure X is not running (through `startx`), change the hostname, and only then restart X.

## Things to Check: Networking

- Hostname resolution  
`dig www.redhat.com`
- IP configuration  
`ifconfig`
- Default gateway  
`route -n`
- Module specification
- Device activation

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



6

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 285 2004 or +1 (919) 754 3700.

Hostname resolution problems can create problems for clients and servers alike. Aside from requiring successful forward lookups, reverse lookups are essential for many host-based security mechanisms. Tools like `host` and `dig` are invaluable for determining whether hostname resolution problems exist.

IP configuration may be checked using the `ifconfig` command, which will print information such as an interface's IP, the subnet mask, and other important settings. The `netstat -r` and `netstat -rn` commands will show if a system's routing table is correct. The absence of a default gateway or the existence of multiple default gateways can create problems. Inability to contact the default gateway (and thus, to reach the gateway to get outside the local network) can also cause networking problems.

It is possible that the kernel module for your particular network interface card has been mis-specified. For example, the Red Hat Enterprise Linux installer sometimes probes a `de4x5`-based card as a tulip-based card. Unfortunately, the tulip module will only work enough to enable the interface, but not enough to work.

Don't overlook the obvious maybe the interface has not been activated, or was deactivated for some reason.

# Order of the Boot Process

- Bootloader configuration
- Kernel
- /sbin/init
  - Starting init
- /etc/rc.d/rc.sysinit
- /etc/rc.d/rc,/etc/rc.d/rc?.d/
  - Entering runlevel X
- /etc/rc.d/rc.local
- X

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 286 2994 or +1 (819) 754-0700.

In order to troubleshoot boot time problems, one must understand the boot process itself, remember how things look when they are working correctly, and narrow down how far into the process a failure is occurring

## No bootloader splash screen or prompt appears

- GRUB is misconfigured
- Boot sector is corrupt
- A BIOS setting, such as disk addressing scheme, has been modified since the boot sector was written

## Kernel does not load at all, or loads partially before a panic occurs

- Corrupt kernel image
- Incorrect parameters passed to the kernel by the bootloader

## Kernel loads completely, but panics or fails when it tries to mount root filesystem and run /sbin/init

- Bootloader is misconfigured
- /sbin/init is corrupted or /etc/inittab is misconfigured
- Root filesystem is damaged and unmountable

## Kernel loads completely, and /etc/rc.d/rc.sysinit is started and interrupted

- /bin/bash is missing or corrupted
- /etc/fstab may have an error – evident when filesystems are mounted or fsck'ed
- Errors in software RAID or quota specifications
- Corrupted non-root filesystems (due to a failed fsck)

## Run level errors (typically services)

- Another service required by a failing service was not configured for a given runlevel
- Service-specific errors
- Misconfigured X or related services in run level 5

# Filesystem Corruption

- Common after crash or improper shutdown
- ext2 mounted for writing marked "dirty"
  - If not mounted or mounted read-only, "clean"
  - if not mounted and "dirty", may be corrupted
  - repair requires exhaustive check
- ext3 usually marked "clean"
  - journal indicates if recovery is needed
  - only need to check files recorded in journal

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



8

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 295 2994 or +1 (919) 754 3700.

Filesystem corruption is one of the more common boot-time problems. It can occur after a system crash causes the machine to shut down without correctly unmounting its filesystems. Filesystem corruption happens because the operating system uses RAM as disk buffers to improve performance; when power fails, information written to memory buffers which is not yet synchronized to the disk is lost.

When a device using an traditional filesystem like ext2 is mounted read-write, the filesystem is marked on disk as "dirty". Only when the `umount` command is issued and all data is safely written to the disk is the disk marked "clean" and unmounted. Therefore, a filesystem that is not mounted but is marked as "dirty" hasn't been properly unmounted and may be corrupted. Since we don't know what files were being written at the time of the crash, the entire filesystem will need to be examined to repair any problems. This can take a lot of time.

The ext3 filesystem adds a transactional journal to ext2. Even when mounted read-write, the system is marked as "clean". However, if at boot time the journal contains information about incomplete writes, corruption may exist. In this case, since we know exactly which files were being written at the time of the crash (from the information in the journal), we can very quickly check and repair just the affected parts of the filesystem. Nonetheless, there are times when it's a good idea to perform an exhaustive check on an ext3 filesystem. It's possible for the hardware to introduce odd errors in rare circumstances. Some examples might be caused by gradual loss of power in a brownout, or severe vibration of the hard drive.

The `fsck` program is a front end to the standard filesystem checking programs on the system. The one for ext2/ext3, `e2fsck`, is used both to repair ext3 filesystems using the filesystem's journal and to exhaustively examine filesystems of either type.

If the root filesystem was not unmounted properly and has a journal, then when the kernel mounts the filesystem read-only, the kernel will read the journal and repair the filesystem. If `rc.sysinit` detects possible disk corruption, you'll be presented with a prompt that reads "Press Y within 5 seconds to force file system integrity check."




# Filesystem Recovery

- If / has journal, kernel examines it at boot
- /etc/rc.d/rc.sysinit runs `fsck` on filesystems marked in /etc/fstab
- `fsck` is a front end to other programs
- A "failed" `fsck` must be run manually

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



9

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 298 2994 or +1 (610) 794 3700.

If you press `y`, then all filesystems marked for checking in /etc/fstab will be treated as if they were "dirty", whether they're actually "dirty" or "clean". This includes ext3 filesystems. Normally it's okay to wait until the prompt times out.

Then, `rc.sysinit` will `fsck` the filesystems listed in /etc/fstab that have a positive integer in the sixth column. Any filesystems that have journals will use them to make rapid repairs, and any filesystems marked "dirty" (or being treated as such) will be exhaustively examined and repaired. If the filesystem is badly corrupted, `fsck` may "fail" and need to be re-run manually. You must provide the root password to get a shell from which you can run `fsck`, or the actual programs `fsck` runs; `e2fsck` (aka `fsck.ext2` and `fsck.ext3`), `dosfsck`, and so on.

`e2fsck` is used to repair both ext2 and ext3 filesystems. It should only be run against filesystems when they are not mounted or when they are mounted read-only. It takes as its argument the partition containing the filesystem to be checked. A number of switches are available, the most important of which are `-y`, which answers "yes" to all questions for which `e2fsck` would ordinarily prompt, and `-b <block>`, which instructs `e2fsck` to repair the filesystem using an alternate copy of the superblock. The appropriate location for the `<block>` argument may be determined using the `dumpe2fs` command. Because a corrupt filesystem may not be readable by `dumpe2fs`, it is advisable to have hard copy output of this information before a problem occurs.

Once filesystems have been repaired, it is possible to remount the root filesystem and bring it up from the `sulogin` state. However, it is preferable to type `exit`, which forces a complete reboot.

Sometimes the root filesystem is corrupted past the point where it is even mountable. Consequently, `/sbin/init` cannot be run. That forces us into the rescue mode described on the following page.

# Recovery Run-levels

- Pass run-level to `init`
  - on boot from GRUB splash screen
  - from shell prompt using: `init` or `telinit`
- Runlevel 1
  - Process `rc.sysinit` and `rc1.d` scripts
- Runlevel s, S, or single
  - Process only `rc.sysinit`
- emergency
  - Run `sulogin` only

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 265 2804 or +1 (616) 754 3700.

In recovery situations, it is often helpful, (and depending on the problem possibly necessary) to boot to a run-level where less services are active. For example, consider if you have a service that causes the machine to panic each time it tries to start. In this case, the road to recovery starts by preventing the service from starting, so you can successfully boot the machine to a stable state and determine the problem with the service. The below listed run-levels are of particular importance in system recovery situations

## Run-level 1

Booting to run-level 1 will cause the system to process the `/etc/rc.sysinit` script followed by each of the `/etc/rc.d/init.d` scripts called in `/etc/rc1.d/*`. By default, RHEL will only call the single script in this runlevel, which after some basic checks and cleanup will `exec init S`.

Switching to run-level 1 from some other run-level (3, 5, etc.) is a convenient way to kill all daemons as each of the `/etc/rc1.d/*` kill scripts will be processed.

## Run-level s, S, single

Booting to run-level single will cause the system to process the `/etc/rc.sysinit` script (if `/etc/inittab` is intact) If `/etc/inittab` is missing or corrupt, you can still boot to single mode, and in that case, you are given the root shell with no scripts processed.

Sometimes going to single user mode is overkill: interactive startup mode, invoked by typing "T" when "Welcome to Red Hat Enterprise Linux" appears at boot time, allows you to choose which services will run.

## "Run-level" emergency


While technically not a run-level, emergency mode shares many characteristics of the above listed run-levels. You can only access emergency mode during boot by passing emergency as a parameter from the grub prompt. No scripts will be processed, and you are given a root shell.

# Rescue Environment

- Required when root filesystem is unavailable
- Non-system specific
- Boot from CDROM (`boot.iso` or CD #1)
- Boot from `diskboot.img` on USB key

Rev RH133 RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

If the root filesystem is available and mountable, then you should be able to use it to fix problems that may occur. When it is not, then you must use a rescue environment. A rescue environment is a streamlined RHEL system that does not require the installed OS to run. Rather than working on the broken system itself, you work outside of the system in an environment that, while more limited than single user mode (or even `su` login mode), should provide enough tools to recover root.

There are several ways to boot into the rescue environment:

- Boot from CDROM, then type `linux rescue` at the `isolinux` prompt
- Boot from a `diskboot.img` USB drive, then type `linux rescue` at the prompt

# Rescue Environment Utilities

- Disk Maintenance Utilities
- Networking Utilities
- Miscellaneous Utilities
- Logging: /tmp/syslog Or /tmp/anaconda.log

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 268 2994 or +1 (919) 754 3700.

The rescue environment exists within a ramdisk image (referenced as `/dev/root`). Because of limitations on size and the number of inodes, many familiar utilities and device nodes are not available. However, tools related to disk maintenance (the probable reason for being in the rescue environment) and network connectivity are provided.

The following is an (incomplete) list of utilities provided by the rescue environment:

**Disk Maintenance Utilities, including:** a complete set of LVM utilities, for managing physical volumes, volume groups, and logical volumes; software RAID tools; swap commands; disk partition utilities; filesystem creators, checkers, debuggers, and labelers for ext2, ext3, jfs, msdos, vfat, and reiser filesystems.

**Networking Utilities, including:** network debuggers (`ifconfig`, `route`, `traceroute`, `host`); network connectivity tools (`ftp`, `rcp`, `rlogin`)

**Miscellaneous Utilities including:** shell commands (`bash`, `chroot`); process management tools (`ps`, `kill`, `killall`); editors (`vi`); mttools commands; kernel module mangement commands; archiving and compression tools (`dd`, `tar`, `cpio`, `gzip`); rpm; file manipulation commands (`cd`, `ls`, `mkdir`, `cp`, `mv`, `rm`)

Within the rescue environment, system logging information can be found in the file `/tmp/syslog`. Booting information is in `/tmp/anaconda.log`. Some configuration files (`modprobe.conf`, `netinfo`, and device files (`[sh]da`, `loop0`) are located in `/tmp` as well.

## Rescue Environment Details

- Filesystem reconstruction
  - Anaconda will ask if filesystems should be mounted
    - watch for error messages
    - `/mnt/sysimage/*`
    - `/mnt/source`
    - `$PATH` includes hard drive's directories
- Filesystem nodes
  - System-specific device files provided
  - `mknod` knows major/minor #'s

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 866 266 2994 or +1 (919) 754 3700.

The rescue environment will attempt to reconstruct the harddisk's filesystem under the mount point `/mnt/sysimage`. Since the rescue environment is often used on systems with damaged or misconfigured filesystems, however, this operation might or might not work. A corrupted partition table will appear to hang the rescue environment (a shell with `fdisk` is available under `Alt-F2`, however.) Using `linux rescue nomount` as the boot prompt directive disables automatic mounting of filesystems and circumvents the hanging caused by bad partition tables. Careful inspection of the output of the `mount` command should determine the state of the reconstructed filesystem.

Because the standard installation provides nearly 7000 device nodes, administrators seldom need to create device nodes directly. In the rescue environment, device nodes are only provided for the most basic devices, including any fixed disks the kernel was able to autodetect.

In order to access any other devices, such as a floppy drive, the relevant device node must be created with `mknod`. Fortunately, the rescue environment's version of `mknod` automatically associates the appropriate device driver major/minor numbers with well-known device names. For example, the device node for the hard disk on the secondary IDE controller can be created with `mknod /dev/hdc`.

## End of Unit 11

- Questions and Answers
- Summary
  - What are some things to check for
    - X problems?
    - Services problems?
    - Networking problems?
    - Boot problems?
  - How might you repair an ext2 filesystem?
  - What are some alternate boot methods?

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



14

For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and international copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 800 266 2904 or +1 (319) 754 3700.

### Important files covered in this Unit:

/tmp/syslog (in rescue environment)  
/mnt/sysimage/\* (in rescue environment)

### Important new commands covered in this Unit:

chroot  
fsck

# Unit 11 Lab

## System Rescue and Troubleshooting

---

**Goal:** To build skills in system rescue procedures.

### Sequence 1: Repairing the MBR in the rescue environment

The rescue environment provides a last resort for repairing an unbootable machine, even when the bootloader or the root filesystem is damaged or misconfigured. In order to access the rescue environment, you will need either a `boot.iso` cdrom on a network that has the Red Hat installation tree (the RedHat directory) available via NFS, or a Red Hat Enterprise Linux CDROM.

**Tasks:** Damage GRUB, leaving the machine in an unbootable state. Use the rescue environment to reinstall GRUB.

1. Use the following command to overwrite the first stage of GRUB in your Master Boot Record with zeros. Specify the blocksize carefully. If you write too many zeros, you will overwrite your partition table as well, and this will become a much more difficult exercise. (Note that the command below assumes you are using IDE drives. You might need to modify the destination device.)

**After typing the following command, check it three times and hit enter but once.**

```
dd if=/dev/zero of=/dev/hda bs=446 count=1; reboot
```

Congratulations -- you have just wiped out your boot sector, but your primary partition table will still be intact. Attempt a reboot to confirm that your system is unbootable. Use the Red Hat rescue environment to repair the system. A suggested sequence follows.

2. Load the rescue environment by booting from a Red Hat installation media (either CDROM or bootdisk floppy), and typing `linux rescue` at the `syslinux` boot prompt. Proceed with the normal installation defaults. Choose `nfs` image for the media type and use the following `nfs` information:

```
nfs server      : server1.example.com
nfs directory   : /var/ftp/pub
```

3. The rescue environment will ask if you wish to mount the hard drive's filesystems. Select "Continue" to mount the filesystems in read-write mode. Examine the output of `mount` to confirm that the filesystem was correctly reconstructed. You might want to refresh your memory by examining your disk's partitions with `fdisk`.

```
sh-2.04# mount
sh-2.04# fdisk -l /dev/hda
```

4. Note that your hard drive has been reconstructed under the mount point `/mnt/sysimage`. Examine `grub.conf` (on your hard drive) to confirm that it is appropriately configured.

```
sh-2.04# cat /mnt/sysimage/boot/grub/grub.conf
```

5. To reinstall GRUB, you must shift contexts, so that `grub-install` believes that the root of your filesystem is the `/mnt/sysimage` directory. Spawn a chrooted shell, run `grub-install`, and then exit

```
sh-2.04# chroot /mnt/sysimage
sh-2.04# grub-install /dev/hda
sh-2.04# exit
```

(or should the above fail to execute properly)

Type the command: `grub` at the bash prompt. This will place you into grub's command shell where you can enter the following commands:

```
grub> root (hd0,0)
grub> setup (hd0)
grub> quit
```

6. Now exit your rescue shell. Note that the rescue environment will unmount any partitions that you mounted.



## Sequence 2: Installing software in rescue mode

Use the following command to overwrite the `mount` command.

```
[root@localhost]# cp /bin/date /bin/mount
```

Congratulations -- you have just wiped out a key executable on your system. Upon attempting a reboot, you should find your system unbootable. Use the Red Hat rescue environment, along with its version of the `rpm` command and the library of RPMs provided by the installation tree, to repair the system. A suggested sequence follows.

1. Load the rescue environment by booting from a Red Hat installation media (either CDROM or boot iso cdrom), and typing `linux rescue` at the `syslinux` boot prompt.
2. The rescue environment will attempt to automatically mount the hard drive's filesystem. Examine the output of `mount` to confirm that the filesystem was correctly reconstructed.

```
sh-2.04# mount
```

3. Note that your hard drive's filesystems have been mounted under `/mnt/sysimage`. Determine which package contains the `mount` command:

```
rpm -qf --root /mnt/sysimage /bin/mount
```

4. Verify the `util-linux` rpm on your harddrive, using a `chrooted` invocation of `rpm`. **Do not forget to exit the chroot or step 5 will fail.**

```
sh-2.04# chroot /mnt/sysimage
sh-2.04# rpm -V util-linux
sh-2.04# exit
```

5. `rpm` should report that the `/bin/mount` executable has been modified. Reinstall the `util-linux` RPM, pulling the source from your installation tree (which has been NFS mounted under `/mnt/source`), again using a `chrooted` invocation of `rpm`

```
sh-2.04# rpm -ivh --force --root /mnt/sysimage
/mnt/source/RedHat/RPMS/util-linux*
```

Note that the `util-linux` package was installed (the hash marks indicate this), although you may see some errors at the end of the process. As it turns out, this is harmless error, although in a production environment, you would want to test this out fully.

6. Now `exit` your rescue shell. Note that the rescue environment will unmount any partitions that you mounted.

**Sequence 3: Troubleshooting Practice**

**Task:** Practice troubleshooting problems on a Red Hat Enterprise Linux system

1. Turn off iptables and mount the /var/ftp/pub directory from server1 if it is not currently mounted:

```
service iptables stop
chkconfig iptables off
mkdir /mnt/server1
mount server1:/var/ftp/pub /mnt/server1
```

2. Install the Troubleshooting Practice RPM:

```
rpm -ihv /mnt/server1/gls/RPMS/rhce-ts-*
```

3. Execute:

```
cd /usr/share/doc/ts<tab>
```

This is where the documentation for the redhat-ts RPM is located. Ensure that your computer is configured as closely as possible to the following specifications:

- a. Authenticate users from your local /etc/passwd file. That is, do not run any network authentication scheme such as NIS or LDAP.
- b. Use 192.168.0.254 (server1.example.com) as your resolver.
- c. Confirm that /usr/local/bin is part of your PATH environment variable.

The following items are required for some, but not all, troubleshooting problems. You may still do most problems if some of these items are missing.

- d. Change to runlevel 3, not runlevel 5. Confirm that the X server is not running (no startx). Only the X-related problems require this.
- e. Confirm that /home is a separate filesystem from the root filesystem and is local to the system (not an NFS mounted filesystem).

4. The Troubleshooting Practice problems come in two parts, each invoked by a separate command. The sections, commands, and number of problems in each section vary; therefore, run <command> count, where <command> is one of the following:

```
For Local      tslocal count
For Networking tsnetwork count
For Booting    tsboot count
```

Invoke the first local problem by running:

```
tslocal 1
```

This command will set up the problem and will explain the goal. The goal will be stored in the file /etc/ts for later reference. Spend three to eight minutes trying to solve the problem.

5. If you have not yet solved the problem, you may need a hint. Hints can be displayed by running the `tshint` command:

```
tshint local 1 1
```

This will display the first hint for the first `tslocal` problem. Continue to invoke hints until you get enough information to solve the problem or until you run out of hints:

```
tshint local 1 2
tshint local 1 3
[ and so on . ]
```

The `tshint` command will tell you when you have reached the end of the hints. Again, do not spend more than five to ten additional minutes on this problem.

6. Whether or not you have solved the problem, run the `tslesson` command:

```
tslesson local 1
```

This command will tell the lessons intended to be taught by the problem. Some `tslesson` messages also give step-by-step instructions on how to approach a particular problem.

7. If, after reading the hints and the lesson, you are unable to solve the problem, call the instructor for assistance.
8. Proceed to the second local problem:

```
tslocal 2
```

Again, consult the hints if you need assistance:

```
tshint local 2 1
tshint local 2 2
[etc.]
```

Be sure to read the lesson before continuing to the next problem:

```
tslesson local 2
```

Consult the instructor if you need assistance resolving this problem.

Proceed to invoke and solve the remaining local problem, using hints as needed and consulting the lesson before moving to the next problem.

9. Invoke the network problems, one at a time. For example:

```
tsnetwork 1
```

Again, the command will tell you what you need to do in order to resolve the problem. For hints, use the `tshint` command:

```
tshint network 1
```

After the problem is solved or you are ready to go on to the next problem, read the lesson:

```
tslesson network 1
```

Consult the instructor if you need assistance resolving this problem.

Proceed to invoke and solve the remaining network problems, using hints as needed and consulting the lesson before moving to the next problem.


10. Invoke and resolve the `tsboot` problems one at a time. Use the hints if, needed, and always consult the lessons. Consult the instructor anytime you need assistance resolving a problem.

# Sources of Help

- man & info documents
- `/usr/share/doc`
- HOWTOs

Rev RH133-RHEL4-1

Copyright © 2005 Red Hat, Inc.



For use only by a student enrolled in a Red Hat training course taught by Red Hat, Inc. or a Red Hat Certified Training Partner. Any other use is a violation of U.S. and International copyrights. No part of this publication may be photocopied, duplicated, stored in a retrieval system, or otherwise reproduced without prior written consent of Red Hat, Inc. If you believe that Red Hat training materials are being improperly used, copied, or distributed please email [training@redhat.com](mailto:training@redhat.com) or phone toll free (USA) +1 888 288 2994 or +1 (919) 754 3700.

## System man pages

Use system man pages for finding syntax information for various system administration utilities.

Many commands also come with `info` pages.

## Package Documentation

Most packages installed with Red Hat Enterprise Linux include documentation which is installed under `/usr/share/doc`. For example, the Grub boot loader comes with a very informative README file and other documents. The amount and quality of documentation will vary from package to package, but most packages include some files to supplement the `man/info` pages.

Documentation is not guaranteed to come in any specific format. Many developers distribute documentation in plain text files. Increasingly, developers are publishing package documentation in HTML format so it can be read with a web browser. Other common formats include PostScript (view with `ghostview` or send to a PostScript printer) and L<sup>A</sup>T<sub>E</sub>X (a SGML-based layout language).

## HOWTOs

The Linux Documentation Project <http://www.tldp.org> is an organized collection of documents which offer step-by-step instruction on various Linux tasks. These documents are called HOWTOs.

HOWTOs are distributed in a variety of formats including plain text and PostScript. They are also available in many different languages.



1. The first part of the document discusses the importance of maintaining accurate records of all transactions and activities. This is essential for ensuring transparency and accountability in the organization's operations.

2. The second part of the document outlines the various methods and techniques used to collect and analyze data. This includes both qualitative and quantitative approaches, as well as the use of advanced statistical tools and software.

3. The third part of the document focuses on the interpretation and application of the results. This involves identifying key trends, patterns, and insights that can inform decision-making and strategic planning.

4. The final part of the document provides a summary of the findings and offers recommendations for future research and practice. This includes suggestions for how the organization can improve its data management and analysis processes.

.....

.....

.....



The following information is provided for your information only. It is not intended to be used as a basis for any action. The information is provided for your information only. It is not intended to be used as a basis for any action.

The following information is provided for your information only. It is not intended to be used as a basis for any action. The information is provided for your information only. It is not intended to be used as a basis for any action.

The following information is provided for your information only. It is not intended to be used as a basis for any action. The information is provided for your information only. It is not intended to be used as a basis for any action.

The following information is provided for your information only. It is not intended to be used as a basis for any action. The information is provided for your information only. It is not intended to be used as a basis for any action.

The following information is provided for your information only. It is not intended to be used as a basis for any action. The information is provided for your information only. It is not intended to be used as a basis for any action.