

# RHCE

## BOOT CAMP

A whole week of geeky fun!



## ABOUT THE CLASS

- 40 hour boot-camp style prep course
- Monday-Thursday: Lecture and labs
- Friday: Practice exam
- Hours:
  - 8:30am - 5:00pm
  - Lunch: 11:45am - 1:00pm

# ABOUT THE INSTRUCTOR

- Nathan Isburgh
  - [instructor@edgecloud.com](mailto:instructor@edgecloud.com)
  - Unix user 15+ years, teaching it 10+ years
  - RHCE, CISSP
  - Forgetful, goofy, patient :)

# ABOUT THE STUDENTS

- Name?
- Time served, I mean employed, at Rackspace?
- Department?
- Unix skill level?
- First attempt at RHCE?
- What are you most worried/interested about with RHEL?



# EXPECTATIONS OF STUDENTS

- Strong foundation in basic Linux use and administration
- Ask Questions!
- Complete the labs
- Email if you're going to be late/miss class
- Have fun
- Learn something
- Pass your exam next week!

# ABOUT RHCE EXAM

- The RHCE exam is 100% hands-on. You will have a computer and a list of tasks to accomplish. When you are finished, an automatic grader will check your machine to be sure that everything is set up correctly.
- There are no questions, only tasks.
- You will have 3.5 hours and access to all RHEL 5 Server software.
- You will not have internet access.

# RHCT/RHCE

- The exam itself is divided between Technician tasks and Engineer tasks. If you complete an adequate percentage of RHCT tasks, you will get the RHCT certification.
- If you complete an adequate number of RHCT tasks, **AND** an adequate number of RHCE tasks, you will get your RHCE, which includes RHCT certification.
- RHCT = local machine settings, basic network access
- RHCE = network services

# TO PASS EXAM:

- Details specific to RHEL v. 5
- Basic System Administration
- Configuration of all major services / daemons
- Implementation of basic security / traffic filtering technologies
- **Locating Local Reference Materials (--help, man)**



# SCHEDULE

Monday	Bootimg, Packages, System Administration
Tuesday	File Systems, Users, Kernel Features
Wednesday	File Sharing, Web Services, Net Services
Thursday	DNS, Email, Network Security
Friday	Practice Exam

```
slideshow.end();
```

# RHCE

## BOOT CAMP

The Boot Process



# OVERVIEW

- The boot process gets a machine from the useless off state to the feature rich operating system we all know and love
- Requires cooperation between hardware and software to correctly hand off processing
- Akin to the life cycle of a human - birth, newborn, infant, toddler, teen, adult



# BIRTH

- Power switch flipped on
- Electricity flows from wall, through power supply where it gets converted to the levels necessary for the computer, and on to the motherboard, drives, CPU and more
- Completely unaware of the world or even what's attached to the motherboard.

# INFANT

- BIOS - Basic Input/Output System - CPU looks for instructions starting at a specific address, which happens to be where BIOS resides. BIOS initializes and starts the....
- POST - Power On Self Test - A simple set of tests that BIOS performs to verify basic functioning of attached hardware.
- Like an infant, extremely limited understanding of world
- Searches for valid MBR, loads the software found there and transfers control to the...

# TODDLER

- Boot Loader - Special software installed to the MBR of the boot partition which selects and loads the kernel.
- Can be configured to immediately load the default OS, or can offer choice to user
- Slightly better understanding of world - can read linux filesystems, sometimes includes powerful debugging and configuration support.
- Main job: select and load kernel, transfer control to kernel

# TEENAGER

- Dreaded teenager age: knows a lot about the world, but doesn't contribute a thing. Still pretty useless.
- Kernel loads and initializes. Device drivers are loaded and initialized. Basic hardware checks performed.
- The First Process is *created from nothing*: `init`



# ADULT

- init loads the inittab, specifying what software needs to be started. init starts running all of the specified startup scripts
- Services are started by init, including network configurations, X Windows, network services, databases, etc.
- At this point, the machine is finally becoming useful: otherwise, an adult
- Eventually, login processes are started and the boot process is complete!

# MORE ON INIT

- init's configuration file is `/etc/inittab`
- This file details actions taken for certain global events, like ctrl-alt-delete and UPS powerfail and powerrestore alerts.
- This file tells init what needs to be done for a given runlevel as well as what the default runlevel is.
- A runlevel defines what services are running...

# RUNLEVELS

- Runlevels:
  - S: System startup
  - 0: OS stopped, machine halted ( usually powers off as well )
  - 1: Single user mode - for maintenance
  - 2: Multiuser, no NFS shares
  - 3: Full multiuser, TUI
  - 4: Unused
  - 5: Full multiuser, GUI
  - 6: Reboot

# RUNLEVELS

- `telinit`: Signal the init process to change the current runlevel
- Switching runlevels is fairly uncommon - generally only used if system maintenance needs to be performed
- Runlevels can be used to control what services a machine provides, and can sometimes be useful to quickly reconfigure a machine for a new task



# INIT SCRIPTS

- What is actually running in a given runlevel is defined by the init scripts for that level.
- That standard location for the init scripts is:
  - `/etc/rcX.d`
  - Where the `x` corresponds to the runlevel
- For example, `/etc/rc5.d` contains all of the init scripts that, combined, provide runlevel 5 service

# RC DIRECTORIES

- The files in the rc directories start with either an S or a K:
  - S means to start the service, ie run the command with “start” as an argument
  - K means to kill the service, ie run the command with “stop” as an argument
- After the S or K, there is a two digit number which is used for ordering the execution of the scripts

# ENTERING A RUNLEVEL

- So when the init process enters a runlevel, the steps are:
  - Run all of the Kill scripts, in order, with “stop” as an argument
  - Run all of the Start scripts, in order, with “start” as an argument

# INIT SCRIPTS

- If you look closely, you will see that `/etc/rcX.d` actually holds a collection of symbolic links
- The actual script files are stored in `/etc/init.d`
- The main reason for this is so that there is only one copy of each init script, reducing the chance that a script change won't be reflected in all runlevels.
- You can run the scripts directly, or use the `service` command to start/stop various components of the OS.



# MANAGING INIT SCRIPTS

- You can manage the links to the init scripts manually, or you can use the `chkconfig` command to get the job done:
- **`chkconfig --list`**
  - List all processes and display their default status in each run-level.
- **`chkconfig [--level levels] name <on|off|reset>`**
  - This command will modify the `chkconfig` configuration for a particular service, setting it on/off for the given runlevels.

# DAEMONS

- A daemon ( or demon ) is just a persistent process that performs some action or service. Daemons are what make machines useful. Examples:
- `httpd`: Web services
- `mingetty`: Watches terminals and starts login processes
- `mysqld`: Database services
- `syslogd`: Logging services

# GRUB

- Grand Unified Boot Loader
- GRUB's configuration file is `/boot/grub/grub.conf`, which is configured as follows:

```
default=0

timeout=10

splashimage=(hd0,0)/grub/splash.xpm.gz

title RedHat Enterprise Linux

    root (hd0,0)

    kernel /vmlinuz ro root=LABEL=/

    initrd /initrd
```

# GRUB SHELL

- Command mode – Pressing “c” while the boot menu is displayed will provide the user with the GRUB shell, where a limited set of commands can be used to explore the filesystem before booting. A full list of the commands available can be found by pressing Tab while in command mode.
- Editing mode – Pressing “e” while the boot menu is displayed will provide the user with the opportunity to edit a line in GRUB's configuration file.
- Append mode – Pressing “a” while the boot menu is displayed will allow the user to append to the kernel line for the default kernel in GRUB's configuration file
- Esc – can be pressed at any time to return you to the previous menu



```
slideshow.end();
```

# RHCE

## BOOT CAMP

Package Management



# RPM

- Redhat Package Manager
- RPM's provide full software packaging features: pre-install scripts, post-install scripts, dependencies, meta information, and an installed software database to name a few.
- The RPM system maintains a database of all installed software on a machine - this is useful for tracking and updating reasons, as well as dependency verification and software management.



# RPM

- **rpm**: The Redhat Package Manager tool. Provides interface to RPM system, performing queries, installs, upgrades, uninstalls and general database maintenance operations.
  - **-i** option: install the given package
  - **-q** option: query the database
  - **-e** option: erase the given package from the system

# RPM QUERIES

- Below are just a few examples of the types of queries you can run against the RPM database.
  - **rpm -qa** Queries for the names of all installed rpms.
  - **rpm -qi** Queries the rpm database for package information.
  - **rpm -qf** Determines which rpm a file is associated with.
  - **rpm -ql** Queries the rpm database to determine which files are associated with a particular rpm.
- With any of these commands, you can add the **-p** option to run the command against a package before it is installed.

# RPM INSTALLATION VERIFICATION

- In addition to storing information about where a package is installed, rpm also stores permissions, file sizes, md5sums, and ownership information. This information can be easily referenced to see if anything has been changed.
  - **rpm -Va** Verifies all installed packages.
  - **rpm -Vi <package>** Verifies given package.
- Rackspace Best Practice Example
  - `rpm -Va | grep ^..5`

# RPM VERIFY OUTPUT

- **S** File Size differs
- **M** Mode differs (includes permissions and file type)
- **5** MD5 sum differs
- **D** Device major/minor number mismatch
- **L** readLink(2) path mismatch
- **U** User ownership differs
- **G** Group ownership differs
- **T** mTime differs
- **C** SELinux Context differs



# EXTRACT RPM CONTENTS

- Use this technique to make a clean working copy of the files and directories that would be installed with a package.
  - `cd /temp/dir`
  - `rpm2cpio /path/to/package | cpio -i -d -m`
- This would allow you to:
  - Replace one corrupted file without un-installing and then re-installing a package
  - Compare original configuration files versus modified files in the running system to quickly locate changed lines, for example with the 'diff' utility

# YUM

- yum: Yellowdog Updater Modified
  - Supports package installation over the network through repositories.
  - RPM backend
  - Simple interface

# REPOSITORIES

- Repositories of packages must be listed in files in the `/etc/yum.repos.d` directory with names ending in `.repo` and having a format like:
  - `[label-for-repo]`
  - `name = descriptive text`
  - `baseurl = protocol://path/to/directory/of/packages`
- Access to the Red Hat Network, including any Satellite Servers, is implemented through a plugin to the yum tool itself and not as a repository definition in the above format.

# LAB

1. Connect to <http://server1.example.com> and read the information there.
2. Download the OpenOffice archive from `server1` and choose an appropriate location to extract all its RPMs
3. Install the `createrepo` package and use it to turn your collection of OpenOffice packages into a yum repository
4. Add that repository to your local yum configuration
5. Using yum, install the “`openoffice.org3-writer`” package, and/or any others from your new repository



```
slideshow.end();
```

# RHCE

## BOOT CAMP

System Administration



# X WINDOWS

- X Windows was developed in the 1980's to provide an intelligent GUI system for UNIX.
- It is an extremely simple client/server model, that is exceptionally easy to extend, hence it's power and world-wide adoption.



# LAYERS

- X Windows is built on a layered concept:
  - X Server
  - Window Manager
  - Desktop
- Also, a display manager runs to provide login services.

# WINDOW MANAGERS

- Special type of X Clients which encapsulate other clients, allowing them to be moved, resized, or “iconified.” They also provide the desktop theme, configurable menus, panel utilities, and session management. RHEL ships with `metacity`, `kwin` and `twm`. These window managers provide the core functionality of the GUI.
- Generally a desktop is run in addition to the window manager, though `twm` is provided as a fallback if a desktop won't start

# DESKTOPS

- Fully integrated graphical environments, sitting on top of a window manager. Usually provides copy/paste features, lots of extra tools/utilities to run and configure a graphical environment. RHEL ships with GNOME and KDE.
- Configured with `DESKTOP` environment variable, set in `/etc/sysconfig/desktop`

# DISPLAY MANAGER

- X equivalent of the text-based login program. RHEL 5 ships with gdm and kdm. Display managers are usually started in run-level 5 from the `/etc/X11/prefdm` script.
- Configured with `DISPLAYMANAGER` environment variable, set in `/etc/sysconfig/desktop`



# XFREE86

- XFree86 was the first open source clone of the X Window system, released in 1991.
- XFree86 formed the de facto GUI platform for Linux, and indeed all of X Windows development for the '90s and into the early 2000's
- Unfortunately, in 2004 the XFree86 project adopted a license change which GNU did not particularly care for, and almost all distributors switched to X.Org.

# X.ORG

- The X.Org Server stepped into the picture in 2004 as a splinter off of the XFree86 project.
- Since they didn't muck with the license, most distributors jumped over to X.Org for their X Windows needs, and to this day X.Org remains the GUI platform of choice for Linux implementations.

# CONFIGURATION

- Configuring X Windows often requires at least Bachelors in Computer Science with a Minor in Great Luck.
- The main configuration file for X.Org is `/etc/X11/xorg.conf`.
- Reading the associated man pages is a must.
- Relying on the GUI configuration tools to help with X Windows configs is a Good Idea

# CONFIGURING X WINDOWS

- For initial configuration of Xorg, you can use the command `system-config-display`.
- This can be run with the `--noui` and `--reconfig` options to completely removing any existing configurations and replace it with conservative settings.



# NETWORK CONFIGURATION

- There are two main approaches to configuring a machine for network access:
  - Static configuration
  - Dynamic configuration
- Static configuration uses set parameters for the configuration, which is known by the machine and the network and never changes. Generally used with servers.
- Dynamic configuration configures network machines on the fly, where a service on the network provides all configuration parameters to a machine when it joins the network. Generally used with workstations.

# DYNAMIC CONFIGURATION

- Dynamic configuration is the easiest to use.
- The machine just needs to set up its interfaces with the DHCP protocol.
- DHCP: Dynamic Host Configuration Protocol.
- A lease is obtained from the DHCP server, providing all network configuration details for the client. The lease expires after some amount of time and is renewed by the client to maintain network access.

# STATIC CONFIGURATION

- Static configuration requires four configuration parameters in order to allow full network functionality:
  - IP Address
  - Netmask
  - Default Gateway or Router
  - DNS Server(s)

# DNS?

- Domain Name Service: This is the glue between network names and IP addresses.
- Remember: Humans like names, computers like numbers. DNS is a service like so many others, mapping names to numbers and numbers to names. Mostly a convenience.
- Also provides for email functionality, geographic load balancing and limited service failover capabilities.



# STATIC CONFIGURATION

- The first two components of static configuration are IP address and netmask.
- These provide LAN-level access
- `ifconfig`: Network Interface configuration tool
- Basic idea:
  - `ifconfig eth0 192.168.0.100 netmask 255.255.255.0`

# GATEWAYS

- The third configuration parameter is the default gateway.
- Provides access to *inter-networking*, or moving from just the local LAN to other LAN's
- `route`: Kernel routing table tool
  - Displays and manipulates network routing table
  - `route add default gw 192.168.0.1`

# DNS SERVERS

- Final piece of configuration information.
- List of one or more IP addresses which provide the DNS service, allowing name to IP address mapping
- Very simple to configure. Add `nameserver` lines to `/etc/resolv.conf`:
  - `nameserver 192.168.7.15`
- Also consider `/etc/nsswitch.conf`

# STATIC CONFIGURATION

- Once all four pieces of information are configured on the system, full network service will be available.
- Best practice:
  - Configure IP address and netmask. Check LAN connectivity.
  - Configure default gateway. Check intra-LAN connectivity.
  - Configure DNS: Check name resolution.



# ONE MORE THING...

- `ifconfig` and `route` directly manipulate the running kernel, and are not permanent changes to the system. After a reboot, changes will be lost.
- To make IP address, netmask and gateway changes permanent, you have to edit two configuration files:
  - `/etc/sysconfig/network-scripts/ifcfg-eth0`
  - `/etc/sysconfig/network`

## /ETC/SYSCONFIG/NETWORK

`NETWORKING={yes|no}`

`HOSTNAME=<fqdn>`

`GATEWAY=<gateway ip>`

`NISDOMAIN=<nis domain name>`

# IFCFG-\* FILES

- To configure a device to use dhcp, the ifcfg file should contain the following:

```
DEVICE=eth0
```

```
BOOTPROTO=dhcp
```

```
ONBOOT=yes
```

# IFCFG-\* FILES

- To configure a device with static settings, the ifcfg file should contain the following:

```
DEVICE=eth0
```

```
BOOTPROTO=static
```

```
IPADDR=<ip>
```

```
NETMASK=<netmask>
```

```
ONBOOT=yes
```

```
GATEWAY=<gateway ip> * Only if different than default
```



# LAB

1. Determine your current network settings (which were assigned by DHCP) and change your machine to a static network configuration using these settings.
2. When you are satisfied with your configuration, restart the network service to put your changes into effect.
3. Test your connectivity to `server1` to make sure you are still online.
4. Refer back to DHCP settings if necessary to correct any mistakes in your static configuration.

# PRINTING

- The new, and preferred printing system for Linux, is CUPS - the Common Unix Printing System.
- CUPS supports IPP protocol (based on HTTP/1.1) and can communicate with LPD print servers
- CUPS can be administered several ways:
  - by manually editing `/etc/cups/cupsd.conf` and `/etc/cups/printers.conf`
  - by using the `system-config-printer` tool, or
  - by using the web administration interface on port 631

# CUPS

- CUPS tools and commands:
  - `lpstat`: used to view status of configured printers
  - `lp`: Create a print request
  - `cancel`: Cancel a pending print request
  - `lprm`: Cancel a pending print request
  - `lpadmin`: printer access control

# PRINTER CONTROL

- Printing under CUPS is a two-step process.
  - First, a job is *spooled* or *queued* for printing in the print spool.
  - Second, the cups daemon pulls jobs from the print queue and feeds them to the appropriate printer.
- Access to the print queue is managed with the `accept` and `reject` commands
- Whether cupsd hands print jobs to the printer is controlled with the `enable` and `disable` commands.



# LAB

1. Configure a “dummy” printer which doesn't correspond to any actual device, and make this the default queue on your machine.
2. Submit a text file to be printed from the command line.
3. View the status of the print job (which should never complete) using both the CLI and Web interfaces.

# CRON

- `cron` is the cron daemon. Cron provides for the ability to execute commands on a regular basis.
- Generally used to run hourly, daily and weekly type system maintenance scripts.
- Also useful to run reports, cleanup jobs and much, much more.

# USING CRON

- Cron is controlled through crontab files.
- There are system-wide crons, accessible under  
`/etc/cron.*`
- Every user has their own crontab, accessible through the `crontab` command

# SYSTEM CRONS

- `/etc/crontab` defines the system cron jobs.
- Many distributions use the `run-parts` script to execute all scripts found in `/etc/cron.hourly`, `/etc/cron.daily`, etc on the appropriate schedule.
- `/etc/crontab` defines the times for each schedule: hourly, daily, weekly, monthly



# CRONTAB

- `crontab`: View, edit or remove crontabs
  - The `-l` option prints the crontab. The `-e` option opens the crontab for editing. The `-r` option removes the crontab.
  - Root can work with the crontab for any user by specifying the username on the command line:
    - `crontab -e -u bob`

# CRONTAB SYNTAX

- There are two main components to a crontab entry:
  - The timespec specifies when the command should be run
  - The command is what gets executed every time the timespec is matched

# CRONTAB TIMESPECS

- The timespec is broken down into 5 fields, separated by spaces:
  - minute hour day-of-month month day-of-week
- Each field can contain a number, a range of numbers, a comma-separated list of numbers, an asterisk or a number slash division rate
- Mostly self-explanatory - some examples will help...

# TIMESPEC EXAMPLES

- 0 23 \* \* \* *11pm every day*
- 30 \* \* \* 1-5 *30 minutes after every hour, M-F*
- 0 7 1 \* \* *7am, first of every month*
- \* \* \* \* \* *Every single minute*
- 0,10,20,30,40,50 \* \* \* \* *Every 10 minutes*
- \*/5 8-17 \* \* 1-5 *Every 5 minutes, 8am-5pm, M-F*



# EXAMPLE CRONTAB

```
01 4 * * * /usr/local/bin/restart-webserver
00 8 1 * * /usr/bin/mail-report boss@mycompany.com
*/5 * * * * /monitor/bin/check-site -e admin@mycompany.com -o /var/log/check.log
```

- There are various additional options and features available to the cron system. Check the man pages for reference:
  - cron, crontab ( sections 1 and 5 )

# LAB

1. Create a cronjob for the user root that checks the amount of available space on the system every Friday at 12:34pm.
2. Create a cronjob as a regular user that lists the contents of /tmp at 3:54am on Sunday, January 2.

# LOGS

- One of the easiest places to find the cause of a problem is in the log files.
- Log files store informational messages from software. The types of messages include debug information, status information, warnings, errors and more.
- Some applications manage their own log files. Others use the system-wide logging package...

# SYSLOG

- `syslog` - The system logger. A framework consisting of a library, a daemon, a configuration file and logs.
- Any application can use the library and log messages through `syslog` with simple function calls.
- Log messages consist of 3 parts:
  - Facility
  - Level
  - Message



# SYSLOG

- The facility describes what part of the operating system generated the message, and is selected by the software:
  - `auth`, `authpriv`, `cron`, `daemon`, `ftp`, `kern`, `lpr`, `mail`, `news`, `security`, `syslog`, `user`, `uucp`, `local0`–`local7`
- The level represents the importance of the message, and is also chosen by the software:
  - `emergency`, `alert`, `critical`, `error`, `warning`, `notice`, `info`, `debug`

# /ETC/SYSLOG.CONF

- `/etc/syslog.conf` defines where all of the log messages should go. Destinations include files, screens of logged in users, console, other syslog servers.
- Basic file format:
  - `facility.level destination`
- Examples:
  - `*.err /dev/console`
  - `mail.* /var/log/maillog`
  - `*.info;mail.none;authpriv.none /var/log/messages`

# /VAR/LOG

- `maillog`: messages from the email subsystem
- `secure`: authentication and security messages
- `cron`: cron messages
- `boot.log`: boot messages
- `messages`: catch-all
- `dmesg` : hardware and kernel events generated before `syslogd`

# LOGS

- As mentioned earlier, not all software uses the `syslog` framework to handle its logging. Quite a bit of software manages its own logs.
- This can make it difficult to track down all of the log locations on an unfamiliar system. The best way to handle this is to start from the init scripts...



# LOCATING APPLICATION LOGS

- To track down the log file location for an application, you need to find its configuration file so you can see where the logs are being written.
- Of course, finding the configuration file might be just as difficult, so it's best to start at the source.
- `init` starts all of the system services, and so there is an `init` script somewhere that is starting up the application in question.
- The `init` script almost always references the configuration file

# LOCATING APPLICATION LOGS

- Now that the configuration file location is known, it only takes a few moments to scan through it and find out where logs are being written.
- As for the format of the log file, that's completely dependent on the application. Some will be similar to `syslog`, others, like `Apache` or `Qmail`, will be completely foreign looking.
- Fortunately, a little common sense and judicious application of Google Ointment will get the information you seek.

# MAINTAINING LOGS

- `/etc/logrotate.conf`
  - This is the main configuration file for logrotate.
- `/etc/logrotate.d/`
  - EVERYTHING in this directory will be parsed as if it is a logrotate configuration file. Usually, applications such as Apache and Sendmail will have configuration files in this directory to control how their logs will be rotated.
- `logrotate -vf /etc/logrotate.conf`
  - Can be run as root at any time to force log rotation and check for errors.

# LAB

1. Take a few minutes to browse through the various logs in `/var/log`. Familiarize yourself with the kinds of information available.
2. Browse the man page for `syslog.conf`
3. Find where the audit service keeps its log and add a corresponding new entry to your logrotate configuration.



```
slideshow.end();
```

# RHCE

## BOOT CAMP

Filesystem Administration



# PARTITIONING

- What is partitioning?
  - Splitting up a hard drive into organizable chunks
- Why?
  - Isolates filesystem corruption
  - Simplifies/speeds backups
  - Allows optimizing filesystems to tasks



# FDISK

- `fdisk`: partitioning tool
  - Works on one disk at a time, allows for viewing and manipulating partition table.
  - Online help ( hit 'm' ) makes tool easy to use
- At boot, the kernel loads a copy of the partition table into memory. Most partition editing commands only update the partition table on the drive, and not in memory. As such, the command `partprobe` can be run to update the information that the kernel has in memory.

# MKFS

- `mkfs`: format a device to create a new filesystem
  - “Paints the parking stripes” for the filesystem structure
  - Creates superblock, block groups, superblock copies, bitmaps and inode tables and creates basic structure on disk
  - Through `-t` option, `mkfs` can create different types of filesystems

# EXT<sub>2</sub>

- Benefits
  - Default file system for pre - 7.x versions of Red Hat
  - Heavily tested / Rock solid stability
- Drawbacks
  - Does not have a journal
  - File system check (fsck) required to mount a “dirty” file system
  - System offline and unavailable while fsck is running

# EXT<sub>3</sub>

- Benefits
  - Default file system of 7.x Red Hat releases and newer
  - Based on proven stability of Ext2
  - Has journal for increased reliability
- Drawbacks
  - Inodes allocated when file system is created (other file systems create them dynamically as they are needed)
  - Not as efficient as other file systems when dealing with lots of small files



# JOURNALING

- Journaling - How does it help?
- Deleting a file in Linux requires two steps:
  1. The file's directory entry must be removed.
  2. The file's inode must be marked as free in the free space map.
- If step 1 happens before a crash, an inode will be orphaned and the file will be lost.
- If step 2 happens first before a crash, the inode will be marked free and will possibly be overwritten.
- Journaling keeps a journal of the changes that are planned for the file system ahead of time. The journal can then replay the changes in the journal at any time to keep the file system clean.

# FILESYSTEM INTEGRITY CHECKS

- `fsck`: Filesystem Check
  - Generally only run when a filesystem needs it:
    - Mount count
    - Last check
    - Dirty
  - Checks all of the filesystem structures for accuracy and completeness

# FILE SYSTEM TOOLS

- `e2label`: View/set filesystem label
- `tune2fs`: View/set filesystem attributes
- `mount/umount`: You better know these already. :)

# FSTAB

- `/etc/fstab` is parsed during boot by `rc.sysinit` to determine what file systems should be mounted and how. After boot, this file is referenced by the `mount` command.

- The file is space delimited and organized as follows:

```
device    mount_point  fs_type  options  dump  fsck
```



# LAB

1. Using `fdisk`, create a new 100MB partition.
2. Create a new filesystem on this partition using `ext3`, a blocksize of `1k`, and a reserve space of 2%. Confirm settings with `tune2fs`. Mount the new filesystem as `/u01` and set it to mount at boot.
3. Un-mount the `/u01` filesystem and force an integrity check. Re-mount the `/u01` filesystem. Use `e2label` to set the filesystem label on `/u01` to `'u01'`.

# AUTOMOUNT

- The `autofs` service can be configured to monitor certain directories and automatically mount a file system when a call is made to files in that directory.
- When `autofs` starts, it parses the configuration file `/etc/auto.master` to determine which directories it should be monitoring. Each directory can then have its own configuration file determining how each file system should be mounted, or the default file `/etc/auto.misc` can be used.

# AUTO.MASTER

- Basic format for `auto.master`:
- Path      Config file
- `/misc      /etc/auto.misc`
- This tells `automountd` to “watch” the `/misc` pathname for activity, and if activity is observed, consult `/etc/auto.misc` for instructions.

# AUTOMOUNT PATH CONFIG FILES

- Basic syntax:
- path    options                    mount device
- `nfs    -fstype=nfs,ro    nfsserver:/share/nfs`
- This tells `automountd` to dynamically mount the `nfs` share “`/share/nfs`” on `nfsserver` when access is attempted on the “`nfs`” pathname under a watched pathname ( `/misc` for example )



# LAB

1. Configure your server to automatically mount `/share` as an NFS share from `server1` to `/server1/share` when a process changes directories there.

# EXTENDED ATTRIBUTES

- Ext2 and Ext3 support attributes that affect how data can be manipulated.
- The `chattr` command can change these file system attributes.
- The `lsattr` command will list the file system attributes.
- Extended attributes can only be set by the root user, unless the `user_xattr` mount option is in effect.

# COMMON EXTENDED FILE ATTRIBUTES

- `i` Immutable. The file can not be changed. By anyone. Period.
- `a` Append-only. File can only be opened for appending.
- Most of the others are experimental and/or esoteric. Surprising? ;)

# ACL'S

- Ext3 supports access control lists, which allow for more flexible permissions than standard file system permissions.
- ACL's can be listed with the `getfacl` command.
- They can be modified with the `setfacl` command.
- To use ACLs, a file system must have the `acl` mount option.
- Use `dumpe2fs -h <block device node>` to see default mount options.



# ACL EXAMPLES

- `setfacl -m u:bob:w memo.txt`
- `setfacl -x g:ru report.txt`
- `setfacl -m g:ru:r another-report.txt`

# QUOTAS

- Quotas are used to limit how many filesystem resources are available to a user.
- Inodes and space are controllable.
- Hard and soft limits are available, with grace periods.
- Enabling quotes is an involved process...

# ENABLING QUOTAS

- `usrquota` and `grpquota` options must be enabled on the filesystem mount
- Two files will be created at the root of the filesystem: `aquota.user` and `aquota.group`
- Run `quotacheck -mavug`
- Turn on quotas by running `quotaon` with the mount point as argument.
- Now you can use `edquota` to set up the quotas
- See man pages: `quota`, `repquota`, `edquota`, `quotaon`, `quotacheck`

# LAB

1. Create a quota for the user `student` with:
  - a block soft limit of 100M and a hard limit of 150M
  - a soft inode limit of 30 and a hard inode limit of 100
2. Create a quota for the group `gdm` so that its members collectively have:
  - a block soft limit of 200M and a hard limit of 300M
  - a soft inode limit of 50 and a hard inode limit of 200



# SELINUX

- Every process or object has an 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

# SELINUX

- Adding the `-Z` option to several commands will show how they are running in regards to SELinux:
  - `ps -Z` lists the process contexts
  - `ls -Z` lists the file contexts
- To change the context of a file, you can use the `chcon` command:
  - `chcon -R --reference=/var/www/html <file>`
- SELinux will log all policy violations to `/var/log/messages` as AVC (access vector cache) denials.

# CONTROLLING SELINUX

- The tool `system-config-selinux` can be used to configure SELinux.
- The file `/etc/sysconfig/selinux` can be edited.
- The command `getenforce` will show the current SELinux mode, and `setenforce` will allow you to change the mode.
- To change the SELinux mode during boot, you can pass the `enforcing=0` option to the kernel in GRUB.
- See also the members of the “`polycoreutils`” and “`setroubleshoot`” packages.

## LAB

1. With SELinux enforcing, configure a website to be served from `/srv`
2. Don't focus on advanced Apache settings, accomplish this in the simplest way possible: just change the global `DocumentRoot`. Further Apache exercises will follow later in class.
3. Populate a simple `index.html` file. Plain text is acceptable.
4. The `setroubleshoot` tool is useful here. Don't be confused by any typos in its output.



```
slideshow.end();
```

# RHCE

## BOOT CAMP

Users and Groups



# USERS AND GROUPS

- Users and Groups define access to the operating system through the file permission scheme.
- Root is the super user, and the only user with special permissions
- Every user is a member of at least one group, which is called their primary group. The main purpose of this primary relationship is to define group owner of created files.
- Users can have a secondary group membership in as many groups as needed. These secondary relationships exist to broaden a user's access to the files on the system.



# CONFIG FILES

- User information is stored in two files:
  - `/etc/passwd`
  - `/etc/shadow`
- Group information is stored in one file:
  - `/etc/group`

## /ETC/PASSWD

- List of user records, one per line, with columns separated by colons. Format:
- `login:x:userid:groupid:gecos:homedir:shell`
- Examples:
  - `root:x:0:0:root:/root:/bin/bash`
  - `mysql:x:27:27:MySQL Server:/var/lib/mysql:/bin/bash`

# /ETC/SHADOW

- Similar colon-separated-column list of records:
- `login:password:password aging fields`
- Aging fields track dates for password resets, locks, etc
- Examples:
  - `root:pB8msP1fCbCqc:13904:0:99999:7:::`
  - `nisburgh:vRoPw6a/jQsp.:14466:0:99999:7:::`

# /ETC/GROUP

- Same colon-separated-column list of records format
- `groupname:grouppassword:groupid:secondarymembers`
- Group passwords allow temporary management to a group, are rarely used and not set up by default
- Examples:
  - `daemon:x:2:root,bin,daemon`
  - `apache:x:48:jack,nisburgh`



# MANAGEMENT

- While it is possible to edit the three files directly, it's easier and safer to use the management commands to create, modify and delete users and groups:
  - `useradd`, `usermod`, `userdel`
  - `groupadd`, `groupmod`, `groupdel`

# USERADD

- `useradd`: Add a new user to the system
- Accepts various arguments to control the settings on the user account. Most common is the `-g` option to specify the primary group of the user, and the `-G` option to list secondary group memberships. Examples:
  - `useradd lisa`
  - `useradd -g clowns -G trouble,simpson bart`

# USERMOD, USERDEL

- `usermod`: Modify a user's settings. Example:
  - `usermod -G detention bart`
- `userdel`: Remove a user from the system. Main option to consider is `-r`, which tells `userdel` to remove the user's home and spool directories. Example:
  - `userdel moe`

# GROUP COMMANDS

- `groupadd`: Adds a new group to the system. Example:
  - `groupadd bullies`
- `groupmod`: Mainly used to rename a group. Example:
  - `groupmod -n mktg mkg`
- `groupdel`: Remove a group. Example:
  - `groupdel microsoft`



# PASSWORDS

- `passwd`: Change login password.
- Root can change the password for any user on the system
- Root can also setup password aging, allowing for timed password resets and account disabling ( or use `chage` )
- `passwd` is also the preferred way to lock a user account:
  - `passwd -l mary`

# PASSWORD AGING

- To set the maximum lifetime for a user's password:
  - `passwd -x days login`
- When a user's password has expired, you can set the number of days it can remain expired before disabling the account completely:
  - `passwd -i days login`

# IMPORTANT USER ENVIRONMENT FILES

- `/etc/skel` default template for a newly-added user's home directory
- `/etc/profile` sets environmental variables used by all users
- `/etc/profile.d` contains scripts specific to certain rpms
- `/etc/bashrc` contains global aliases and system settings
- `~/.bashrc` contains user aliases and functions
- `~/.bash_profile` contains user environment settings and can be set to automatically start programs at login

## LAB

1. Create a new group 'dev'. Create a new user 'alice' as a member of the 'dev' group, with a description of "Alice from Dev" and a default shell of '/bin/csh'. Use the `passwd` command to set a password for alice, then log in as alice and verify her access.
2. Set a maximum password lifetime of 4 weeks for the alice account. Look at the `passwd`, `shadow` and `group` files.
3. Configure the users `guido`, `linus`, and `richard`. Set all their passwords to "linux".
4. Make these users part of the `ru` group.
5. Configure the directory `/home/linux` so that each user from the `ru` group can read, create, and modify files.
6. Configure the directory `/home/linux/work` so that each user can create and read files, but only the file's owner can delete.
7. Use ACL's to allow `alice`, not in `ru`, access to the `work` folder.



# PAM

- Applications which are compiled against `libpam.so` may use PAM's modules to customize how individual applications verify their users. Each application has its own configuration file in `/etc/pam.d`
- The first field of the configuration file indicates how the module will be used:
  - **Authentication management (auth)** Establishes the identity of a user.
  - **Account management (account)** Allows or denies access to the account.
  - **Password management (password)** Enforces password management policies.
  - **Session management (session)** Starts, stops, and records each session.

# PAM

- The second field of the configuration file indicates the effect that the module will have on the application:
  - **Required** If this module fails, access will not be granted, but all other modules will still be run.
  - **Requisite** If this module fails, access will not be granted and no other modules will be run.
  - **Sufficient** If this module succeeds, access will be granted and no other modules will be run.
  - **Optional** The result of this module is ignored.

# PAM

- The third field of the configuration file indicates the name of the actual PAM module to be used for the config line in question.
- Side note:
  - The config file `system-auth` is a collection of many PAM modules commonly used by many authentication services. You will see it included by many of the other configuration files. *Do not modify this file directly.*

# PAM

- **pam\_unix**            Authenticates users by UNIX password
- **pam\_securetty**    Only allows root to log in from secure terminals listed in `/etc/securetty`
- **pam\_nologin**       Will not allow any non-root user to login if `/etc/nologin` exists
- **pam\_time**           Can be configured to allow/deny access based on the system time
- Helpful PAM documentation can be found in:
  - `/usr/share/doc/pam-<version>`



# LAB

1. Using PAM, prevent “guido” from being able to login on Virtual Console 2. Guido should still be able to login elsewhere.

Hint: Configure the `pam_access` module.

2. Set up the `pam_time` module to restrict `linus` so he can only login between 8am and 5pm Monday through Friday, and block out all non-root users from logging in midnight to 2am Sundays for a maintenance period.

# NIS

- NIS Servers can be configured to centrally manage system and account information. These servers can share the contents of `/etc/passwd`, `/etc/shadow`, `/etc/group`, and several other files among any number of clients.
- To configure a client, you must install the `ypbind` and `portmap` RPMs, and then you can run `system-config-authentication`.
- This command will make the proper entries in:
  - `/etc/sysconfig/network`
  - `/etc/yp.conf`
  - `/etc/nsswitch.conf`
  - `/etc/pam.d/system-auth`

# LAB

1. Configure your server to authenticate against `server1.example.com`
2. You should then be able to log in to your server as `station#` (where `#` is your station number) with the password: `redhat`
3. Next, configure the `automounter` service to automatically mount the user's home directory from `server1` at login
4. Finally, configure `automounter` to automatically do this for ANY `station#` account

Hint: Search for "Wildcard Key" in `man 5 autofs`

# LDAP

- LDAP Servers can also be configured to centrally manage system and account information. LDAP is much more secure and flexible than a default NIS configuration, and as such is becoming much more popular.
- To configure a client, you must install the `nss-ldap` and `openldap` RPMs, and then you can run `system-config-authentication`.
- This command will make the proper entries in:
  - `/etc/ldap.conf`
  - `/etc/openldap/ldap.conf`
  - `/etc/nsswitch.conf`
  - `/etc/pam.d/system-auth`



```
slideshow.end();
```

# RHCE

## BOOT CAMP

Kernel Features



# IMPORTANT KERNEL DIRECTORIES

- `/boot` contains the `vmlinuz` and `initrd` required to boot the system
- `/usr/src/kernels` directory for kernel sources, RHEL5
- `/proc` virtual file system for seeing “into” the kernel



# /PROC/\*

- The `/proc` folder contains copious amounts of information useful for troubleshooting. Some examples:
  - `/proc/meminfo`    Memory utilization breakdown
  - `/proc/devices`    Mapping major numbers to drivers
  - `/proc/dma`            dma channel assignments
  - `/proc/ioports`    io port assignments
  - See the manpage for `proc` for more information and descriptions

# /PROC/\*

- Also in the `/proc` folder is detailed information on every process on the system.
  - Details on process status, environment, commandline, and more can be obtained
- Read the `proc` manpage - tons of information available through `/proc`

# SYSCTL

- `sysctl`: Get/set kernel parameters
  - `sysctl -w kernel.pid_max=65535`
  - `sysctl -a`
  - `sysctl -w vm.swappiness=100`
- Also, you can view/edit runtime values under `/proc/sys`
- To make changes permanent, edit `/etc/sysctl.conf`

# LAB

1. Configure your server to have an open file limit of 524288 files.
2. Configure your server to refuse any ping requests.
3. Configure your server to forward ipv4 packets.
4. Make all of these changes persistent across reboots.



# MODULAR

- The Linux kernel is modular, allowing functional blocks of software to be added and removed on the fly via the modules mechanism.
- Modules encompass functions such as:
  - Device drivers
  - Kernel features - firewalls, RAID, LVM
  - Filesystems

# LSMOD

- lsmod: Prints all of the currently loaded modules

```
[root@dev1 ~]# lsmod
Module                  Size  Used by
ip6v6                   264608  20
binfmt_misc             14096   1
dm_multipath            21136   0
parport_pc              31724   0
lp                      16576   0
parport                 42252   2 parport_pc,lp
usbcore                 129724   1
ext3                    125968   1
jbd                     61928   1 ext3
raid10                  23808   0
raid456                 119840   0
xor                     10512   1 raid456
raid1                   24064   0
raid0                   10752   0
multipath               11776   0
linear                  9088    0
dm_mirror               23016   0
dm_snapshot             18872   0
dm_mod                  55752   3 dm_multipath,dm_mirror,dm_snapshot
processor               26412   0
fuse                    42160   1
[root@dev1 ~]#
```

# RMMOD

- `rmmmod`: Removes (unloads) a loaded modules
  - Can not unload a module that is a dependency of another module
  - Can not unload in-use modules

# INSMOD

- `insmod`: Loads a module into the kernel.
  - Full pathname required
  - Does not handle dependencies automatically



# MODPROBE

- modprobe: Intelligent module handler
  - Can load/unload modules
  - Automatically handles dependencies
  - Only need to specify name of module, not full path, when loading
- depmod: Rebuilds module dependency lists

# SOFTWARE RAID

- Software RAID can all be configured, monitored, and modified with the mdadm command.
- To create a RAID array, you can run the following command:
  - `mdadm -C <RAID dev> -l <LEVEL> -n <# DISKS> <partitions>`
- To verify the RAID array, use either of the following commands:
  - `mdadm --detail <RAID device>`
  - `cat /proc/mdstat`

# LAB

1. Create a RAID-5 array on your machine consisting of:
  - 4 partitions
  - each 512MiB in size
  - one of which should be reserved for use as a hot spare
2. Format this array with ext3 and mount it with support for user quotas so that it will persist across reboots.

# LVM

- The Logical Volume Manager
  - Abstracts the physical hardware into logical drive spaces which can be dynamically grown/shrunk and span disparate physical devices
  - Simplifies hard drive management as it abstracts away the details of the underlying storage devices.
  - Adds a small amount of overhead to the VFS layer, slightly reducing performance.



# LVM TERMINOLOGY

- **Physical Volume (pv)** A physical volume is simply the partition/RAID device for the LVM space.
- **Physical Extent (pe)** A physical extent is a chunk of disk space. Can be any size, but default to 4M.
- **Volume Group (vg)** A volume group is a collection of physical volumes.
- **Logical Volume (lv)** A logical volume is a grouping of physical extents from your physical volumes. This logical volume is where you can format a file system.

# LVM BASIC IDEA

- To create a space suitable for `mkfs`, three steps must occur:
  - `pvcreate`: Create a physical volume
  - `vgcreate`: Create a volume group on PV
  - `lvcreate`: Create a logical volume on VG
- See also `pvdisplay`, `vgdisplay`, `lvdisplay`

# PVCREATE

- Easiest of the LVM tools:
- `pvcreate /dev/sda4`

# VGCREATE

- In basic form, you need to provide a name:
- `vgcreate VolGroup00 /dev/sda4`
- Note that `/dev/sda4` is actually a physical volume created with `pvcreate` - not just a device



# LVCREATE

- `lvcreate -n myvol -L 10G VolGroup00`
- Creates a new logical volume called `myvol`, 10 gigs in size pulled from the `VolGroup00` Volume Group.

# RESIZING LV'S

- `vgextend <volume group name> <new PV path>`
  - Add a new physical volume to a volume group
- `lvextend {-l <+extents>| -L <+size>} <lv>`
  - Grow a logical volume
  - NOTE: Use the + to give the amount of additional space added, otherwise specify the total desired size to end up with.

# RESIZING LV'S

- `resize2fs <logical volume>`
  - Once the lv has been extended, you will need to extend the file system
  - You can grow the file system while it is mounted, but before shrinking it must first be unmounted.
- `lvresize -r {-l <+extents>| -L <+size>} <lv>`
  - Resizes logical volume **and** filesystem at same time!

# LAB

1. Add logical volume management on top of your raid array.
2. Use half the available space for a logical volume formatted with ext3 and mounted persistently across reboots.
3. Take a snapshot of this logical volume and check the file system for errors.
4. Assuming none are found, reset the counter for days and mounts until a check is forced on the original file system.
5. Copy some data onto the LV, then expand it and the filesystem by 50MB. `fsck`, then re-mount the filesystem and verify it's contents.



# SWAP SPACE

- Swap space allows the kernel to better manage limited system memory by copying segments of memory onto disk
  - Performance gains
  - “Expanded” memory space
- `mkswap`                      Create a new swap space for use by the kernel
- `swapon/swapoff`            Enable/disable a swap area
- `/proc/swaps`                Lists current swap areas

# LAB

1. Add 500MB of swap space to your system using a device.
2. Add 500MB of swap space to your system using a swap file.

```
slideshow.end();
```



# RHCE

## BOOT CAMP

File Sharing Services



# NFS

- The Network File Service, or NFS, is used to share data with other servers.
- For this service to work properly, `portmap` and `nfs-utils` rpms must be installed, and `portmap` and `nfs` must be running.
- The command `rpcinfo` can be run to confirm that these services are running on a remote server:
  - `rpcinfo -p server1`
- The directories to be shared are listed in `/etc/exports`

# /ETC/EXPORTS

- The directories to be shared are listed in `/etc/exports`
- `/etc/exports` should be configured as follows:
  - `<shared directory> <who>(<how>)`
- Note the **lack** of space between the `who` and the parenthesis for `how`. Be very careful about this!
- Example:
  - `/to/be/shared     station*.example.com(rw)`

# EXPORTS NETWORK SPECIFICATIONS

- The host/network to be shared to can be specified in a number of convenient ways:
  - **Host**                      Just a single host ( given by name/ip )
  - **Netgroup**                  NIS netgroup, expressed as `@group`
  - **Wildcards**                  Using the asterisk, match based off hostnames plus wildcards, as `*.example.com`
  - **IP Networks**              Specify with IP/netmask or CIDR notation:  
`192.168.1.0/24`  
  
`192.168.1.0/255.255.255.0`



# EXPORTFS

- **exportfs -r** refreshes the server share list
- **exportfs -a** exports all shares in /etc/exports
- **exportfs -u** un-exports a share name
- **showmount -e server1** shows shares on server1

# NFS PERSISTENCE

- NFS mounts can be made persistent across reboots by adding the following to /etc/fstab:
  - `server1:/share /server1/share nfs defaults 0 0`

# LAB

1. Create a new user.
2. Configure your anonymous NFS user to use this new UID.
3. Grant read/write access to a directory this user owns to our class network, except for server1, who should get read-only access. (Bonus: Does the user name matter?)
4. Mount the NFS share from your neighbour, and add it to their `fstab`.

# VSFTPD

- VSFTPD is the default ftp server
- The primary configuration file is `/etc/vsftpd/vsftpd.conf`
- Provides two levels of user access:
  - **Anonymous:** by default these users are chrooted to `/var/ftp` for security
  - **User:** these users authenticate with a username/password and can download any file they can read and can upload into any directory in which they have write access
- Individual users can be denied by placing their names in:
  - `/etc/vsftpd/ftpusers`



# LAB

1. Configure VSFTPD to only allow the user `richard` to ftp to your server.
2. Make sure that `richard` is chrooted to his home directory upon login.
3. Configure your FTP server to allow anonymous access, chrooted to `/srv`

# SAMBA

- SAMBA is an open source implementation of Windows networking protocols. With SAMBA, it is possible to:
  - Provide file and print services for various Microsoft Windows clients
  - Integrate with a Windows Server domain as a Primary Domain Controller (PDC) or as a Domain Member.
  - Be part of an Active Directory domain.

# SAMBA

- SAMBA provides the following services in Linux:
  - Authentication and authorization of users (Active Directory)
  - File and printer sharing
  - Name resolution
  - Browsing (Wins or NetBios)

# GETTING SAMBA GOING

- Four packages must be installed for SAMBA to work as desired:
  - **samba** provides basic software for sharing files and printers
  - **samba-client** allows server to connect to windows shares (also includes the smbclient command, which functions like a command-line ftp client)
  - **system-config-samba** GUI configuration tool
  - **samba-common** contains samba configuration files



# GETTING SAMBA GOING

- For SAMBA to work properly, the following services must be running:
  - `smbd` (SMB/CIFS Server) for authentication and authorization and file and printer sharing
  - `nmbd` (NetBIOS name server) for resource browsing and possibly as a wins server

# CONFIGURING SAMBA

- The main configuration file for SAMBA is:
  - `/etc/samba/smb.conf`
- This file is **very** well commented and has examples for just about anything that you need to do.
- Once you have made a configuration change, you can test it with the `testparm` command.

# SAMBA USERS

- To have a SAMBA user, that user must first be created in `/etc/passwd`
- The command `smbpasswd -a` can then be used to add a user to `/etc/samba/smbpasswd` for SAMBA authentication.

# SAMBA SHARES

- To see the SAMBA shares a user has access to, you use `smbclient` as follows:
  - `smbclient -L <server> -U <user>%<passwd>`
- To mount a share, you use the UNC path:
  - `mount.cifs //server/share /mount/point -o username=<user>`
- To configure this mount to happen at boot, add the following to `fstab`:
  - `//server/share /mount/point cifs credentials=/etc/samba/pub.cred 0 0`
- (where `/etc/samba/pub.cred` is a file that only root can read which contains usernames and passwords)



# LAB

1. Configure SAMBA to share your `/srv` directory only to one neighbor who must log in with the SAMBA username of `richard`.
2. Make this share read-only for the SAMBA user `guido`.
3. Mount the share from your neighbor. Configure it to mount automatically at boot time. Use a credentials file to store the account information securely.

```
slideshow.end();
```

# RHCE

## BOOT CAMP

Web Services



# APACHE

## CONFIGURATION

- The main apache configuration file is `httpd.conf` and is found in `/etc/httpd/conf/`. This configuration file stores the core configuration of the web server.
- In Apache 2, the `/etc/httpd/conf.d` directory stores configurations that are specific to a particular Apache module. All files in this directory ending in `.conf` will be parsed as a configuration file.



# APACHE CONFIGURATION

- You can find this example Apache VirtualHost definition at the bottom of `httpd.conf`:

```
<VirtualHost _____>
    ServerName name
    ServerAlias alias
    DocumentRoot path
    CustomLog /path/to/access_log combined
    ErrorLog /path/to/error_log
</VirtualHost>
```

- The `NameVirtualHost` directive **must be used** to specify an IP that can host multiple websites.

## LAB

1. Configure two websites on your server. “X” represents your station #.
2. `wwwX.example.com` should be served from `/var/www/html` and should also respond to requests for the short hostname `wwwX`.
3. `vhostX.example.com` should be served from `/home/linus/html` and should also respond to requests for the short hostname `vhostX`.
4. Both should be listening on your primary ip address, but `wwwX.example.com` should be the default site.

# SECURING APACHE

- Apache support access control through `allow` and `deny` directives:
  - `allow from <host|network|ALL>`
  - `deny from <host|network|ALL>`
- These can be applied in the given order:
  - `order allow,deny`      Allows explicitly allowed clients and **denies everyone else**. Anyone matching both the `allow` and `deny` are denied.
  - `order deny,allow`      Denies explicitly denied clients and **allows everyone else**. Anyone matching both the `allow` and `deny` are allowed.

# SECURING APACHE

- These access control directive are applied through a per-Directory or per-File basis.
- The `allow`, `deny` and `order` directives are placed inside one of the following tags:
  - `<Directory>`
  - `<File>`



# LAB

## 1. Reconfigure your two websites such that:

- `wwwX.example.com` is accessible to everyone except for the person sitting to your left.
- `vhostX.example.com` is only accessible to the person sitting to your right.

# SQUID

- Squid is designed to cache internet objects and can act as a proxy server for HTTP, FTP, and many other types of requests.
- Squid is highly flexible and powerful, but for the RHCE exam, you only need to demonstrate the ability to set it up and proxy web services, possibly denying access to a given subnet.
- The configuration file for Squid is

`/etc/squid/squid.conf`

# KEY SQUID SETTINGS

- **http\_port** *3128 by default*
- **visible\_hostname** *the hostname Squid broadcasts*

# KEY SQUID SETTINGS

- Access control in squid is handled via ACL definitions coupled with access definitions, as:

```
acl mynet src 192.168.0.0/255.255.255.0
```

```
acl yournet src 192.168.1.0/255.255.255.0
```

```
http_access allow mynet
```

```
http_access deny yournet
```



# LAB

1. Configure your server to offer Squid proxy service to the person sitting on your right, but not to the person sitting on your left.
2. This service should listen on port 8080.

```
slideshow.end();
```

# RHCE

## BOOT CAMP

Network Services



# OPENSSSH

- OpenSSH is the open source version of SSH, and is used by most UNIX variants for secure remote administration.
- This service is configured in `/etc/ssh/sshd_config`.



# OPENSCH CONFIGURATION

- OpenSSH Configuration Parameters of interest:

- Port
- ListenAddress
- PermitRootLogin
- PubkeyAuthentication
- Subsystem sftp

# XINETD

- xinetd is the extended internet services SUPER daemon. :)
- This service acts as a super daemon by listening on key ports for certain types of requests.
- When a request is received, xinetd starts the appropriate service and then hands off the request so that it can be handled correctly.
- xinetd is configured in `/etc/xinetd.conf`, the services that it controls are configured in `/etc/xinetd.d/`

# LAB

1. Configure your box to allow both the 'root' and 'student' users to login locally, but not over ssh.
2. Configure an anonymous rsync service to share the contents of your /srv directory.

# NTP

- The Network Time Protocol is a very useful and accurate method to keep your system clock synchronized with time servers around the world. This is important because:
  - Timestamps in log files across machine will line up, allowing for proper analysis and comparison
  - Cron jobs run at the right time
  - Knowing the correct time just makes for a happy server



# LAB

1. Enable NTP on your machine, and use `0.pool.ntp.org` and `1.pool.ntp.org`.
2. Use the `ntpq` command to figure out how far off your machine's clock is from true time.

```
slideshow.end();
```

# RHCE

## BOOT CAMP

Extra Labs



# LAB

1. Change your default desktop environment to KDE from GNOME
2. Create a new partition 512MB in size and extend VolGroup00 with it. Add an extra 256MB to LogVol00 and resize the filesystem to include the new space.
3. Disable ssh access from sample1.example.com
4. Disable ssh access for user linus from 1pm-4pm daily.



# LAB

1. Configure Apache to serve the contents of `/home/linux/work` when accessed via `wwwX.example.com`. Disallow access from `sample1.example.com`
2. Configure a new RAID 5 array with three 512MB members, one of which is missing, thereby creating a degraded array. Create a new `ext3` filesystem on the array and mount it to `/r1`. Copy some data to the array, then add the third device back in to the array and let it rebuild to a clean state.

# LAB

1. Make KDM the default window manager.
2. Test the quotas you set up on `/home` for the student account. Note how the system responds as approach cross soft and hard limits.
3. Configure SAMBA to provide access to user's home directories
4. Implement `pam_motd.so` for `sshd`. Set the MOTD to "I love this RHCE class!"

# LAB

1. Boot into single user mode and change the root password to "rhce"
2. Set up an at job to create a file named "at.ran" in the /tmp folder 5 minutes from now. Use date to note the current time, then verify later on that the file was created when expected.
3. Verify the httpd package. Extract the original httpd.conf from the rpm and diff it with the live httpd.conf

```
slideshow.end();
```



# RHCE

## BOOT CAMP

BIND



# CONFIG FILES

- BIND basically has two types of configuration files:
  - BIND configuration file, specific to BIND and its features
  - Database files, or zone files, which contain DNS resource records used to describe all of the DNS information needed in a domain
- Next, a discussion of resource records and the six most common ones: SOA, NS, A, CNAME, PTR, MX

# RESOURCE RECORDS

- A resource record contains the DNS *information* about a domain.
- There are several types of resource records, including address records (A), mail exchangers (MX) and name servers (NS).
- Every domain has at least 2 resource records, an SOA and an NS. But that wouldn't be a very useful domain, so there are usually quite a few more records, defining addresses, mail exchangers, canonical names and more.

## SOA

- Start Of Authority: This resource record defines authority for a zone.

```
domain IN SOA nameserver adminemail (  
    serial  
    refresh  
    retry  
    expire  
    negativettl  
)
```

We'll discuss these later!



# SOA EXAMPLE

```
rackspace.com. IN SOA ns1.rackspace.com. dnsadmin.rackspace.com. (  
    2009123004      ; Serial number  
    3h              ; Refresh interval  
    1h              ; Retry interval  
    1w              ; Expires  
    1h              ; Negative TTL  
)
```

# NS

- Name Server: Defines authoritative nameservers for the zone.

```
zone IN NS nameserver
```

- Example:

```
rackspace.com. IN NS ns1.rackspace.com.
```

# A

- Address: Maps hostnames to IP addresses

*hostname IN A ipaddress*

- Example:

*ns1.rackspace.com. IN A 192.168.1.5*

# CNAME

- Canonical Name: Maps *alias* hostnames to their *canonical* counterparts.

*aliashostname IN CNAME canonicalhostname*

- Example:

*ns.rackspace.com. IN CNAME ns1.rackspace.com.*



# CANONICAL?

- In layman terms, canonical is another way of saying “real”, “absolute” or “official”.
- So a canonical name refers to the “official” name for a host.
- Creating an alias for a host means that you have to decide on the canonical name, which would be some A record.
- When a resolver performs an A query on a CNAME, the nameserver looks up the canonical name to find out the address to return.

# WHY NOT JUST USE A?

- There are basically two reasons to use CNAME records instead of just lots of A records.
  - First, ease of maintenance. If you need 10 names for one machine, defining them with CNAME is easiest if you then need to change the IP address of the machine. Only one change instead of 11.
  - Second, *canonicalization*. Some services, most notably sendmail, will convert all aliases into their canonical names. This simplifies mail configuration.

# PTR

- Pointer: Maps an IP address back to a name, specifically the canonical name.

*ipaddressdomain* **IN PTR** *canonicalhostname*

- Example:

5.1.168.192.in-addr.arpa. IN PTR ns1.rackspace.com.

# PTR

- Remember, there is only one PTR record for a given IP address, and it should always point to the *canonical hostname*.
- Also, as a side note, make sure your mail servers map both directions exactly. This is important for proper authentication:

- mailer.mydomain.com. -> 192.168.1.50

- 50.1.168.192.in-addr.arpa. -> mailer.mydomain.com



# MX

- Mail Exchanger: Defines hosts responsible for incoming email for the named zones.

*zone IN MX preference mailhandler*

- Example:

`rackspace.com. IN MX 10 mail1.rackspace.com.`

# MX RECORDS

- MX records allow for enhanced mail routing functionality.
- When an email is shipped out, the server canonicalizes the delivery address. So, for example, bob@ns.rackspace.com becomes bob@ns1.rackspace.com.
- Then the server looks up the MX records for ns1.rackspace.com, choosing the record with the **lowest** preference and attempting delivery there. If delivery fails, the next lowest is attempted.
- This allows for backup email servers!

# PUTTING IT ALL TOGETHER

db.rackspace.com:

```
$TTL 1h
rackspace.com. IN SOA ns1.rackspace.com. dnsadmin.rackspace.com. (
    2009123004 ; Serial number
    3h         ; Refresh interval
    1h         ; Retry interval
    1w         ; Expires
    1h         ; Negative TTL
)

rackspace.com. IN NS ns1.rackspace.com.
ns1.rackspace.com. IN A 192.168.1.5
maill.rackspace.com. IN A 192.168.1.20
ns.rackspace.com. IN CNAME ns1.rackspace.com.
rackspace.com. IN MX 10 maill.rackspace.com.
```

# TOGETHER...

db.192.168.1:

```
$TTL 1h
1.168.192.in-addr.arpa. IN SOA ns1.rackspace.com. dnsadmin.rackspace.com. (
    2009123004 ; Serial number
    3h         ; Refresh interval
    1h         ; Retry interval
    1w         ; Expires
    1h         ; Negative TTL
)

1.168.192.in-addr.arpa. IN NS ns1.rackspace.com.
5.1.168.192.in-addr.arpa. IN PTR ns1.rackspace.com.
20.1.168.192.in-addr.arpa. IN PTR maill.rackspace.com.
```



# SO WHAT ELSE?

- In addition to the zone files for your domains, you need a couple more zone files to get BIND up and running.
  - Loopback address
  - Root hints

# LOOPBACK ADDRESS?

- Someone has to take responsibility for loopback address requests! It's simple enough. `db.127.0.0:`

```
$TTL 1w
0.0.127.in-addr.arpa. IN SOA ns1.rackspace.com. dnsadmin.rackspace.com. (
    2009123004 ; Serial number
    3h         ; Refresh interval
    1h         ; Retry interval
    1w         ; Expires
    1h         ; Negative TTL
)

0.0.127.in-addr.arpa. IN NS ns1.rackspace.com.
1.0.0.127.in-addr.arpa. IN PTR localhost.
```

# ROOT HINTS!

- The root hints tell the nameserver where those DNS Root Servers are located, so that requests for hosts outside of your authoritative zones can be resolved.
- This one is the simplest to put together. You don't even have to write it!
- Simply ftp the db.cache file from `ftp.rs.internic.net/domain`

# WHEW

- Finally. All of the zone files are put together and ready. Final step? Configuring BIND.
- The config file is generally `/var/named/chroot/etc/named.conf`



# NAMED.CONF GENERAL

```
acl "clients" { 192.168.0.0/24;};  
  
acl "servers" { 192.168.1.2; 192.168.1.3;};  
  
options {  
  
    directory "/var/named";  
  
    forwarders { 192.168.0.254 ;};  
  
    allow-query { clients; };  
  
    allow-transfer { servers; };  
  
};
```

# NAMED.CONF ZONES

```
zone "example.com" IN {  
  
    type master;  
  
    file "example.com.zone";  
  
};  
  
zone "example.com" IN {  
  
    type slave;  
  
    file "slave.example.com.zone";  
  
    masters { 192.168.2.254; };  
  
};
```

# LAB

1. Configure your machine to act as the authoritative nameserver for a “demoX.example.com” domain and a “rhceX.example.com” domain, where X is your station number.
2. For both domains, configure these records:
  - “www”, “mail”, and “ns” should all resolve to your IP address
  - “web” should resolve to “www”
  - “mail” should be listed as the primary MTA for the domain
  - “ns” should be listed as the DNS server for the domain
3. Also configure your machine to respond to reverse DNS lookups, such that your own IP address will resolve to “www.rhceX.example.com”

```
slideshow.end();
```



# RHCE

## BOOT CAMP

Email Services



# SENDMAIL

- What is Sendmail?
- Sendmail is an extremely popular mail transfer agent (MTA) used by default on many UN\*X distributions to handle the receipt and delivery of emails.
- Sendmail has been around a very long time, and still carries some configuration thorns from previous decades
- Namely, using m4 for a configuration “language”

# SENDMAIL

- MUA versus MTA
  - A mail user agent (MUA) is a program that users run to read, reply to, compose, and dispose of email (such as Outlook, Mozilla Mail, Eudora, etc). You can have many different MUA's installed and running on one machine.
  - A mail transfer agent (MTA) is a program that delivers mail and transports it between machines. Usually, there is only one MTA running on a machine at any particular time.
- LDA (Local Delivery Agent)
  - Once the MTA receives a message, it determines if the message is intended for a local or remote recipient. If the message is intended for a remote location, the message is then passed off to the appropriate MTA. If the message is local, it will be passed to the LDA. **The LDA on RedHat is procmail.**

# SENDMAIL CONFIGURATION

- **/etc/mail/sendmail.cf**
  - Main configuration file. This file is parsed at each successful start-up.
- **/etc/mail/sendmail.mc**
  - Configuration changes should always be written here.
- **/etc/mail/local-host-names**
  - This file contains a list of domain names that the server will handle mail for.
- **/etc/aliases**
  - This file specifies redirects for one user to another address or group of addresses.



# ALTERNATIVES!

- Alternatives can be used when many packages provide the same service.
- The executable that the Sendmail init script invokes is really just a symbolic link to another symlink in the `/etc/alternatives` directory.
- For example take a look at `/usr/sbin/sendmail`.
- In order to choose between Sendmail and postfix, we just change the symlink.
- This can be done with the following commands:
  - `alternatives --display mta`
  - `alternatives --config mta`
  - `alternatives --set mta`

# POSTFIX

- Postfix was designed from the ground up to be a replacement for Sendmail.
- The Postfix group has the following goals for their product:
  - It should be more efficient than Sendmail.
  - It should be more secure than Sendmail.
  - It should be easier to administer than Sendmail.
  - It should be 100% Sendmail compatible.
- To accomplish these goals, Postfix is made up of many individual programs which each handle a particular aspect of mail transfer. These programs are managed by a supervisory `master` daemon.

# CONFIGURING POSTFIX

- Postfix's configuration file is `/etc/postfix/main.cf`
- The directives in this file can be changed manually, or `postconf -e` can be run to apply them from the command line. For example, the following are the most common of the changes that can be made:
  - `postconf -e "myorigin = redhat.com"`
  - `postconf -e "mydestination = redhat.com mail.redhat.com"`
  - `postconf -e "my_networks = 192.168.0.0/24, 127.0.0.1"`
  - `postconf -e "inet_interfaces = all"`
- `postconf -n` can then be called to check your configuration for errors

# DOVECOT

- dovecot is the default POP/IMAP server for RHEL 5.
- The configuration file is `/etc/dovecot.conf`
- Usually, the only changes that need to be made are the enabling of the desired protocols.



# MUTT

- `mutt` is a full-featured MUA for your terminal. You can use it to test `pop3s` and `imaps`:
- `mutt -f protocol://server`

# LAB

1. Configure Postfix to receive mail for `stationX.example.com`, and store user mailboxes in Maildir format.
2. When mail is received for `ru@stationX.example.com`, that mail should be forwarded to the users `richard` and `linus`.
3. Configure `dovecot` to serve user Maildirs on both `imap` and `pop3`.
4. Generate a new key and self signed certificate for use with `ssl` encrypted `imaps` and `pop3s` and then enable those protocols.
5. Test your secure mail server with `mutt`.

```
slideshow.end();
```



# RHCE

## BOOT CAMP

Network Security



# TCP WRAPPERS

- TCP Wrappers was originally written to provide host based access control for services which did not already include it.
- It was one of the first “firewalls” of a sort. :)
- Before you can set up tcp\_wrappers on a service, you have to check if the service supports it...

# CHECKING TCP WRAPPER SUPPORT

- Determine which binary the application runs as. Check `init` script or:

```
# which sshd
```

```
/usr/sbin/sshd
```

- Check for `libwrap` support in the binary.
- If you see `libwrap` support in the output, then you can configure access to the service with `tcp_wrappers`.

```
# ldd /usr/sbin/sshd | fgrep wrap
```

```
libwrap.so.0 => /usr/lib/libwrap.so.0 (0x009c5000)
```

# TCP WRAPPER OPERATION

- If an application is compiled with support for `tcp_wrappers`, that application will check connection attempts against the `tcp_wrappers` configuration files:

- `/etc/hosts.allow`

- `/etc/hosts.deny`



# TCP WRAPPER OPERATION

- These files are parsed in the following order:
  - The file `/etc/hosts.allow` is consulted. If the configuration of this file permits the requested connection, the connection is immediately allowed.
  - The file `/etc/hosts.deny` is consulted. If the configuration of this file does not permit the requested connection, the connection is immediately refused.
  - If the connection is not specifically accepted or rejected in either file, the connection is permitted.

# TCP WRAPPER CONFIGURATION

- The basic syntax for these files is:
  - `<daemon>: <client>`
- For example, to disable `ssh` connections from `192.168.2.200`, add this line to `/etc/hosts.deny`:
  - `sshd: 192.168.2.200`

# IPTABLES

- IPTables works at the kernel level, just above the network drivers, to provide several useful features.
  - Extremely powerful and flexible Layer 2 filtering engine.
  - NAT support
  - Port forwarding
  - And many more

# IPTABLES RULE MATCHING

- The IPTables configuration is parsed from top to bottom.
- IPTables will respond based on the first match that it finds.
- If there is no specific match, the chain policy will apply.



# IPTABLES TOOLS

- **iptables:** View/modify current firewall rules
- **iptables-save:** Script to save current firewall rules for use with iptables-restore
- **iptables-restore:** Restores iptables-save format firewall rules - useful to set up firewalls at boot
- Consider iptables init script for save/restore. Config file:

`/etc/sysconfig/iptables`

# IPTABLES RULES

- When creating a new rule, considerations include:
  - What chain should the rule apply to?
  - What is the traffic pattern to look for?
  - What should happen with the traffic?

# IPTABLES CHAINS

- **INPUT**

- This chain is responsible for filtering traffic destined for the local system.

- **OUTPUT**

- This chain is responsible for handling outbound traffic.

- **FORWARD**

- This chain is responsible for controlling traffic routed between different interfaces.

# IPTABLES RULES

- Below are a few examples of possible IPTables match criteria:

- incoming interface            **-i**
- protocol                        **-p**
- source ip address            **-s**
- destination ip address       **-d**
- destination port              **--dport**



# IPTABLES RULES

- Finally, some examples of what to do with traffic when matched:
  - **DROP**            Do not deliver, do not respond
  - **REJECT**        Do not deliver, send reject notice
  - **ACCEPT**        Deliver
  - **LOG**            Just log the packet

# IPTABLES RULES

- So to summarize the syntax:
  - `iptables`
  - What chain should the rule apply to?
    - `-A INPUT`
  - What is the traffic pattern to look for?
    - `-s 192.168.2.100`
  - What should happen with the traffic?
    - `-j REJECT`

# LAB

1. Using `iptables`, configure your mail server to NOT accept connections from the `192.168.1.0/24` network, EXCEPT for the `192.168.1.2` host.

```
slideshow.end();
```



# RHCE

## BOOT CAMP

Extra Labs



# LAB

1. Work with a neighbor to implement slave DNS servers.
2. One of you should slave the demoX.example.com domain and remain master for the rhceX.example.com domain.
3. The other should slave rhceX.example.com and master demoX.example.com.
4. Examine logs on both machines and find the zone transfer messages

```
slideshow.end();
```