

# Advanced Linux System Administration on Red Hat

Kenneth Ingham

September 29, 2009

## 1 Course overview

This class is for people who are familiar with basic Linux administration (i.e., they know users, packages, filesystems, disk layout, network configuration, etc). The class covers information they will need to perform complex administration tasks on Linux.

## 2 Course objectives

Students attending this class will learn:

- logical volume management
- disk quotas
- software RAID
- kernel and filesystem tuning
- network monitoring, troubleshooting, and tuning
- how to build Linux clusters for high availability, load balancing, and high performance.
- how to automate the OS installation process.

## 3 Student background

If you are attending this class, then we assume that

- You know how to verify, install, update, and remove software packages.
- You are familiar with the structure of data on disks, including partitions and filesystems and how to work with them.
- You know how to configure logging.
- You are familiar with the **/proc** filesystem and how to work with it.
- You know how to create and delete users.
- You can write shell scripts.
- You understand how the system boots and how to customize it.
- You know how to configure the network to use static or dynamic IP addresses.

- You can set up network servers and cause them to start at boot time.
- You know how to set up NFS as a client and server.

## 4 Logistics

The class lasts four days. The student computers need to run CentOS 5 or Red Hat Enterprise 5. The OS needs to be pre-installed, and set up with 1GB of unallocated space on the disk and a 1GB /test partition, as well as the CentOS/Red Hat tools needed for this class. We can provide a kickstart config file that will do the proper setup and install the proper packages. The class uses the following software:

- A CD DVD or USB memory to mount; the installation CDs/DVD work fine for this
- Access to Internet NTP servers or a local NTP server
- Bonnie or IOzone filesystem benchmarks (on class web site)
- CentOS 5.x or Red Hat Enterprise 5.x
- DHCP server (on distribution media)
- Kernel source (from Internet or source CDs or DVDs; **not** on the normal distribution media)
- Kernel source from kernel.org; not on distribution media
- LVM tools (on distribution media)
- Linux Virtual Server software (on distribution media)
- NFS tools (on distribution media)
- TFTP client and server (on distribution media)
- The Xen kernel and other tools (on distribution media)
- *gkrellm* (on class web site)
- *iptraf* (if covered in this class)
- *nmap* (on distribution media)
- *ttcp* (on class web site)
- *xosview* (on class web site)
- a DHCP and DNS server for the class
- a second disk or USB drive
- an NIS server on the instructor machine with an account the students will log into (if this class covers NIS)
- bonnie or bonnie++ (on class web site)
- disk quota utilities (on distribution media)
- free space to create a volume
- lmbench (on class web site)
- mdadm (on distribution media)
- static IP addresses (these do not need to be routable)
- the ability to be isolated from any other networks so the students can set up a local DHCP server OR enough crossover cables so every pair of class machines can be connected together

The class computers need at least a 100Mbps Ethernet connecting them; for the clustering, Gigabit Ethernet would be even better. Similarly, for some of the cluster exercises, a few machines with two Ethernet ports and an extra switch would be useful, but not required.

The classroom needs the ability to be isolated from any other networks so the students can set up a local DHCP server OR enough crossover cables so every pair of class machines can be connected together.

DHCP network configuration is useful, but not necessary if the instructor is informed about the network configuration parameters. Internet access is useful, but not required.

The machines should be safe for connection to an internal network.

The class needs a web server for the class web site. The instructor's laptop may be this web server; otherwise the machine provided in the classroom for the instructor is a good choice. This machine obviously will need web server software installed.

## 5 Class outline

1. Introduction (Lecture: 15; Lab: 0)
  - (a) Class Introductions
  - (b) Class Logistics
    - i. Class schedule
    - ii. Breaks
    - iii. Question policy
    - iv. Break room and restroom locations
    - v. Assumptions about your background
  - (c) Typographic conventions
  - (d) What the class covers
2. Logical Volume Management (Lecture: 45; Lab: 75)
  - (a) LVM overview
  - (b) Using LVM
    - i. At install time
    - ii. After OS install
    - iii. *system-config-lvm*
  - (c) Extending logical volumes
    - i. Extending ext3 filesystems
    - ii. Extending Reiser filesystems
  - (d) Adding a disk to a volume group
  - (e) Removing a physical volume
  - (f) Reducing a logical volume
  - (g) Removing volumes
  - (h) Linear versus striping of extents
  - (i) Snapshots

- (j) Other LVM info
  - (k) Summary
  - (l) Lab
3. Software RAID (Lecture: 35; Lab: 80)
- (a) Overview
    - i. RAID levels
    - ii. Linux software RAID features
  - (b) RAID-linear Configuration (using *mdadm*)
  - (c) RAID-1 Configuration (using *mdadm*)
  - (d) RAID-5 configuration (using *mdadm*)
  - (e) Handling drive failure
    - i. Hot swap (as close as it gets with SW RAID)
  - (f) Other RAID info
    - i. Booting and RAID
  - (g) RAID and LVM
  - (h) Software RAID versus hardware RAID
  - (i) Summary
  - (j) Lab
4. Performance monitoring (Lecture: 25; Lab: 25)
- (a) Introduction
  - (b) *ps* and threads
    - i. Example
  - (c) *sar*
  - (d) *free*
  - (e) *vmstat*
  - (f) *iostat*
  - (g) *top*
  - (h) Graphical tools
  - (i) Troubleshooting
  - (j) Summary
  - (k) Lab
5. Kernel tuning (Lecture: 55; Lab: 45)
- (a) Overview
  - (b) Standard tips for improving performance
  - (c) Benchmarks
    - i. Kernel statistics
  - (d) Device driver parameters
  - (e) Systems with very large memory ( $\geq 8\text{GB}$ )
  - (f) Kernel parameters

- i. Memory, buffers, and paging
  - ii. Threads
  - iii. System V IPC
  - iv. POSIX IPC
    - v. Table sizes
    - vi. Process limits
  - (g) The future
  - (h) Summary
  - (i) Lab
- 6. Disk and filesystem tuning (Lecture: 25; Lab: 50)
  - (a) Benchmarks
  - (b) Standard tips for improving performance
  - (c) I/O scheduler
  - (d) Disk I/O parameters
    - i. Read ahead
    - ii. *hdparm*
  - (e) Ext2/3 Filesystem tuning
  - (f) LVM parameters
  - (g) RAID parameters
  - (h) Summary
  - (i) Lab
- 7. Network monitoring, testing, and tuning (Lecture: 45; Lab: 80)
  - (a) Network monitoring and testing tools
    - i. *tcpdump*
    - ii. *wireshark*
    - iii. *telnet*
    - iv. *netcat*
    - v. *netstat*
    - vi. *nmap*
  - (b) Tuning
    - i. Prerequisite assumptions
    - ii. Benchmarking
      - A. Example
    - iii. Parameters
      - A. Maximum Transmission Unit
      - B. Example
      - C. Bandwidth-Delay product
      - D. TCP parameters
      - E. IP fragmentation parameters
      - F. Other kernel parameters

- (c) NFS performance
  - (d) Summary
  - (e) Lab
8. Configuring a DHCP server (Lecture: 25; Lab: 60)
- (a) Introduction
  - (b) Configuration file
    - i. Parameters
      - A. Options
    - ii. Subnet declarations
    - iii. Host declarations
    - iv. Group declarations
    - v. Example ***dhcpcd.conf***
  - (c) Dynamic DNS (DDNS)
  - (d) Testing and troubleshooting
  - (e) Summary
  - (f) Lab
9. Automating installs with Kickstart (Lecture: 60; Lab: 120)
- (a) Introduction
  - (b) Kickstart configuration
    - i. Commands
    - ii. Commands to specify the installation source
    - iii. Commands related to disk partitioning
    - iv. LVM commands
    - v. RAID commands
    - vi. Commands related to network configuration
    - vii. Commands related to the X window system configuration
    - viii. Example commands section
    - ix. Package selection
    - x. Pre- and post-install scripts
    - xi. GUI configuration tool
  - (c) Specifying the kickstart configuration file to use
  - (d) Building an automatic installation distribution
  - (e) Network installations
  - (f) Troubleshooting
  - (g) Summary
  - (h) Lab
10. PXE boot and install (Lecture: 20; Lab: 90)
- (a) Introduction
  - (b) Setup
    - i. The TFTP server

- ii. The DHCP server configuration
  - (c) Troubleshooting
  - (d) Summary
  - (e) Lab
- 11. Virtualization with Xen (Lecture: 60; Lab: 60)
  - (a) Introduction
    - i. Why use a virtual machine (VM)
    - ii. When to not use a VM
  - (b) Xen
  - (c) How Xen treats various parts of the machine
    - i. Xen networking
  - (d) Setting up Xen
    - i. Setting up a VM
  - (e) Virtual machine configuration files
    - i. Configuration file variables
      - A. Variables for both virtualization types
      - B. Variables for paravirtualization
      - C. Variables for full virtualization
  - (f) *xm*
  - (g) Xen notes
  - (h) Accessing data on Xen disk images
  - (i) Xen errors
    - i. Other possible troubles and troubleshooting tips
  - (j) Summary
  - (k) Lab
- 12. Clustering on Red Hat: an Overview (Lecture: 30; Lab: 20)
  - (a) Why cluster
  - (b) Cluster hardware
  - (c) Cluster software
    - i. Conga
    - ii. Conga notes
    - iii. *system-config-cluster*
  - (d) Fencing
  - (e) Distributed storage
    - i. Storage Area Network (SAN)
    - ii. Network-attached Storage (NAS)
  - (f) Summary
  - (g) Lab

13. The Linux Virtual Server (LVS) (Lecture: 25; Lab: 60)
  - (a) Introduction
    - i. Vocabulary/acronyms
  - (b) The director
  - (c) The realservers and the service
  - (d) *ipvsadm*
  - (e) Examples
    - i. LVS-NAT
    - ii. LVS-DR
    - iii. LVS-Tun
  - (f) LVS versus DNS load balancing
  - (g) Troubleshooting
  - (h) Summary
  - (i) Lab
14. Time synchronization with NTP (Lecture: 20; Lab: 15)
  - (a) Introduction
  - (b) How NTP works
  - (c) Setting up NTP as a client
  - (d) Server pools
  - (e) Authentication
  - (f) Troubleshooting
  - (g) Summary
  - (h) Lab
15. SELinux (Lecture: 35; Lab: 30)
  - (a) Overview of SELinux
  - (b) Discretionary versus mandatory access control
  - (c) SELinux vocabulary
  - (d) Security contexts
    - i. Overview
    - ii. Example
    - iii. Processes
    - iv. Files and filesystems
  - (e) Security policies
  - (f) Enabling SELinux
  - (g) Working with file contexts
  - (h) Handling SELinux denials
  - (i) SELinux notes
  - (j) Summary
  - (k) Lab
16. Disk Quotas (Lecture: 15; Lab: 35)

- (a) Overview
- (b) Enabling quotas
  - i. */etc/fstab* changes
  - ii. Remount the filesystem
  - iii. Create the quota database
- (c) Setting up user quotas
- (d) User interaction with the quota system
- (e) Other quota info
- (f) Summary
- (g) Lab

## Appendices

### A. Kernel configuration (Lecture: 35; Lab: 180)

- (a) Introduction
- (b) Ways to configure the kernel
  - i. `make config`
  - ii. `make menuconfig`
  - iii. `make xconfig`
- (c) Configuration options
  - i. Important configuration options
- (d) Other *make* targets
- (e) Kernel modules
- (f) Troubleshooting
- (g) Summary
  - i. Steps for building a kernel
- (h) Lab

### B. Network Configuration (Lecture: 30; Lab: 45)

- (a) Network configuration
  - i. DHCP client configuration
  - ii. Static network configuration
- (b) DNS lookups
  - i. */etc/resolv.conf*
  - ii. *host*
- (c) Virtual network interfaces
- (d) *mii-tool* and *ethtool*
  - i. Examples
- (e) *system-config-network*
- (f) Troubleshooting
- (g) Summary

- (h) Lab
- C. Network services (Lecture: 30; Lab: 45)
  - (a) *xinetd*
  - (b) *ssh*
    - i. Public key authentication
    - ii. Tunneling
  - (c) NFS
    - i. Client
    - ii. Server
  - (d) Automounter
  - (e) Troubleshooting
  - (f) Summary
  - (g) Lab
- D. Useful sysadmin tools (Lecture: 40; Lab: 65)
  - (a) *lsof* and *fuser*
  - (b) *rsync*
    - i. Some *rsync* options
  - (c) *cron*
    - i. The ***crontab*** file
  - (d) Logging and log files
    - i. The syntax of ***syslog.conf***
  - (e) Log file rotation
    - i. *logrotate* directives
  - (f) The ***/proc*** filesystem
  - (g) *sysctl*
  - (h) Summary
  - (i) Lab