



---

## Red Hat Enterprise Performance Tuning

**Duration: 4 Days**    **Course Code: RH442**

---

### Overview:

Red Hat Enterprise System Monitoring and Performance Tuning (RH442) is designed to teach senior Linux system administrators the methodology of performance tuning and capacity planning for Red Hat Enterprise Linux. This course discusses system architecture with an emphasis on understanding the implications of system architecture on system performance, methods for testing the effects of performance adjustments, open source benchmarking utilities, methods for analyzing system and networking performance, and tuning configurations for specific application loads.

---

### Target Audience:

Experienced Linux system administrators responsible for maximizing resource utilization through performance tuning.

---

### Objectives:

- Elements of monitoring and tuning
  - Simple network monitoring (SNMP)
  - Graphical reporting
  - Kernel-level profiling
  - Queuing theory
  - Compensating for physical disk characteristics
  - Reducing disk visit count
  - Processes and the scheduler
  - Kernel timing and process latency
  - Memory addressing and allocation
  - Memory caches
  - Memory reclamation
  - Essential network tuning
- 

### Prerequisites:

- RHCE certification or equivalent experience
  - Students without an RHCE certification are encouraged to check their experience levels by taking a free pre-assessment test at [redhat.com/explore/pre-assessment](https://redhat.com/explore/pre-assessment)
-

## Content:

### Unit 1 - Elements of Monitoring and Tuning

- Performance tuning is...
- Before you begin
- How much is how much?
- Calculating unit conversions
- Monitoring
- Monitoring vs. Profiling
- Whole-system view
- Sample application model
- Sample hardware model
- Before tuning the kernel
- Kernel tunables
- Using system documentation
- Recovering from problems

### Unit 2 - Simple Network Monitoring

- SNMP
- SNMP versions
- Management Information Base
- The MIB hierarchy
- Referring to MIB objects
- Reading a MIB file
- Installing SNMP packages
- Finding MIB objects
- Using SNMP v1 for queries
- Using SNMP v3 for queries
- Configuring the SNMP client
- Enabling the SNMP agent
- Profiling SNMP host access controls
- Configuring SNMP v1 access
- Configuring SNMP v3 access
- Beyond RH442: Extending snmpd

### Unit 3 - Graphical Reporting

- MRTG
- Configuring MRTG
- Allowing access to MRTG
- Creating a dashboard in MRTG
- Ad-hoc utilities
- Installing iostat and sar
- Using iostat and sar
- Configuring sar
- Awk
- Using awk
- Gnuplot
- Using Gnuplot
- Creating a custom script
- Scheduling a custom script

### Unit 4 - Kernel-Level Profiling

- OProfile
- Installing OProfile
- Installing kernel-debuginfo
- Setting up OProfile
- Running a test
- Reviewing results
- OProfile architecture
- SystemTap
- Required packages

### Unit 6 - Compensating for Physical Disk Characteristics

- Physical factors affect disk IO
- Disk storage density
- Choosing a peripheral interconnect
- SCSI bus considerations
- Electro-mechanical positioning
- Block IO requests and cache effect
- Tuning sequential read access
- Tuning the disk queue
- Tuning the deadline scheduler
- Tuning the anticipatory scheduler
- Tuning the noop scheduler
- Tuning the (default) cfq scheduler
- Fine-tuning the cfq scheduler
- Physical block device interfaces
- Virtual block devices
- Tuning virtual block devices
- Logical volumes and VBDs
- Implementing snapshot storage

### Unit 7 - Reducing Disk Visit Count

- The virtual file system (VFS)
- Layout of ext2/ext3
- Fragmentation
- Viewing fragmentation
- Tuning fragmentation
- Filesystem limits
- Journaling
- Improving journal performance
- Tuning journal placement
- Other mount-time options
- Finding lock contention
- Reducing visit count with RAID
- Calculating chunk size
- Calculating filesystem stride
- Tuning round-robin RAID
- Write overhead for RAID5 and RAID6
- Improving RAID1 performance
- Tuning RAID1
- Tuning RAID in SysFS

### Unit 8 - Processes and the Scheduler

- Characterizing process states
- Getting ready to run
- Types of CPU cache
- Locality of reference
- Improving locality of reference
- Multitasking and the run queue
- Preempting the current process
- Sorting the run queue
- SCHED\_OTHER
- Tuning scheduler policy
- Viewing CPU performance data

### Unit 9 - Kernel Timing and Process Latency

- How Linux tracks time
- Tuning system ticks
- Tuning processor speed

### Unit 11 - Memory Caches

- Strategies for using memory
- A closer look at demand paging
- Tuning page allocation
- Tuning overcommit
- Slab cache
- ARP cache
- Tuning ARP cache
- Page cache
- Tuning page cache
- Anonymous pages
- SysV IPC
- Tuning SysV IPC
- Viewing memory with free
- Other commands to view memory usage

### Unit 12 - Memory Reclamation

- Characterizing page status
- Calculating dirty and clean memory
- Reclaiming dirty pages
- Tuning pdflush
- Reclaiming clean pages
- Out-of-memory killer
- Tuning OOM policy
- Detecting memory leaks
- What is swap?
- Improving swap performance
- Tuning swappiness
- Tuning swap size
- Tuning swap for think time
- Tuning swap visit count
- Monitoring memory usage

### Unit 13 - Essential Network Tuning

- Simplified transmit model
- Simplified receive model
- Kernel socket buffers
- Calculating total buffer size
- Calculating per-socket buffer size
- Tuning core buffer size
- Tuning TCP buffer size
- Tuning DMA buffer size
- Is packet fragmentation a problem?
- Tuning fragmentation buffers
- Network interrupt handling
- Improving interrupt handling
- Tuning interrupt handling
- Network sockets
- TCP sockets
- Viewing network sockets
- Tuning TCP socket creation
- Tuning TCP socket keepalive

- SystemTap scripts
- The stap command

#### Unit 5 - Queuing Theory

- Introduction to queuing theory
- Little's Law
- Queue length
- Queue length vs wait time
- Wait time
- A closer look at wait time
- How much user time is needed?
- Profiling time with time
- Finding hot spots in code
- Completion rate
- Arrival rate vs completion rate
- Finding a valid observation period
- Predicting system-wide limits
- Predicting resource limits
- Summary of strategies

- IRQ balancing
- Tuning IRQ affinity
- Equalizing CPU visit count
- Tuning process affinity with taskset
- Tuning run queue length with taskset
- Hot-plugging CPUs
- Scheduler domains
- Configuring the root cpuset
- Configuring a child cpuset
- Important files for scheduler domains
- Virtual CPUs
- Tuning VCPUs at domain creation
- Tuning VCPUs dynamically
- Tuning VCPU affinity

#### Unit 10 - Memory Addressing and Allocation

- Overview of memory addressing
- Virtual address space (32-bit)
- Viewing process address space
- Tuning process address space
- Physical address space
- Mapping virtual addresses (x86)
- Uniform memory architecture (x86)
- Overview of memory allocation
- Improving RAM performance
- Improving MMU performance
- Tuning the NUMA allocator
- Improving TLB performance
- Tuning TLB performance
- Viewing system calls
- Virtual domain memory
- Tuning memory at domain creation
- Tuning domain memory dynamically
- Recovering unassigned memory

---

#### Further Information:

For More information, or to book your course, please call us on 0800/84.009

[info@globalknowledge.be](mailto:info@globalknowledge.be)

[www.globalknowledge.com/en-be/](http://www.globalknowledge.com/en-be/)