

Red Hat World Tour 2004

Technology Road Map





Overview

- Major New Features
 - GCC 3.4
 - SELinux
 - 2.6 Kernel
 - Sistina GFS
 - Resource Management
 - Desktop
- Red Hat Enterprise Linux 4 Overview





Major New Features

- Hardened 2.6 Kernel
- Virtualization and Storage Management
 - Clustered File System
 - LVM
- Security (SELinux, Auditing)
- Class-based Resource Management
- Toolchain
 - GCC 3.4 offers performance improvements over 2.96 and 3.2
 - GDB v6
 - Eclipse
 - Full open-source Java 1.4.1 implementation
- Performance and Stability
- Desktop Improvements





GCC 3.4

- Improved Usability of Profile Feedback and Coverage Testing
 - Large C++ test case on IA32 hardware showed 11% speedup with -O0 and 7.5% speedup with -O2 compilation
- Unit-At-A-Time Compilation Scheme for C, Objective-C, C++ and Java
 - Whole file parsed first and optimized later
 - Schema produced 1.3% improvement for SPECInt2000 on IA32
- C/Objective-C/C++ Improvements
 - Support for precompiled headers
 - Preprocessor file handling rewritten; *#import* and *#pragma once* undeprecated
 - Many C math library functions now recognized as built-ins and optimized
 - Compilers can now handle source files written in any character encoding supported by host C library





GCC 3.4 (cont.)

- C++ Improvements
 - G++ much closer to full conformance to ISO/ANSI C++ standard
- Objective-C Improvements
 - Updated front-end to include bug fixes and enhancements previously only available in Apple's GCC version
- IA32/X86-64 Improvements
 - Tuning for K8 (AMD Opteron/Athlon64) available via `-march=k8` and `-mtune=k8`
 - Support for data structures exceeding 2GB in 64-bit mode
- IA64
 - Tuning code for Itanium2 added via `-mtune=itanium2`
 - 3% SPECInt2000 improvement from addition of DFA processor pipeline descriptors
 - 60% compiler speedup on SPECInt2000 C programs from rewrite of instruction bundling using DFA pipeline hazard recognizer





SELinux

- Red Hat Enterprise Linux 4 and Fedora will both provide SELinux as a component in the base OS
- Provides flexible, comprehensive, full integrated Mandatory Access Control (MAC)
 - Built upon 10 years of NSA operating system research
 - Application of NSA's Flask security architecture
 - Cleanly separates policy from enforcement using well-defined policy interfaces
 - Fine-grained control over kernel services
 - Transparent to applications and users
 - Role-Based Access Control, Type Enforcement





SELinux: DAC Versus MAC


■ Discretionary Access Control (DAC)

- File-level r, w, x triples
- Access control lists (ACL)
- Privileged application can compromise security
- Privileged application's configuration errors can compromise security (e.g. running apache as 'root')


■ Mandatory Access Control (MAC)

- System partitioned into security contexts
- Security policies specify only permitted actions between security contexts (no permission means no execution)
- Superuser account cannot be used to compromise machine (see <http://www.coker.com.au/selinux> for root accounts on public SELinux machines)

Security Rings



Security Contexts





2.6 Kernel Features

- Clinix code has been largely accepted and merged into mainline
 - Support for Hitachi H8/300, NEC v850 and Motorola m68k embedded processors
- New support for NUMA (non-uniform memory access) servers
- Addition of support for NCR's Voyager (support for up to 32 IA32 processors) and NEC's PC-9800 platform
- Improved hyperthreading support
- Scalability Improvements
 - Intel's PAE (physical address extension) allows 32-bit IA32 machines to access up to 64GB of RAM
 - Number of unique users and groups bumped from 32-bit to 64-bit (65,000 to over 4 billion)
 - Number of PIDs before wraparound increased from 32,000 to 1 billion
 - Improved 64-bit block device support allows file systems up to 16TB
 - Device limits increased from 255 major and 255 minor to 4095 major and > 1 million minor





2.6 Kernel Features (cont.)

- Kernel is Preemptible
- Rewrite of I/O Scheduler
- Support for New Futuxes
- Improvements in Module Subsystem
 - Drive modules no use “.ko” suffix
 - Now possible to disable unloading of module
- Better Support for Modern Laptop and Desktop Features
- Improved Support for ACPI (Advanced Configuration and Power Interface)
- Improved Hot-Plug PCI and Power Management
- Support for Multiple AGPs
- Support for USB2.0 Devices
- Bluetooth Device Support





2.6 Kernel Features (cont.)

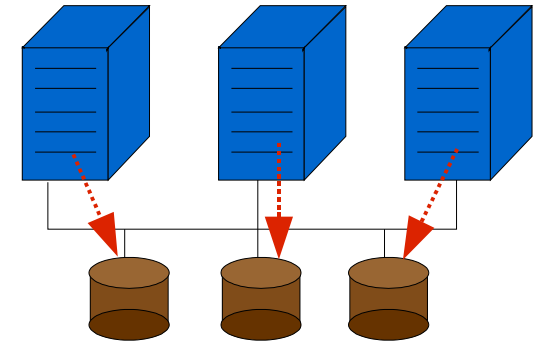
- Rewrite of IDE Subsystem
 - IDE CD/RW devices can now be written to directly
 - SATA (serial-ATA) support added
- Experimental Support for NFSv4 Protocol for Both Client and Server
- NFS Server Implementation Improvements
 - More scalable, supporting 64 times more concurrent users
 - Support for serving over TCP
 - Support for zero-copy networking on supported interfaces
- Support for Native CIFS (Common Internet File System)
- Support for Software-Suspend-To-Disk






Sistina Global File System

- Sistina's Global File System is the first virtualization technology to be delivered under the Open Source Architecture
- Provides concurrent access, shared or served storage, cluster file system
- Exceptional performance and scalability
- Red Hat is currently converting code license to GPL for delivery to the community during 2004
- GFS will be a component of Red Hat Cluster Suite



Standard File System
(Servers mount partitions privately)



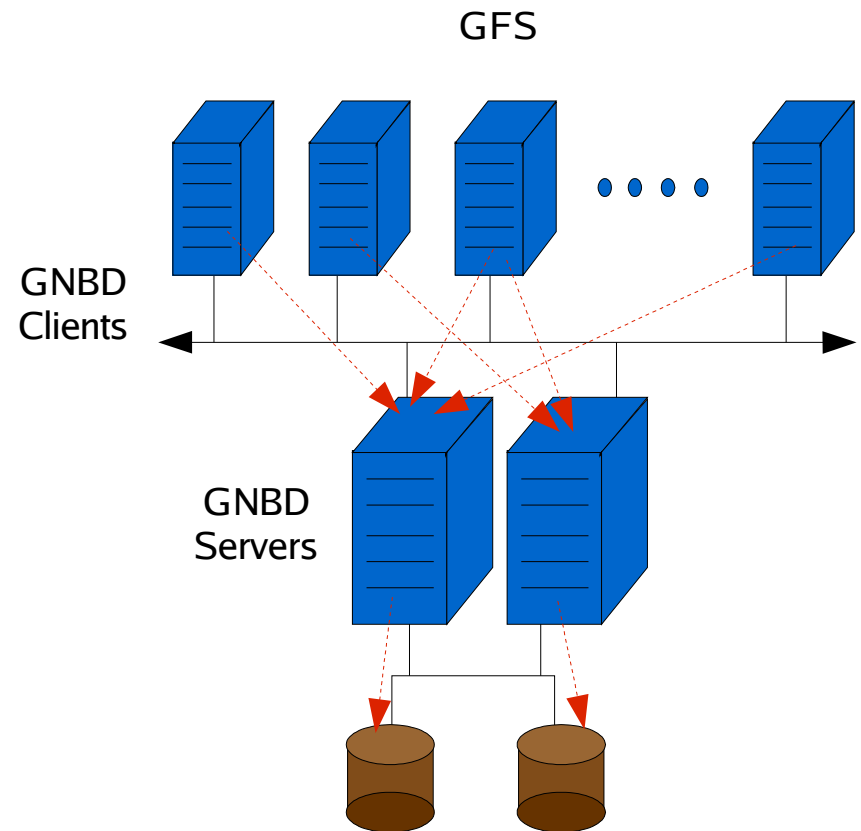
Cluster File System
(Servers mount partitions for concurrent access)





Sistina Global File System (cont.)

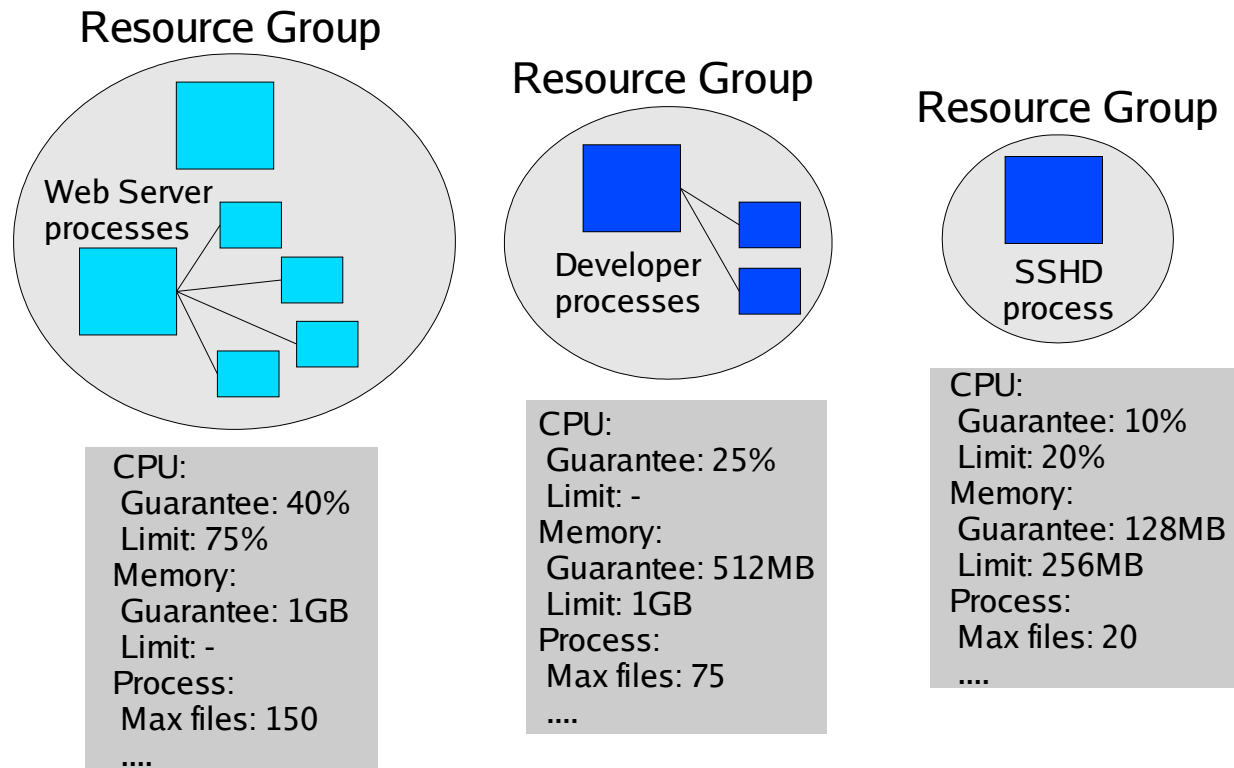
- Features
 - Fully POSIX compliant
 - Multi-node journaling, with recovery by remaining nodes after failure
 - I/O multi-pathing
- Usually configured with fully shared storage (e.g. fibre channel)
- Optionally, using GNBD (Global Network Block Device) it is possible to deploy FS without fully-connected shared storage
 - Reduced cost of deployment
 - Suitable for heavily networked environments such as HPC





Resource Management

- Class-based Kernel Resource Management (CKRM)
 - Provides hierarchical process group definition and resource utilization control
 - Permits applications to execute with a pre-defined Quality of Service (QoS)





Desktop Usability

- Desktop usability is a fundamental requirement for both enterprise and consumer adoption
 - Fit and finish
 - Hardware support
 - Open source Java
 - Productivity applications
 - Profile management and desktop lockdown
 - Hardware and system configuration tools
 - Integrated search
 - Instant messaging
 - Web and media clients





Desktop Interoperability

- Capabilities for working in heterogeneous Red Hat/Microsoft environment
 - Active Directory
 - Exchange
 - Office document format capability
 - ActiveX and Windows Media
- Solutions for customer applications not yet available on linux





Desktop Supporting Services

- Enterprise desktop deployment requires more than just robust client desktop environment
- Additional component of the Open Source Architecture
 - Directory/Identity services
 - Single sign-on integrated into desktop and applications
 - Calendaring, groupware and messaging
 - Remote manageability
 - Thin client interface





Red Hat Enterprise Linux 4

- Continues single code base across all architectures which was introduced in RHEL3
 - Greatly improves code stability and maintainability
 - Support for 7 architectures
 - IA32
 - IA64
 - X86-64 (AMD64 and IA32E)
 - Pseries
 - Iseries
 - Zseries
 - S/390
 - Simplifies ISV application support as only certify once
 - Available in 10 languages (EN, FR, DE, IT, ES, pt_BR, JA, zh_CN, zh_TW, KO)
 - DVD availability

