Red Hat Enterprise Linux version 3
Technical Summary

October, 2003

V4a

Overview

- This presentation provides a summary of the technologies that have been included in the Red Hat Enterprise Linux 3 product family

- Product goals:
  - Continue to improve the high-end enterprise functionality
  - Consolidate Red Hat Enterprise Linux family to provide consistent products and support across all architectures
  - Recognize the rapid maturation of Linux as a practical desktop environment suitable for widespread deployment

Enterprise Linux version 3

- The Red Hat Enterprise Linux 3 product family shipped in October 2003
- Includes a large number of new features
  - Over 100 Priority 1 features; over 350 general enhancements
  - Requests from OEM and ISV partners, and customers
  - Backported features from the Linux 2.5/2.6 kernel trees
- A single source code based is used for all architectures
  - Greatly improves code stability and maintainability
  - 5 new architectures; 64-bit clean implementation
  - Eliminates feature skew; simplifies ISV application support
  - Available in 10 languages (EN, FR, DE, IT, ES, pt_BR, JA, zh_CN, zh_TW, KO)

Enterprise Linux release model

- Extended development of new releases
  - Partners and customers involved in alpha/beta tests delivered through Red Hat Network (RHN)
  - 12-18 month release cycle
- Regular, consolidated updates provided during product lifetime:
  - Maintains compatibility across entire family
  - RHEN plus typically ISO respins
  - Bug fixes
  - Minor enhancements
  - New hardware
  - 5 years

Enterprise Linux 3 features

- Focus on performance, scalability, availability, application development & standards support. Major new features include:
  - Kernel based on 2.4.21 with numerous 2.5/2.6 features
  - Better support for large SMP, memory, and I/O configurations
  - Forward compatibility between RHEL 2.1 and RHEL 3
  - Greatly improved desktop environment
  - 4GB-4GB Kernel/User Memory Split
  - Enhanced standards support
  - Enhanced security features
  - Native Posix Threading Library
  - GCC 3.2 tool chain environment
  - Logical Volume Manager
  - Diskless system support

4GB-4GB Split

- Major new capability to support large physical memories and increased application virtual address space
  - Practical support for very large physical memory configurations
    - Up to 64GB currently in beta
  - Application virtual address space increased ~30% to almost 4GB
    - Enables support for larger user applications
  - This feature is for the X86 architecture only
    - Not required for 64-bit architectures
  - Included in the hugemem kernel

4GB-4GB Split

- A classic 32-bit 4GB virtual address space is split 3GB for user processes and 1GB for the kernel
- The new scheme permits 4GB of virtual address space for the kernel and almost 4GB for each user process

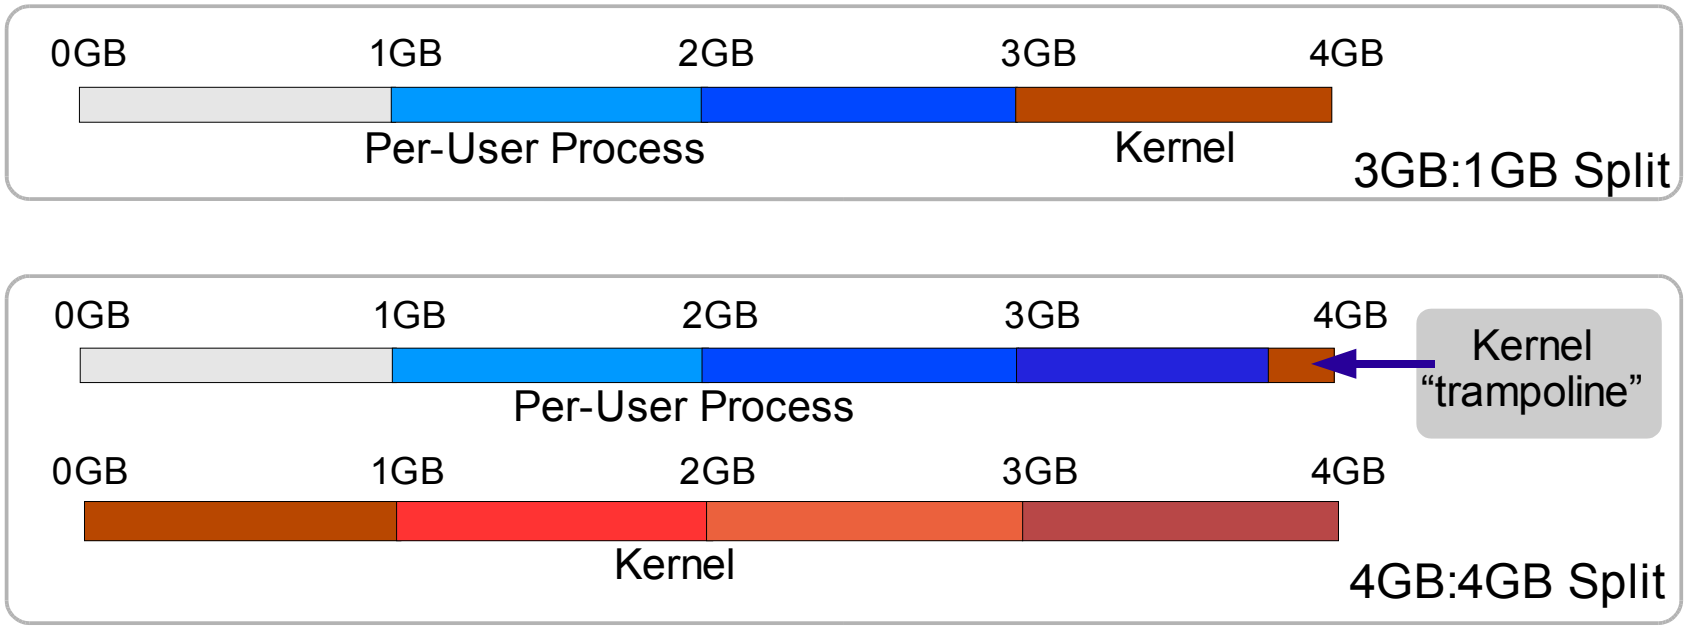

Native Posix Thread Library

- Required for high performance multi-threaded commercial applications, e.g. Java
- Full implementation of POSIX threads
- Major feature that will accelerate Linux adoption in the enterprise
- Highly scalable, native implementation
  - Creation/deletion performance independent of the number of threads running
  - Includes threaded core dumps
  - Informal benchmarks show >50,000 simultaneous thread creations-deletions/second
- Thread Local Storage & Futex APIs

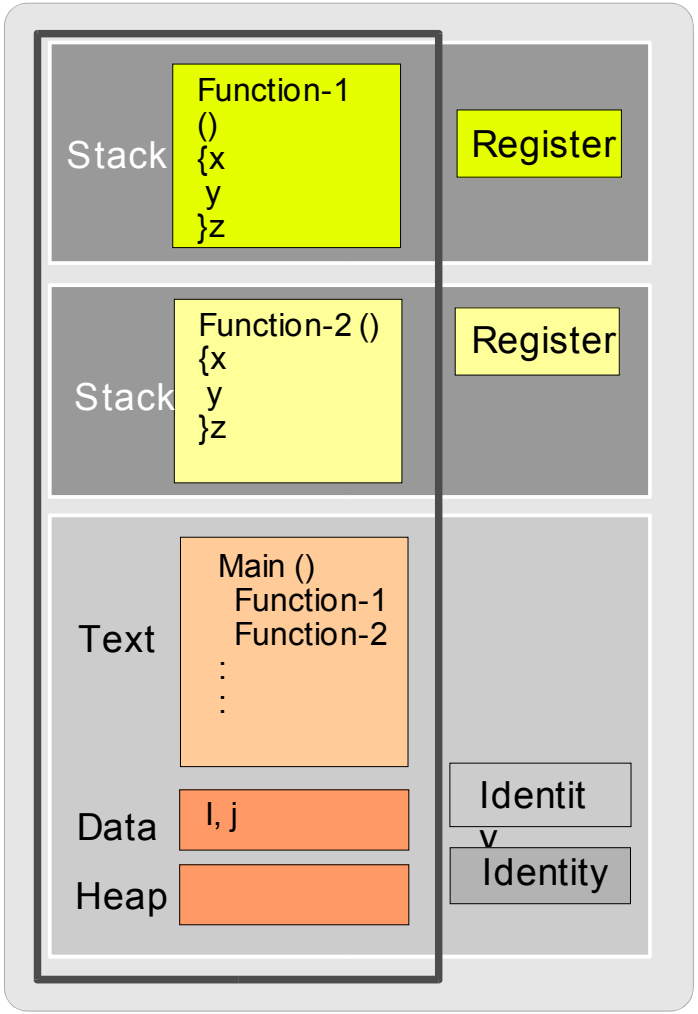

redhat

Compiler Environment

- GCC 3.2 toolchain
  - Full ANSI C++ support
  - ISO C99 Standard support
  - Memory debugging support
  - Architecture optimizations
    - Pentium IV s/w pipelining, etc
    - IA64 instruction scheduler
    - Compiler intrinsics for MMX & SSE (multimedia/streaming instructions)

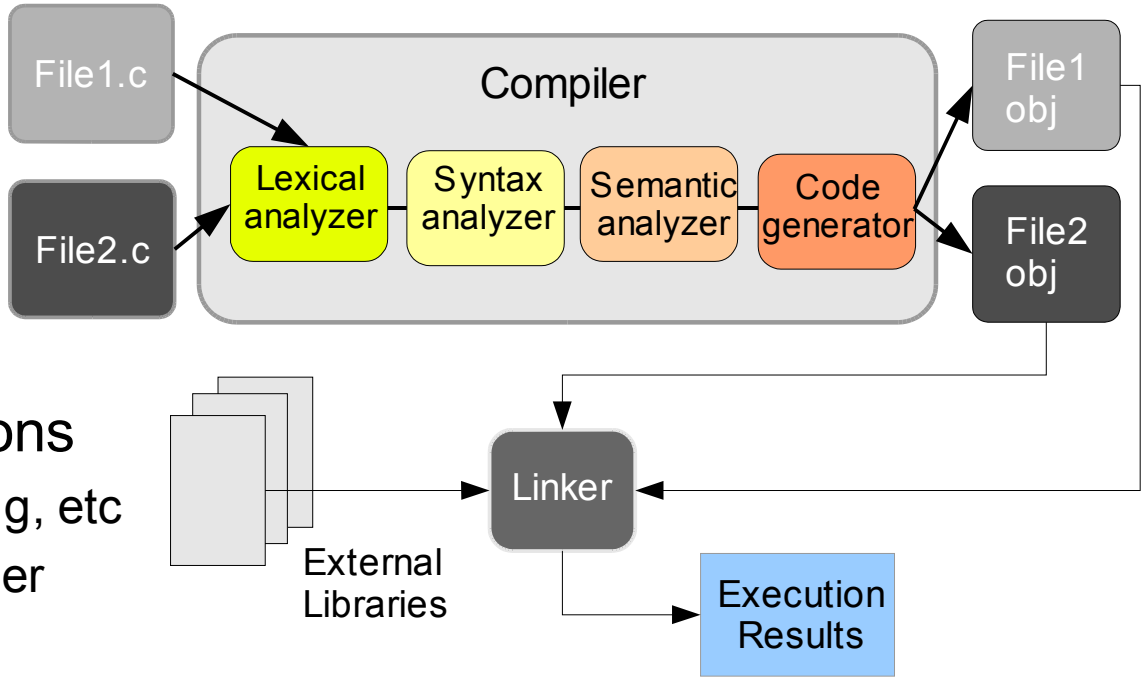

- GCC “ssa” toolchain included as a technology preview
  - Static symbol assignment improves code generation
  - For special purpose apps (Eclipse)

Hyperthreading Scheduler

- Recognizes differences between logical and physical processors
- Optimizes process scheduling to take advantage of shared on-chip cache
- Implements one run queue per physical processor (as opposed to one run queue per processor or per system)
- Support for 16 logical CPUs (or 8 hyperthreaded CPU pairs

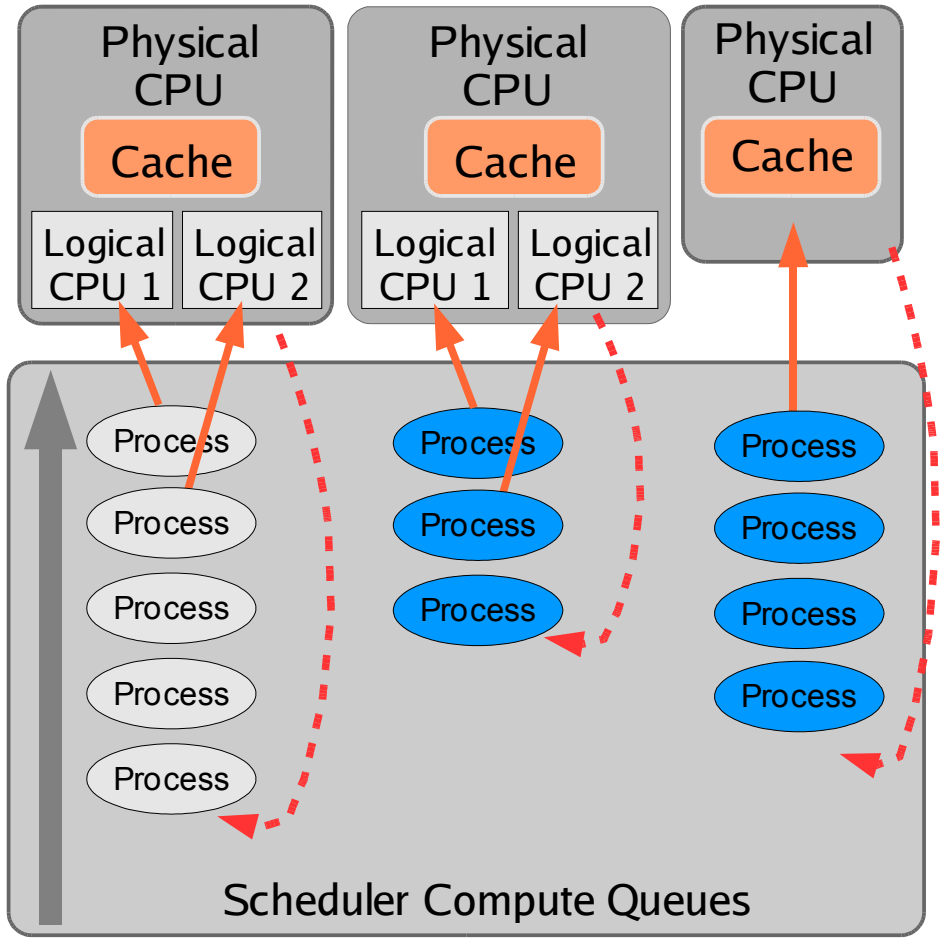

redhat

Logical Volume Manager

- Separate physical and logical devices
- ext2/ext3 filesystems resizable
- Allows flexible storage management
- Compatible with software RAID
- Uses LVM1 implementation (from Sistina)

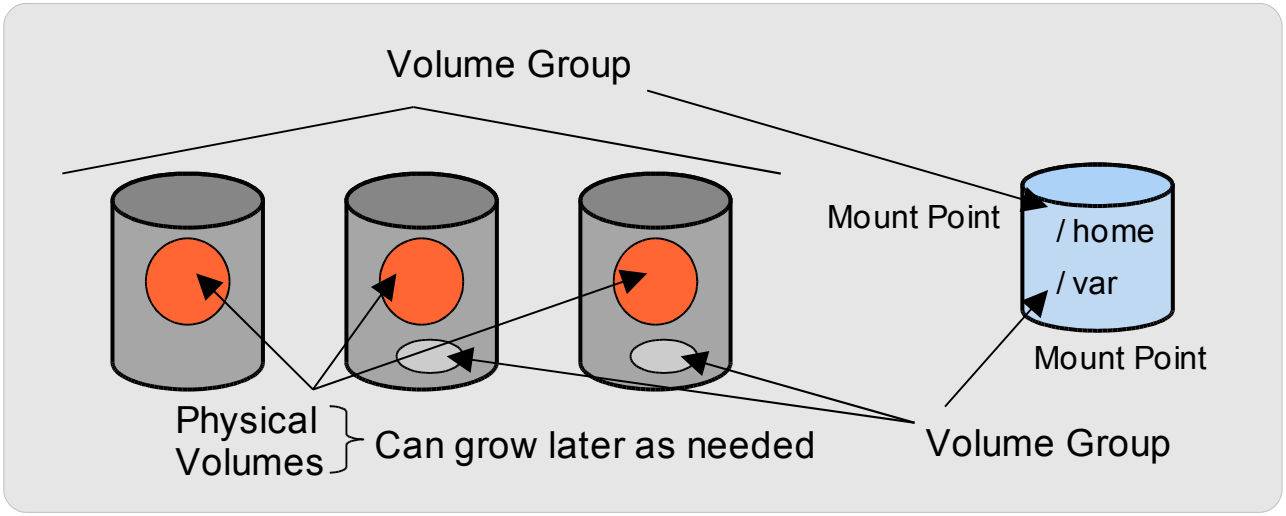

redhat

Standards support

- LSB 1.3 compliance (Linux Standard Base)
  - Standard available at www.linuxbase.org
- NIAP Common Criteria certification expected to be complete by the end of 2003 (National Information Assurance Partnership)
  - Certification to EAL 2 (Evaluation Assurance Level)
  - Internationally accepted standard
  - Specified by US Department of Defense

I/O subsystem improvements

- 64-bit SCSI/Fibre Channel DMA support (improved performance with >4GB memory)
- Up to 256 SCSI devices (permits larger systems to be configured)
- VaryIO support (permits larger I/O transfers)
- Serial ATA support – SATA1 (for Intel PIIX/ICH ATA)

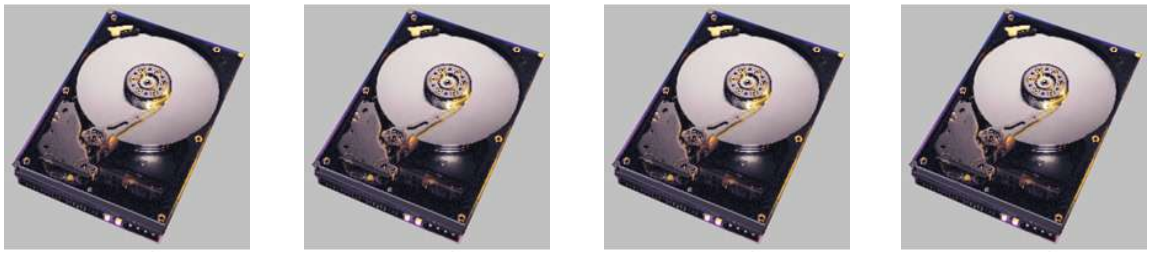

redhat

Desktop environment

- New graphics hardware support
  - XFree86 4.3.x
- Bluecurve graphical user interface
  - Unified GNOME/KDE look and feel
  - Designed for usability
- Bundled productivity applications
  - OpenOffice.org & Ximian Evolution
  - Mozilla 1.4 Web browser

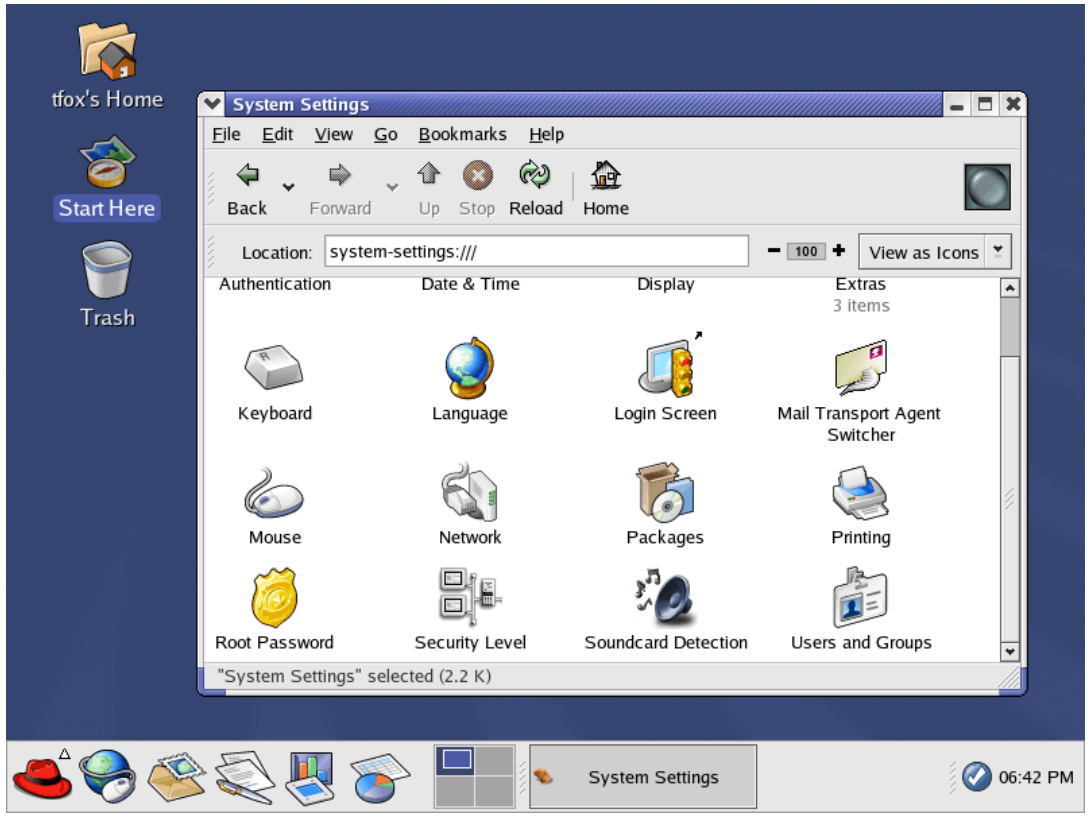

OpenOffice.org

redhat

Serviceability

- Kernel crash dump and analysis enhancements
- Configurable application core dump paths
- Code profiling support included in the kernel – Oprofile
  - System-wide profiler, capable of profiling all kernel/library/application code
  - Uses hardware performance counters in the CPU
  - Includes several post-profiling tools

Diskless system support

- Suitable for HPC, Grid, Telco, and thin-client configurations
- Allows a Red Hat Enterprise Linux server to host other Red Hat Enterprise Linux images, which clients boot over the LAN
- Minimal per-client storage overhead
- Clients can use local disks for swapping and general storage

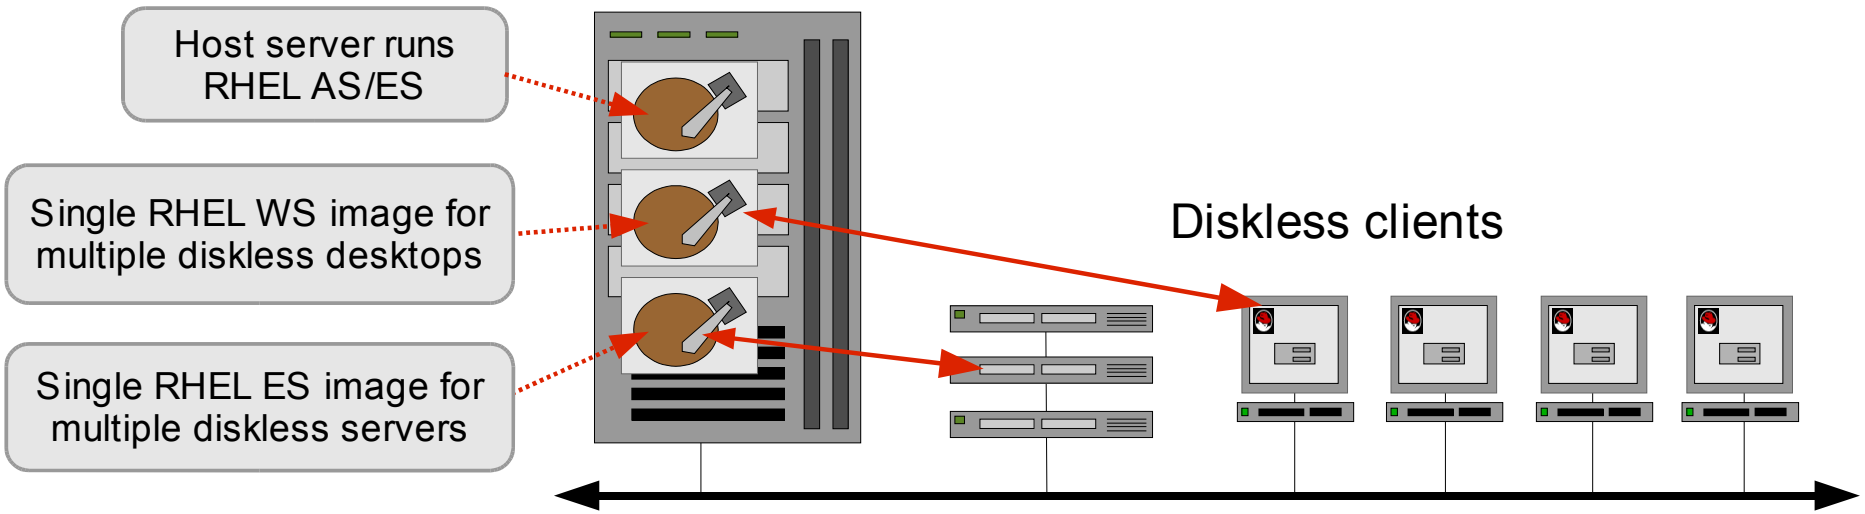

redhat

Java support

- Bundled open source Java environment
  - GCJ/libgcj (Java GCC compiler front-end)
- Third-party Java environments tested and available:
  - IBM Java for x86, IA64, iSeries, pSeries, zSeries, S/390
  - Sun Java for x86 and IA64
  - BEA JRockit for x86 and IA64
- IBM Java provided on 'Extras' CD and RHN channel
- AMD64 supports:
  - x86 Java implementations
  - 64-bit gcj

Networking

- Improvements to channel bonding
  - Failover & bandwidth aggregation for servers w/multiple NICs
- Kernel IPsec – secures IPv4 traffic
  - Tunnel mode builds tunnels between subnets
  - Transport mode secures communication directly between two machines
  - Packets are encrypted, authenticated and anti-replay protected
  - Able to communicate with IPsec devices and OS
- Kernel IPv6 support (more complete implementation than in 2.1
- Kernel support for both IGMP V2 & V3 (Internet group management protocol

NFS

- Significantly improved stability
- Client-side focused performance enhancements
  - NFSv3 readdirplus caches directory information
- Enhanced robustness
- NFS over TCP
- O_Direct support added

Security

- File system ACLs
  - Unix file permissions not always adequate
    - Multiple UIDs, Groups, and set-UID apps proliferate
  - ACLs are additional sets of read/write/execute triplets
  - Can be added to any objects
    - Files, directories, devices, or any other file system objects
  - Highly configurable – fine tune access
    - Without resorting to multiple groups or set-UID apps
  - Includes support for NFS mounted file systems

Security

- Stronghold capabilities included in Enterprise Linux version 3 AS, ES & WS
  - Provides secure web server capabilities
  - Stronghold was previously provided as a separately purchasable product
  - Updated to Apache 2.0 web server
  - With OpenSSL, PHP, mod_perl, etc.

Miscellaneous features

- Red Hat Content Accelerator web cache/accelerator update
- Large Translation Buffer pages – hugetlbfs
  - Large pages (>4KB) to conserve TLB slots
    - Improves performance (esp. database applications)
- Ext3 updates for performance and stability
- SMB file server / client upgraded to Samba 3.0
- IBM x440 Summit integrated into the standard kernel
- Fine-grain process accounting (x86 only)
- Semtimedop – semaphore with time limitation

Miscellaneous features

- Simplified product packaging
  - For each architecture, set of core CDs provide a common Red Hat Enterprise Linux foundation for all family members
  - Each family member has a unique installation CD containing product-specific client/server packages
- Kernel unsupported package
  - Provided as a convenience for non-enterprise drivers and modules
- Memory management enhancements
  - Support for RMAP VM and Large Pages
- ACPI 2.0 (developed with HW partners)
  - Itanium2 and AMD64 only
  - x86 support not planned due to widespread non-compliant BIOS's

Standard x86 kernels

- Kernel-BOOT = Installation only
- Kernel = uniprocessor, highmem disabled, PAE disabled, 4GB/4GB split disabled
- Kernel-smp = multiprocessor (including hyperthreading), highmem enabled, PAE disabled, 4GB/4GB split disabled
  - PAE provides support for >4GB physical memory
    - Slight performance penalty: ~6%
- Kernel-hugemem = multiprocessor, highmem enabled, PAE enabled, 4GB/4GB split enabled
  - For systems with >=16GB physical memory and increased process virtual address space

64-bit landscape

- AMD64 – 32-bit compatibility libraries + hardware execution
- IBM pSeries/iSeries – 32-bit compatibility libraries (ppc64)
- IBM zSeries/s390x
- Itanium2
  - Partial x86 compatibility libraries
  - Intel's execution layer not bundled

Linux kernel 2.5/2.6

- Many Linux kernel 2.5/2.6 features have been backported to 2.4 and included in Red Hat Enterprise Linux 3 to allow earlier enterprise deployments, e.g.:
  - Threads – Per-device locks for block IO – Rmap VM
  - O(1) scheduler – IPv6 – IPsec
- Some features are too intrusive to backport and will be usable in future 2.6-based Red Hat Enterprise Linux products
  - 2.6.0 expected late 2003
  - Stable kernel not expected until sometime in 2004
- Red Hat has not finalized plans for 2.6-based enterprise products

Red Hat Applications

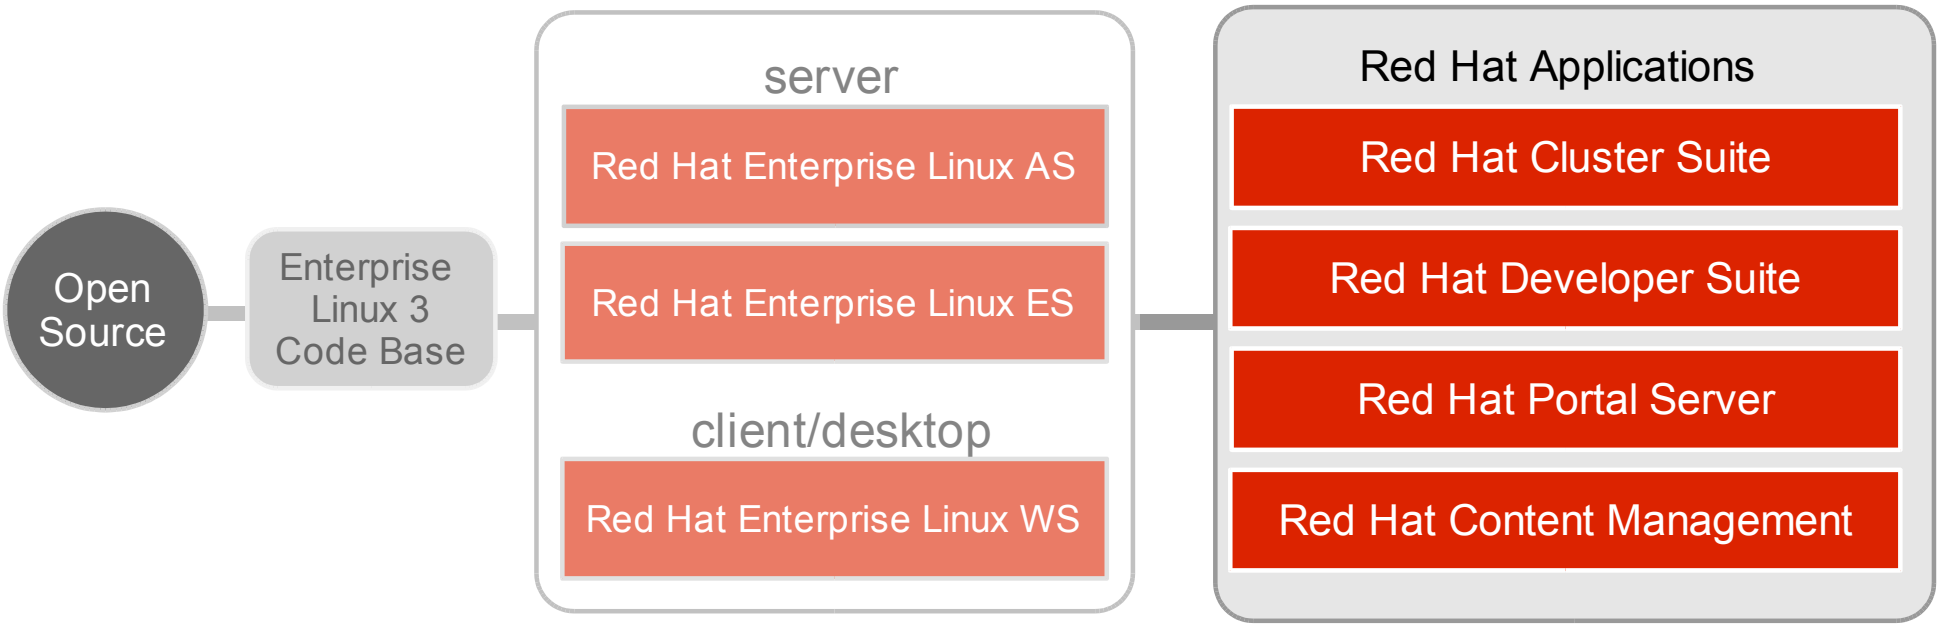

redhat

Red Hat Applications

- Red Hat has extended the value of open source solutions by developing a suite of layered products for Enterprise Linux:
  - Cluster Suite – high availability “failover” clustering
    - Previously available only with Enterprise Linux AS
    - Supported on x86-compatible systems
  - Developer Suite
    - Eclipse-based IDE & developer tools
    - Supported on x86-compatible systems
- Other products are under active development for delivery in 2004

Red Hat Applications

- The new products are delivered with a simple subscription pricing model that includes1 year of maintenance and Red Hat Network access
  - Layered product support level is inherited from the underlying Red Hat Enterprise Linux support level
  - 5 years of maintenance/support
- The new products complement the existing Red Hat Portal Server and Content Management products
- The existing Stronghold product has been retired, and its secure web server capabilities incorporated into Red Hat Enterprise Linux

Red Hat Cluster Suite

- Previously bundled in AS
  - Now available as layered product for AS and ES
- Enhanced to support n-node failover clusters (up to 8 nodes)
- Applications run on any machine
- Shared SCSI or Fibre Channel data
- Allows scalability to large number of applications
- Also includes IP Load Balancing capability (LVS/Piranha)
- Configuration/management GUI

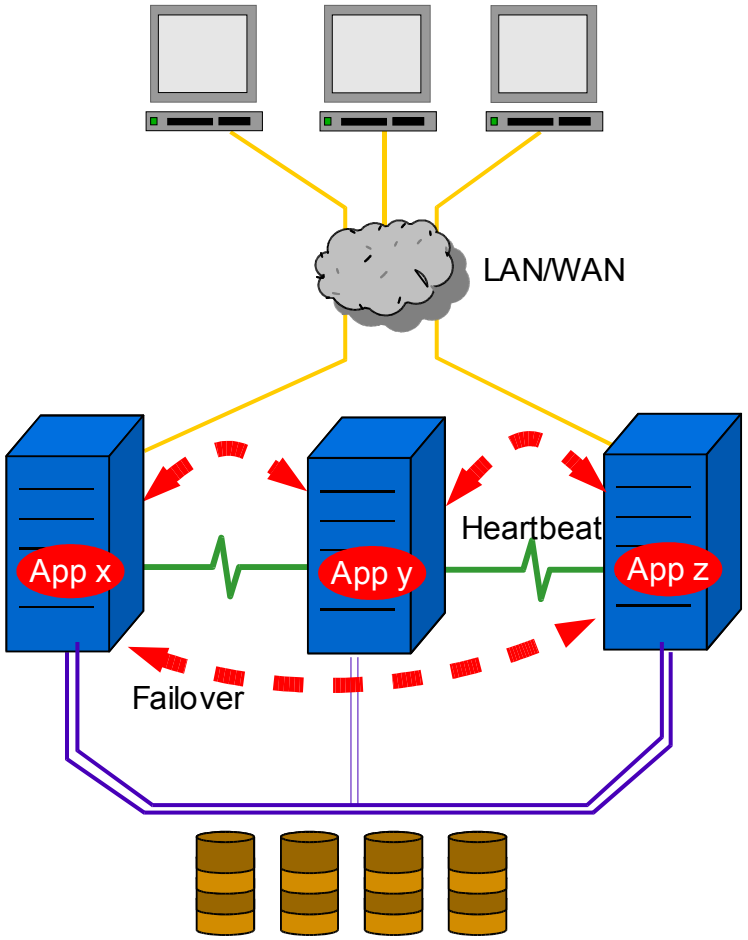

redhat

Red Hat Developer Suite

- Provides a complete development environment for enterprise application developers:
  - Eclipse IDE framework and plugins
- Includes Eclipse plugins
  - C/C++, Java, RPM and profiling
    - Other plugins will be provided in the future

redhat

Questions?

redhat