



Linux Journaling File Systems

Linux For zSeries Journaling File Systems

Volker Sameske (sameske@de.ibm.com)

Linux On zSeries Development
IBM Lab Boeblingen, Germany

Share New York, NY

August 15-20, 2004

Session 9257

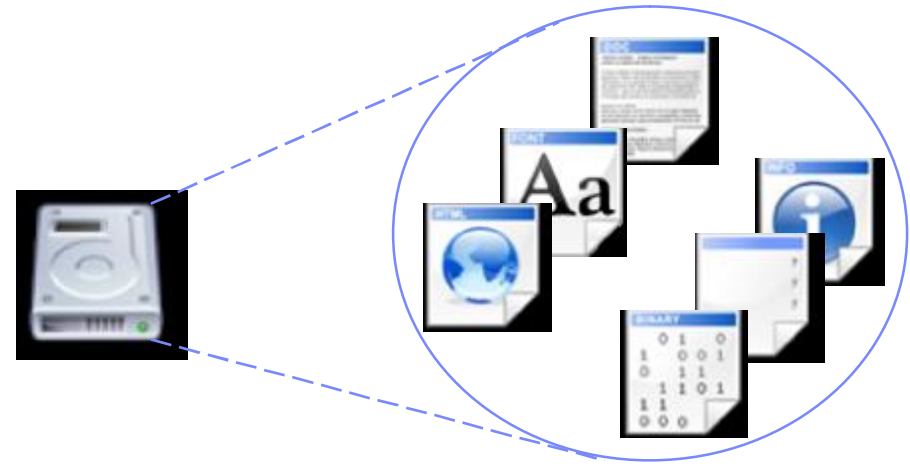


Agenda

- File systems.
 - Overview, Definitions.
 - Reliability, Scalability.
 - Performance.
 - Expand and Shrink.
- Volume Management.
 - LVM, EVMS, MD.
 - Striping.
- Measurement results.
 - Hardware/Software Setup.
 - Throughput.
 - CPU load.



File System - Definition



- Informally
 - The mechanism by which computer files are stored and organized on a storage device.

- More formally,
 - A set of abstract data types that are necessary for the storage, hierarchical organization, manipulation, navigation, access and retrieval of data.


Some File System Terms

- Meta data
 - "Data about the data"
 - File system internal data structure (e.g. inodes)
 - Assures all the data on disk is properly organized and accessible.
- Inode
 - Contains various information about a file
 - E.g. Size or date and time of creation
 - Contains pointers to the data
- Journal
 - On-disk structure containing a log
 - Stores current meta data changes

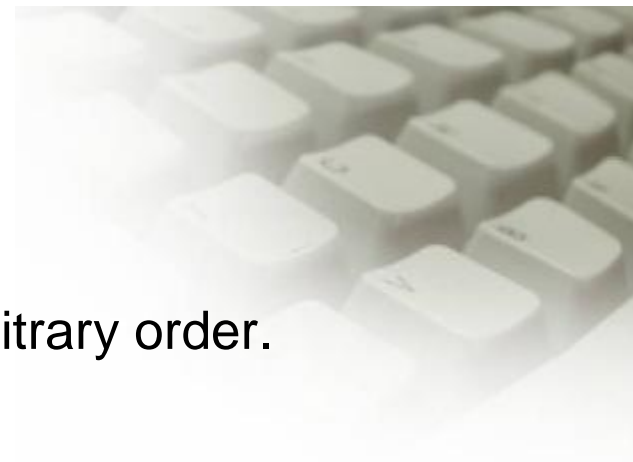


Virtual File System (VFS)

- Software layer in the kernel.
- Provides the file system interface to user space programs.
- Allows coexistence of different file system implementations.
- File system independent.
- All file systems (ext2, ext3, jfs, reiser, ...) provide certain VFS routines.
- Performs standard actions, equal for all file systems.



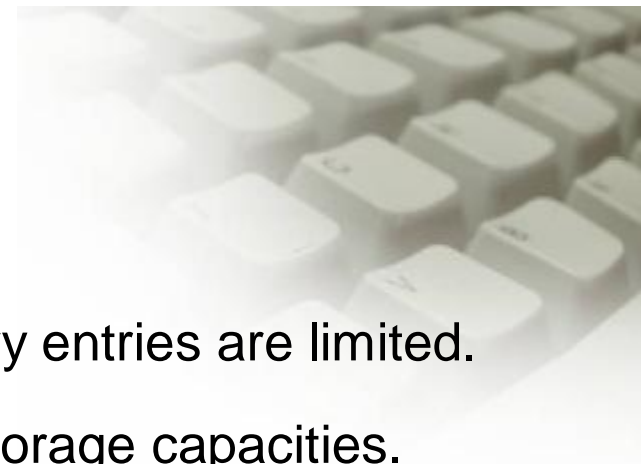
File Systems – Reliability



- Non-journaling file systems.
 - Data and meta-data is written directly and in arbitrary order.
 - No algorithm to ensure data integrity.
 - After crash, complete structure of file system has to be checked to ensure integrity.
 - File system check times depend on size of file system.
 - à Long and costly system outages, risk of data loss.
- Journaling file systems.
 - Data integrity ensured.
 - In case of system crash journal has to be replayed to recover consistent file system structure.
 - File system check time depends on size of journal.
 - à Shorter system outages, no data loss.



File Systems – Performance / Scalability




o Non-journaling file systems.

- File sizes, partition sizes and number of directory entries are limited.
- Inadequate performance when handling huge storage capacities.
- Block size – choose between space waste or performance.
- à Doesn't scale as expected, performance flaw.

o Journaling file systems.

- Large files, larger directories, new organization.
- Scale up easily with thousands of files.
- Good behaviour with small and big files.
- Dynamic i-node allocation (not in ext3 yet).
- à Scales with your file system.

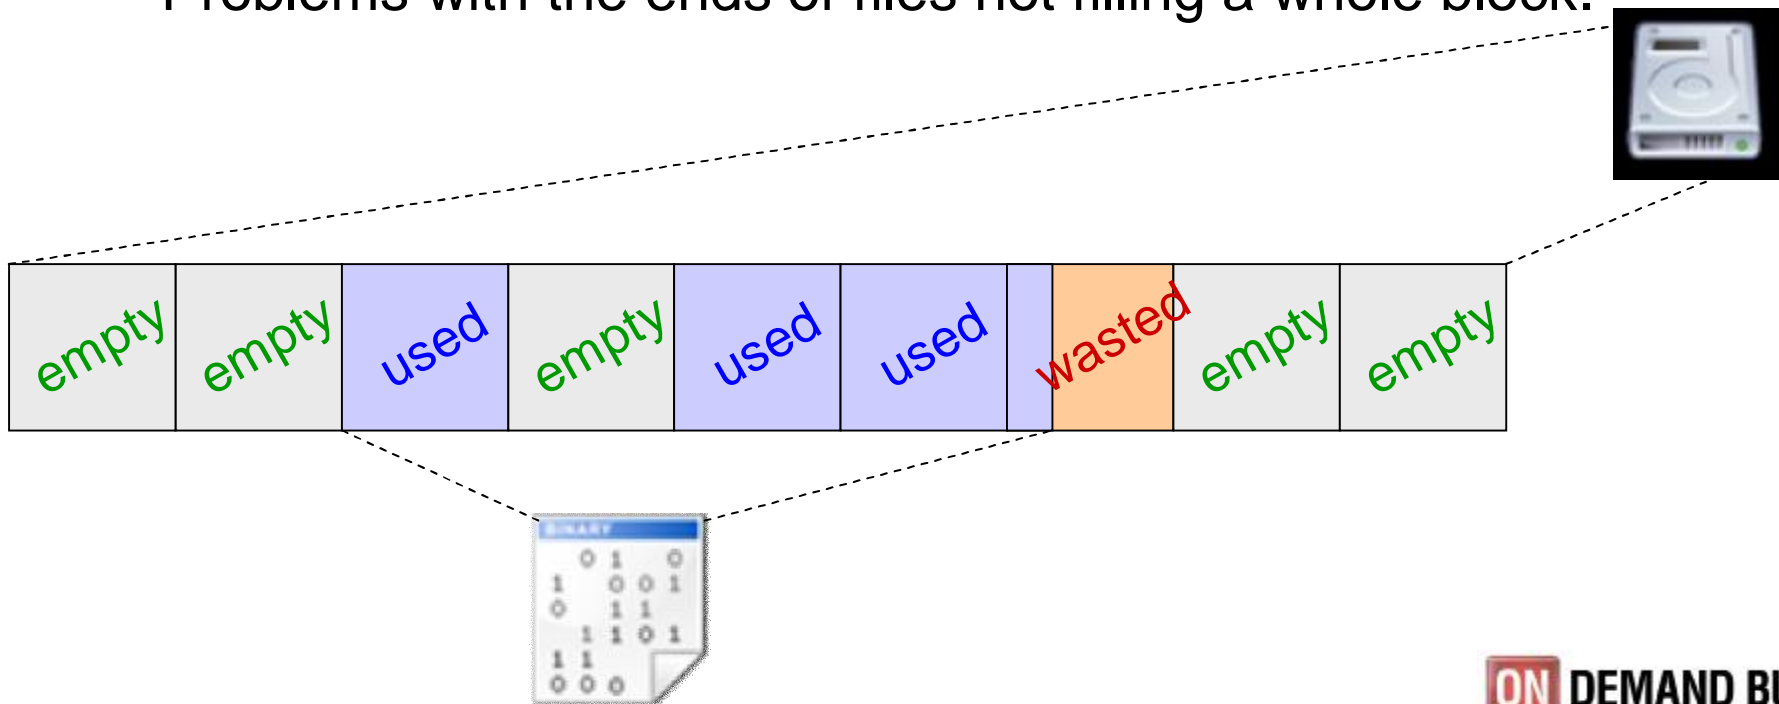
Issues With Journaling File Systems



- Stability.
 - Quite stable, but not so mature as other file systems.
- Migration.
 - Effort to spend.
 - Often one way migration (not ext2 ↔ ext3).
- Performance.
 - Some extra work to do.
 - Room for improvements.
 - à Needs extra resources.

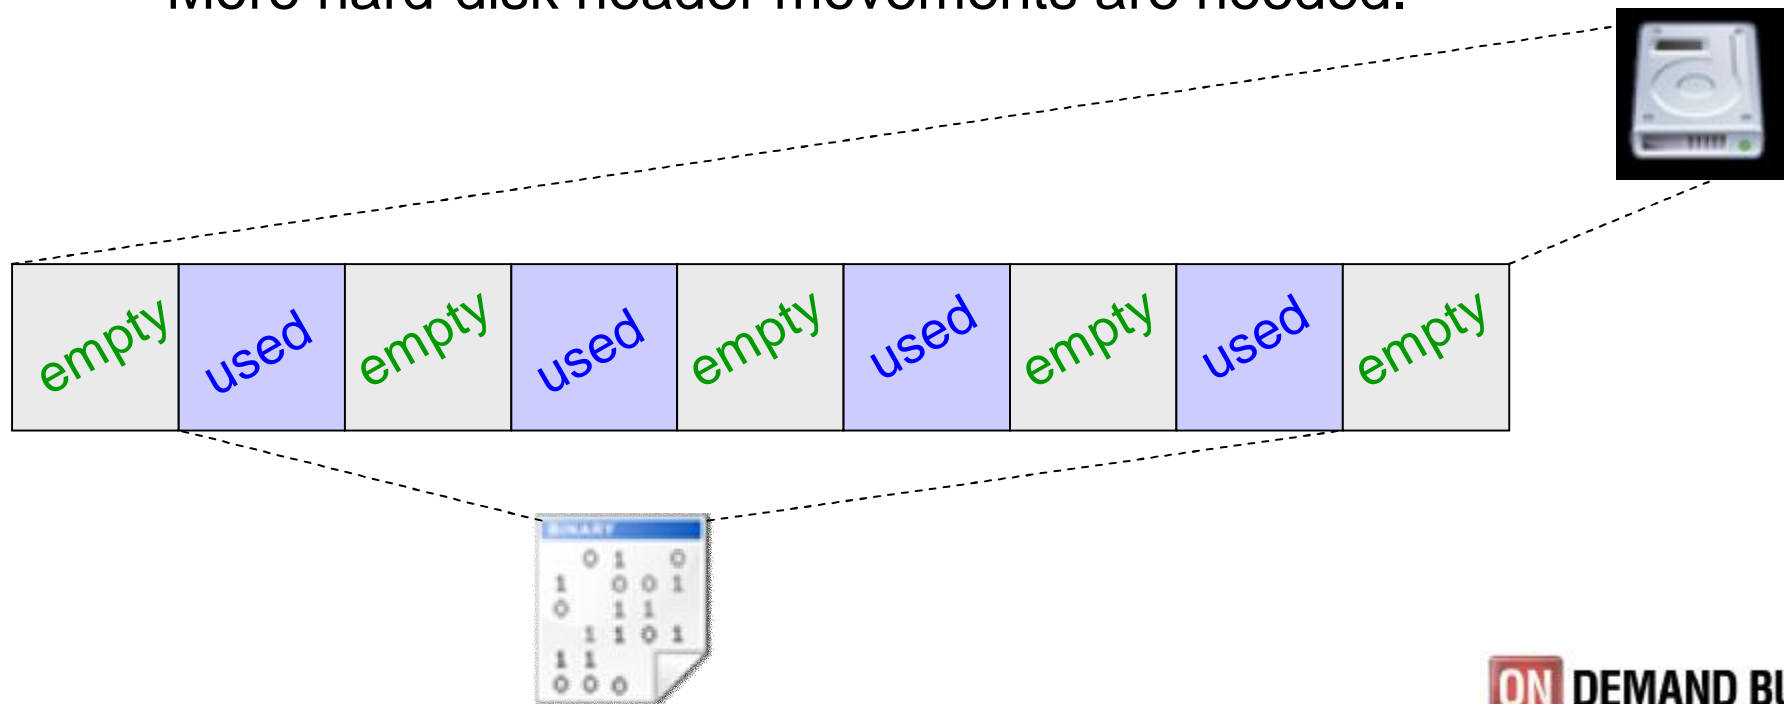
File System Problems

- Internal fragmentation.
 - The logical block is the minimum allocation unit.
 - Average loss of space when blocking data into blocks of fixed size.
 - Problems with the ends of files not filling a whole block.



File System Problems

- External fragmentation.
 - Logical blocks of a file are scattered all over the disk.
 - Operations over that file will be slower.
 - More hard-disk header movements are needed.




Considered File Systems

- Second Extended File System (ext2).
 - Third Extended File System (ext3).
 - Journaling File System (jfs).
 - Reiser File System (reiserfs).
 - Extended File System (xfs).
-
- All file systems have more or less a fairly extensive set of user space tools - for dumping, restoring, repairing, growing, snapshotting, tools for using ACLs and disk quotas, and a number of other utilities.




Second Extended File System (ext2)

- ext2 was most popular Linux file system for years.
- "Old-timer."
- Heavily tested ("rock-solid").
- Faster than other file systems.
- Easy upgradability.
- e2fsck analyzes file system data after system outage.
 - Meta data is brought into a consistent state (lost+found).
 - E2fsck analyzes the entire file system.
 - Recovery takes significantly longer than checking a journal.
 - Recovery time depends on file system size.
- No choice for any server that needs high availability.
- <http://e2fsprogs.sourceforge.net/ext2.html>




Third Extended File System (ext3)

- Developed by Stephen Tweedie and others.
- ext3 is a journaling extension to the ext2 file system.
- Shares ext2 on-disk and meta data format.
- Seamless migration from ext2.
- Large number of existing ext2 systems.
- Supports full data journaling.
- Resizing (only with unmount) possible.
- Easy downgrade from ext3 to ext2 (remount).
- Converting from ext2 to ext3 involves two separate steps:
 - Creating the journal.
 - § `tune2fs -j` creates an ext3 journal with the default parameters.
 - § Other parameters are `size=` and `device=` (man page).
 - Specifying the file system type in `/etc/fstab`.



Third Extended File System (ext3) – Cont.

- 3 modes, specified at mount time or with tune2fs.
- data=journal mode.
 - Maximum security and data integrity.
 - Both meta data and data are journaled.
- data=ordered mode (default).
 - Ensures both data and meta data integrity, but journaling only for meta data.
 - Data blocks are written before meta data update.
- data=writeback.
 - Data blocks are written after meta data update.
 - Performance advantages, but recovery “problems”.
- <http://www.zipworld.com.au/~akpm/linux/ext3/>



Journaling File System (JFS)

- Developed by IBM Austin Lab.
- Ported from OS/2 Warp Server.
- "meta data only" approach.
- Possibility to save data and journal to separate disks.
- Dynamic inode allocation.
- Supports large files and partitions.
- Modifiable journal size.
- Two different directory organizations.
 - For small directories, it allows the directory's content to be stored directly into its inode.
 - For larger directories, it uses B-Trees.
- <http://www.ibm.com/developerworks/oss/jfs/index.html>



Reiser File System (reiserfs)

- Developed by a group around Hans Reiser (Namesys, SUSE).
- Default choice when installing SLES8.
- Only meta data journaling (SLES8).
- Disk space optimization algorithm.
- Online enlargement of file system.
- Three performance tuning options at mount time.
 - no unhashed relocation.
 - hashed relocation.
 - no border allocation.



reiserfs

Reiser File System – Cont.

- Variable journal size and variable journal transactions size.
- Key assets.
 - Disk space utilization.
 - Disk access performance.
 - Fast crash recovery.
- Dynamic allocation of inodes.
- Storage allocation in portions of the exact size.
- File data and meta data are stored next to each other for small files.
- Small files can be stored as part of the meta data.
- <http://www.namesys.com/>



Reiser4 File System – Outlook

- Important parts are rewritten from scratch.
- Atomic file system.
- Based on plugins.
- Squishes small files together.
 - Space efficient.
- Different tree handling.
 - Performance improvements.
 - Dancing trees instead of balanced trees.
- DARPA founded development.
 - Architected for high grade security.
- Is in final testing, and will ship soon.
 - It will take some time until it will become stable as necessary for production.




reiserfs

ON DEMAND BUSINESS™

Extended File System (XFS)

- Originally file system for SGI IRIX OS, early 1990s.
- Good at manipulating large files.
- Only meta data journaling.
- Possibility to use a separate log device.
- Allocation groups.
 - "File systems in a file system."
 - Eight or more linear regions of equal size.
 - Own inode and free disk space management per AG.
 - Allocation groups are independent.
 - Simultaneously addressing.



Extended File System (XFS) – Cont.

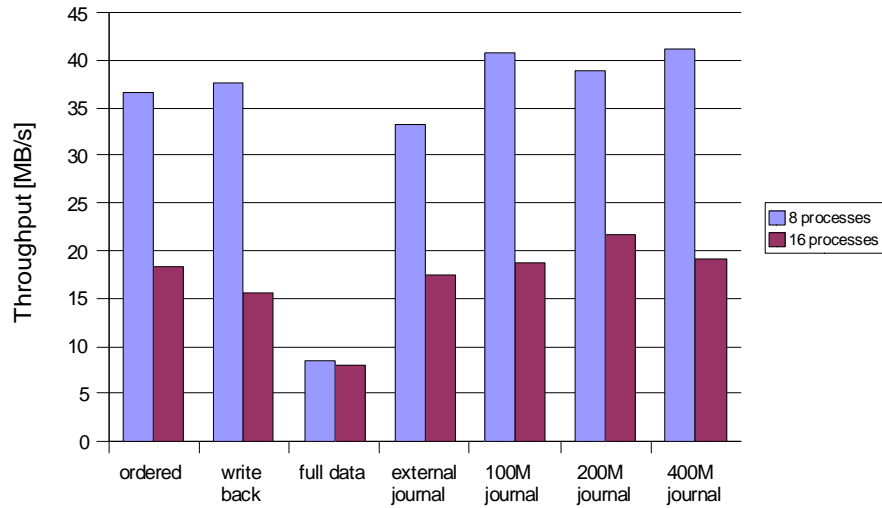
- Use of Btrees for free space and inodes inside allocation groups.
- "Delayed allocation".
 - Pending transaction is stored in RAM.
 - Delayed decision, where to store data.
 - Delayed until the last possible moment.
 - Increased performance.
 - Data loss after a crash during a write is more severe.
- Preallocation to avoid file system fragmentation.
- IRIX XFS disks could be used with Linux.
- <http://oss.sgi.com/projects/xfs/index.html>



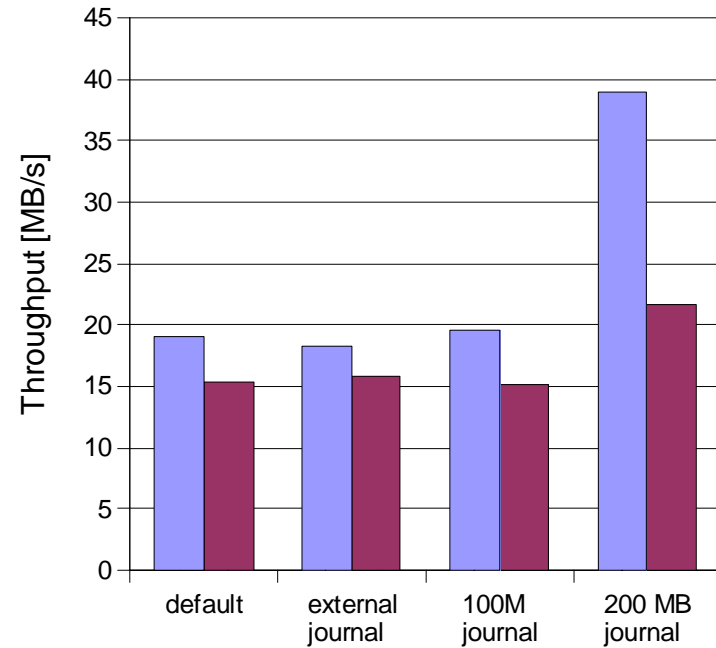
xfs

File system Options

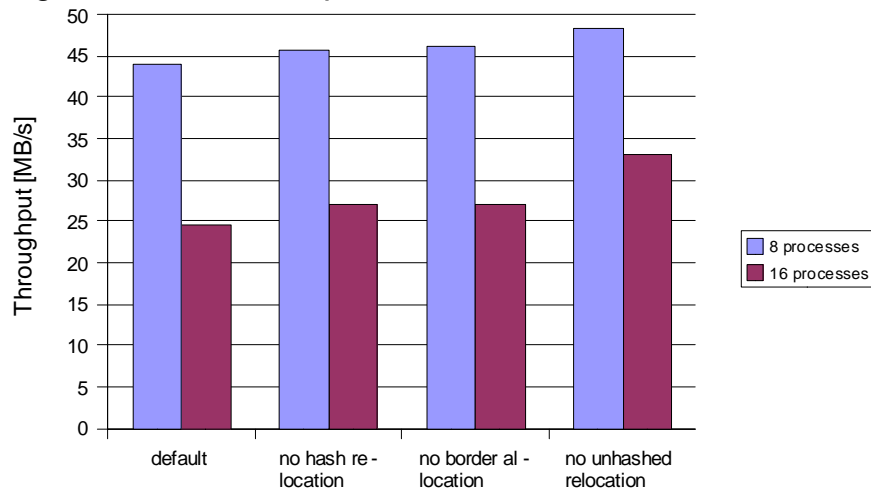
single disk, ext3 optimizations, z/VM, 31bit, 4 CPUs



single disk, jfs optimizations, z/VM, 31bit, 4 CPUs



single disk, reiserfs optimizations, z/VM, 31bit, 4 CPUs



Expand And Shrink

- Expanding and shrinking volumes are common volume operations on most systems.

File System	Shrinks	Expands
ext2	Offline only	Offline only
ext3	Offline only	Offline only
JFS	No	Online only
ReiserFS	Offline only	Offline and online
XFS	No	Online only

Unique File System Features



- ext2:
 - No journaling, but fast.
- ext3:
 - Small, simple, mature, robust, forward/backward compatibility.
 - Data journaling.
- JFS.
 - Inode meta data can contain subdirectory names, symlinks, EAs, etc. if there is space.
 - Two different directory organizations.
- Reiserfs.
 - Space-efficient for small files (tail merging).
- XFS.
 - Designed to support large MP systems.
 - Allocation groups.

Measurement Setup



- SUSE SLES9 RC5
- Dbench 2.1



- **2084-316 (z990)**

- 0.83ns (1200MHz)
- 2 * 32 MB L2 Cache (shared)
- 96 GB
- 4 FCP channels

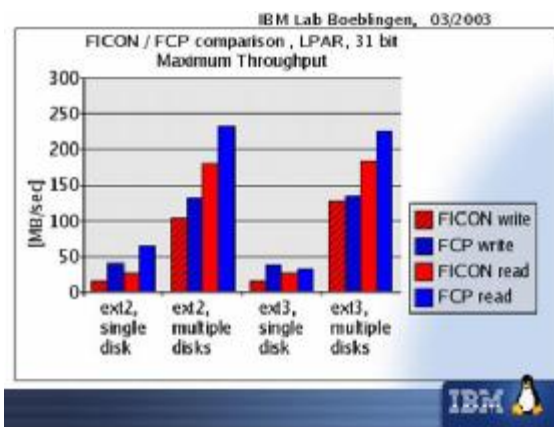
- **2105-F20 (ESS)**

- 384 MB NVS
- 16 GB Cache
- 128 * 36 GB disks
- 10.000 RPM
- FCP



Measurement Setup: Dbench File I/O

- File system benchmark.
- Version 2.1.
- Generates load patterns similar to Netbench.
- It does no networking calls.
- Does not require a lab of load generators to run.
- De-facto standard for generating load on the Linux VFS.



- Author: Andrew Tridgell.
- Released under the GNU Public License.

Things To Keep In Mind




- Dbench is only one benchmark.
 - Dbench throughput is a „mixed“ throughput.
 - We have used only one special hardware/software setup.
There are many possibilities.
 - Disk access under VM is slightly slower.
 - Volume managers with striping are faster than single disks.
 - Lots of different file system options.
- à Many, many switches to tweak.

Volume Management And Multipathing

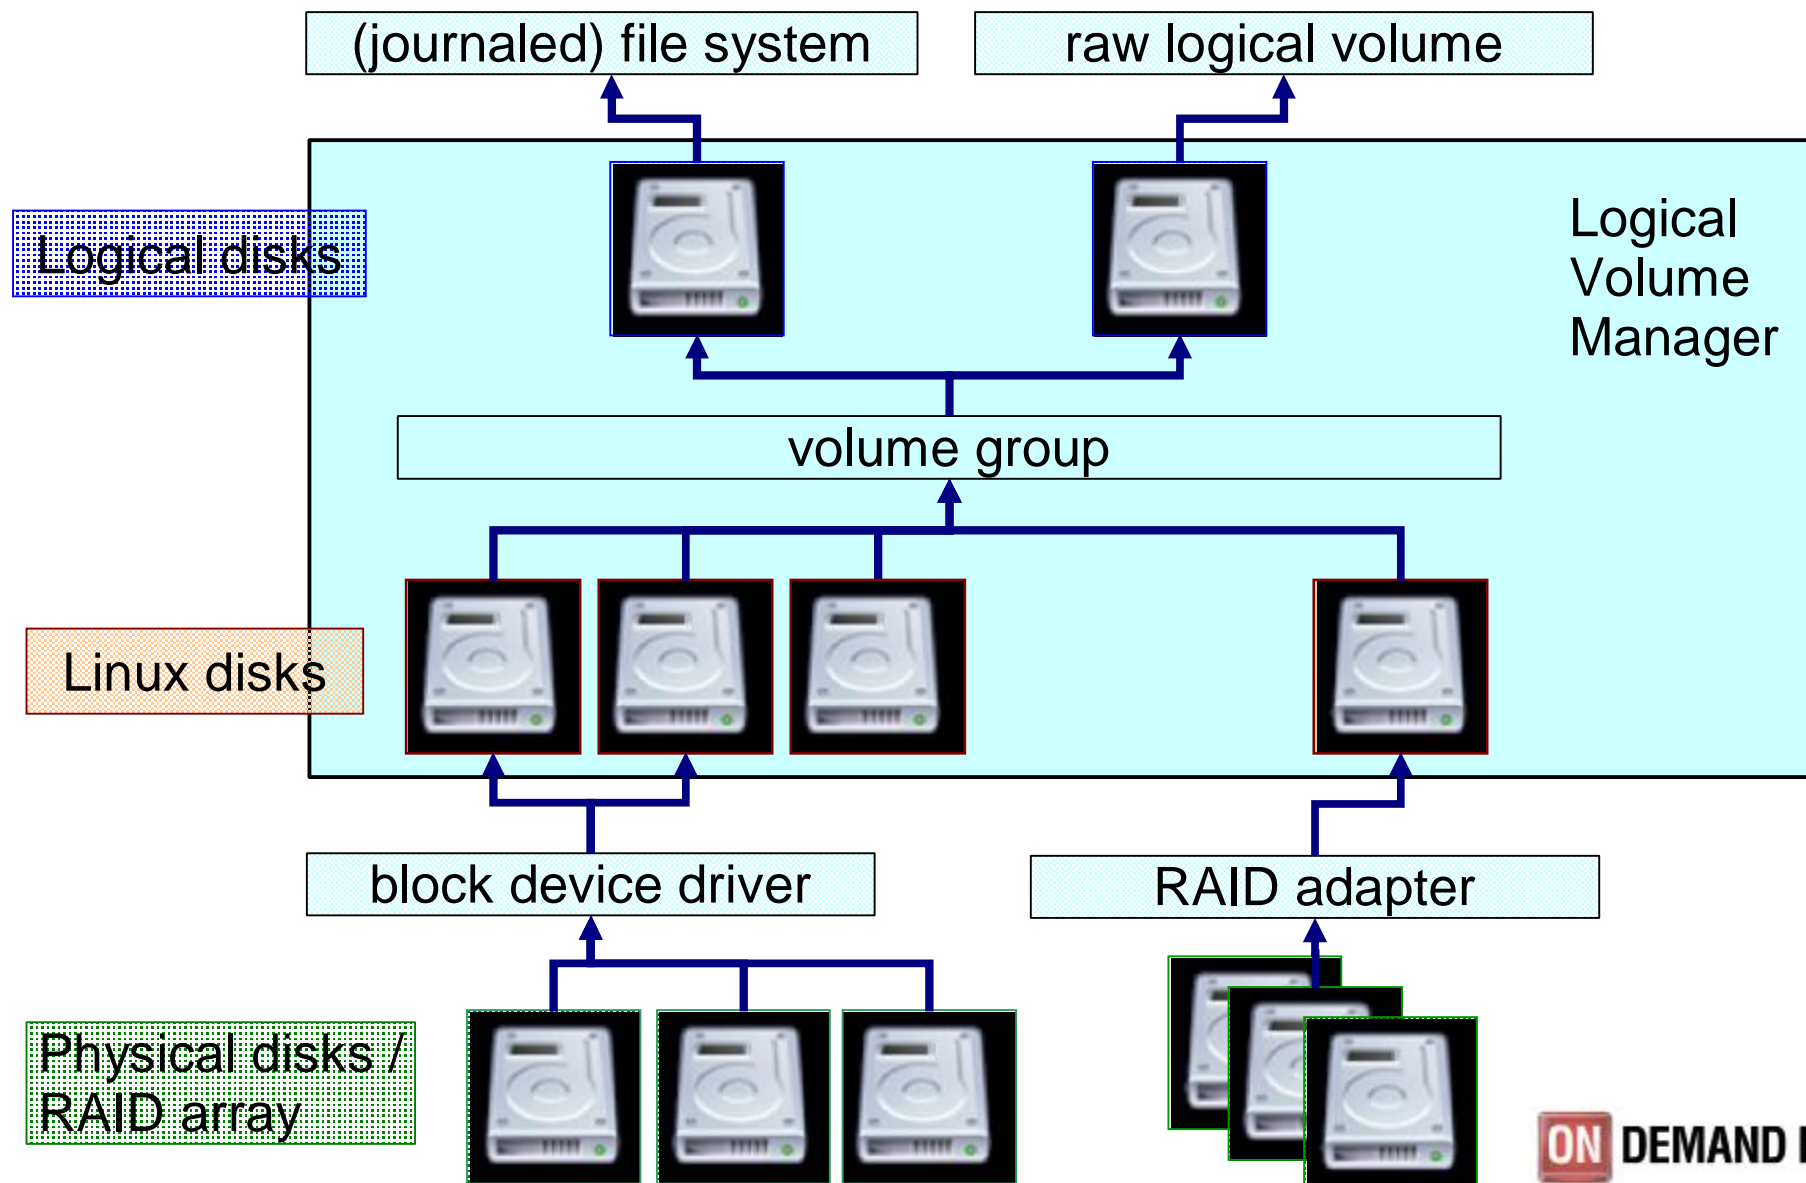
- SLES8
 - LVM – Logical Volume Manager

- SLES9
 - Device Mapper subsystem in 2.6 kernel
 - EVMS – Enterprise Volume Management System
 - LVM – Logical Volume Manager

- RHEL3
 - MD
 - mdadm – multiple device administration



LVM System Structure

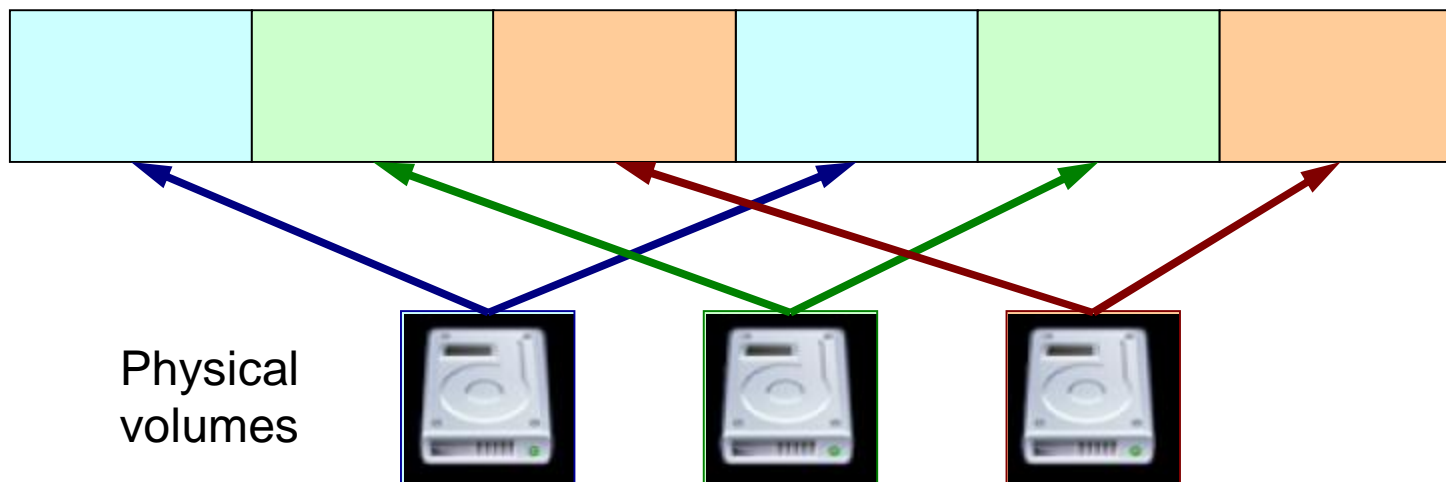


ON DEMAND BUSINESS™

Improving Disk Performance With Striping.


- Technique for spreading data over multiple disks.
- With LVM **and** striping parallelism is achieved.

Striped data stream




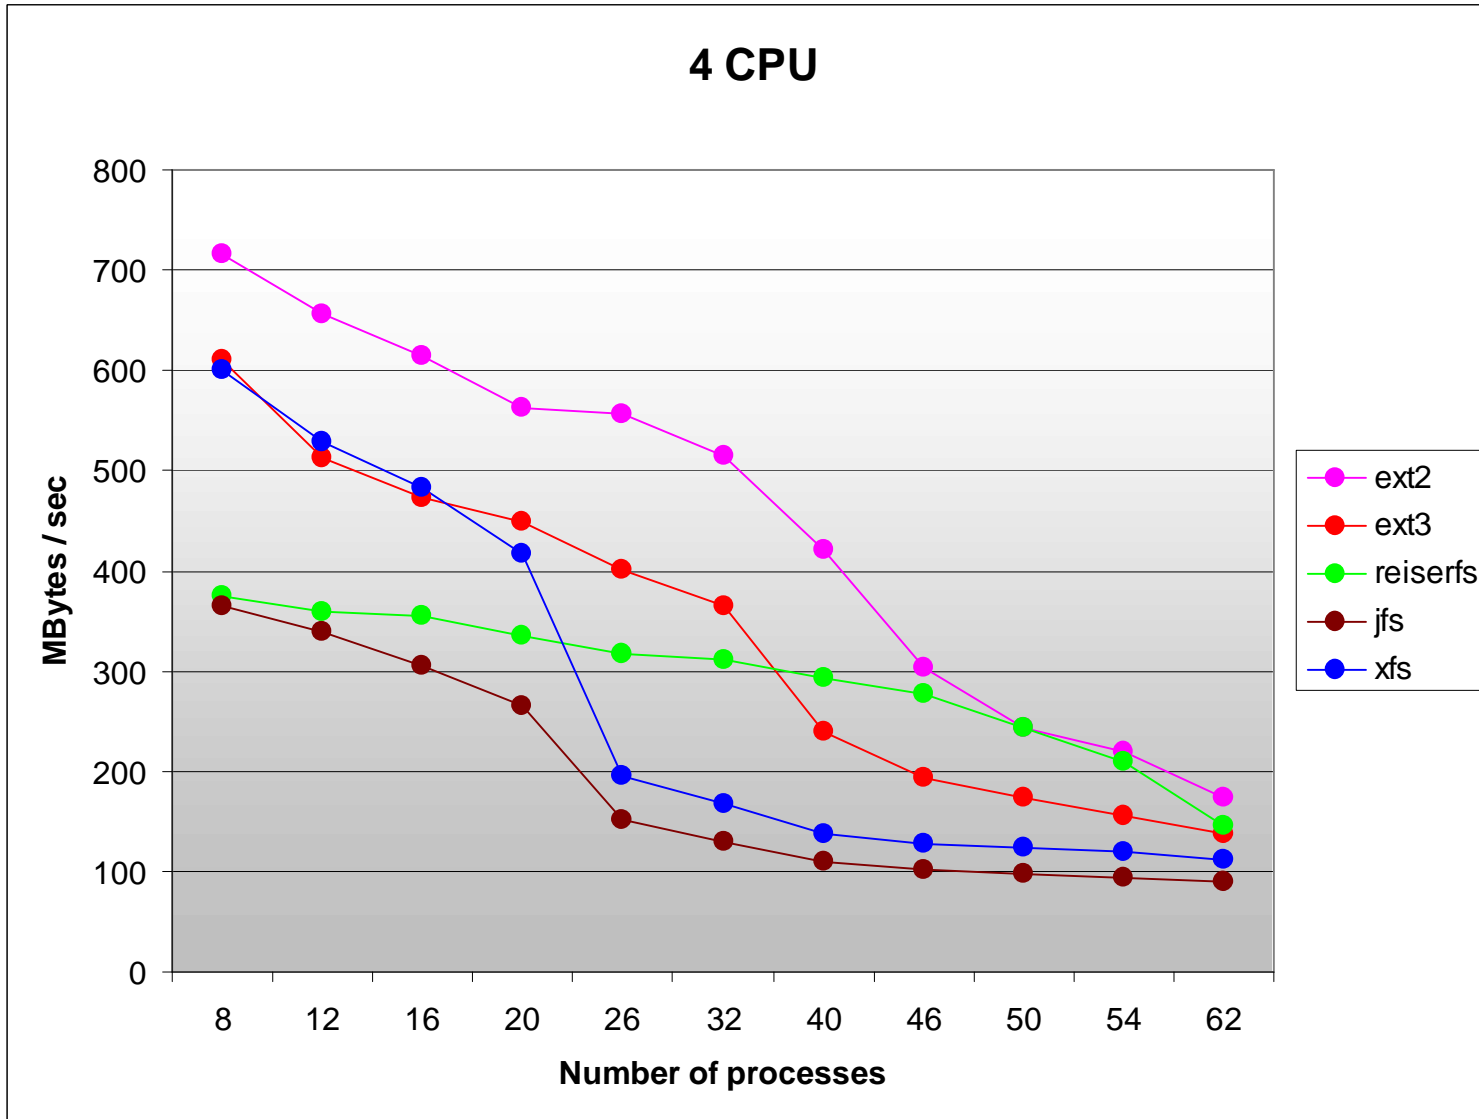


Recovery Times





Recovery of 96gb FS on 4 way w/ Shark







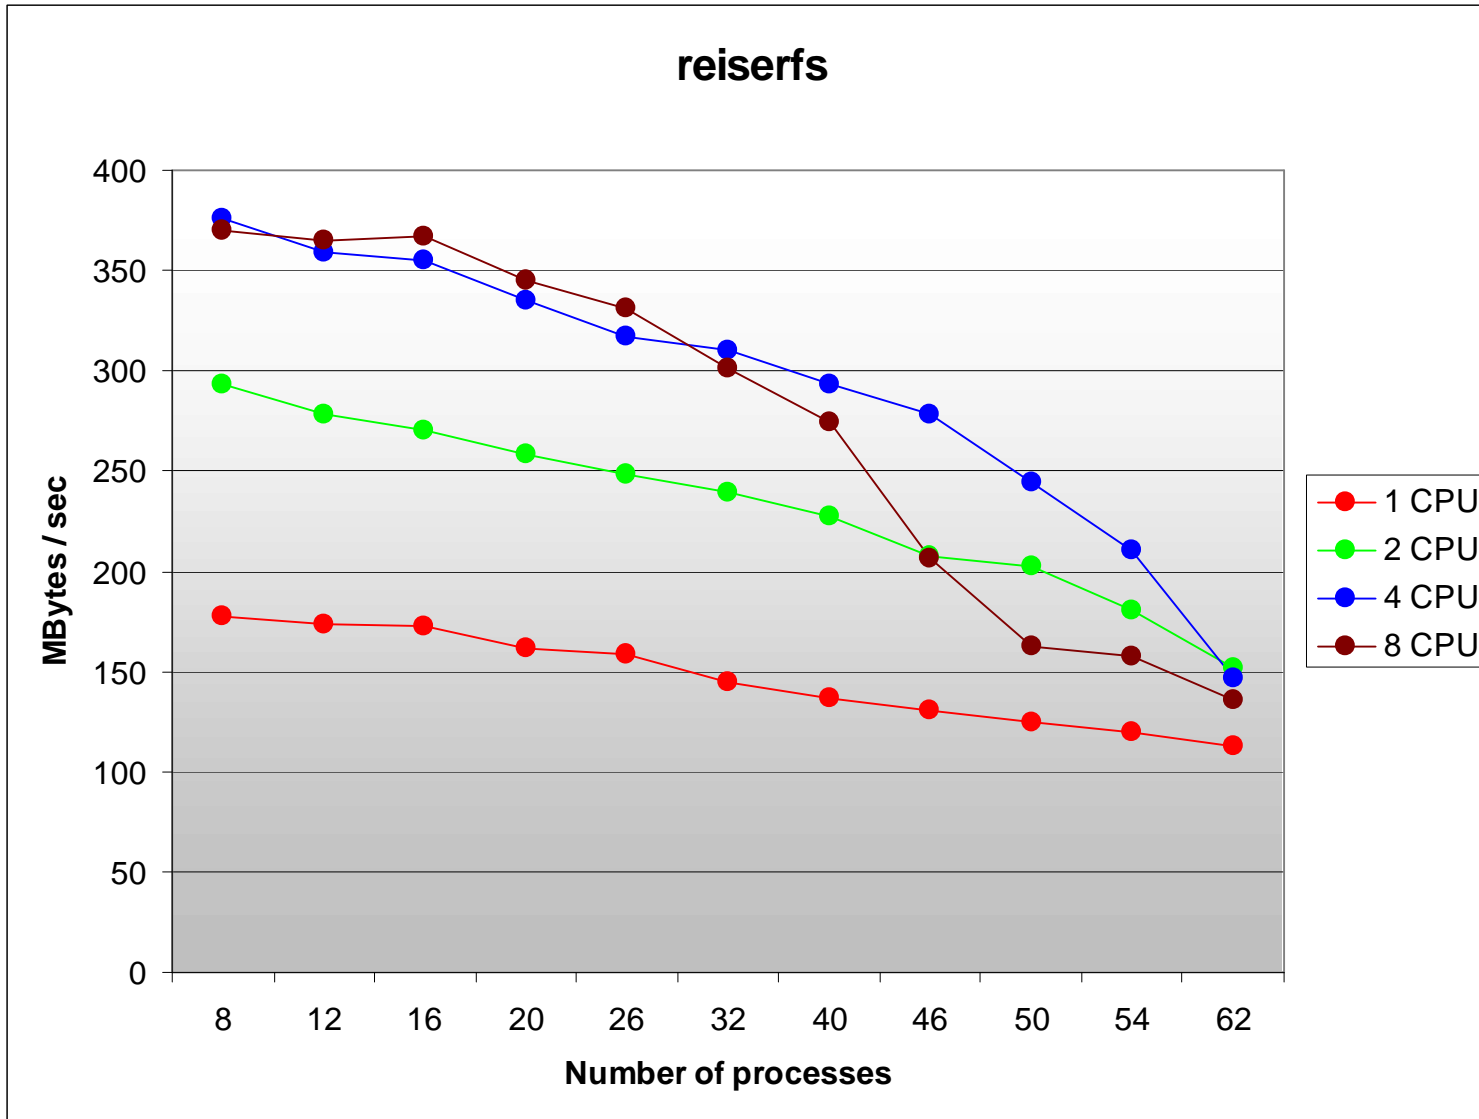
Throughput – FS Comparison – SCSI, LVM, 256MB








Throughput – FS Comparison – SCSI, LVM, 2 GB








Throughput – SCSI, LVM, 256 MB







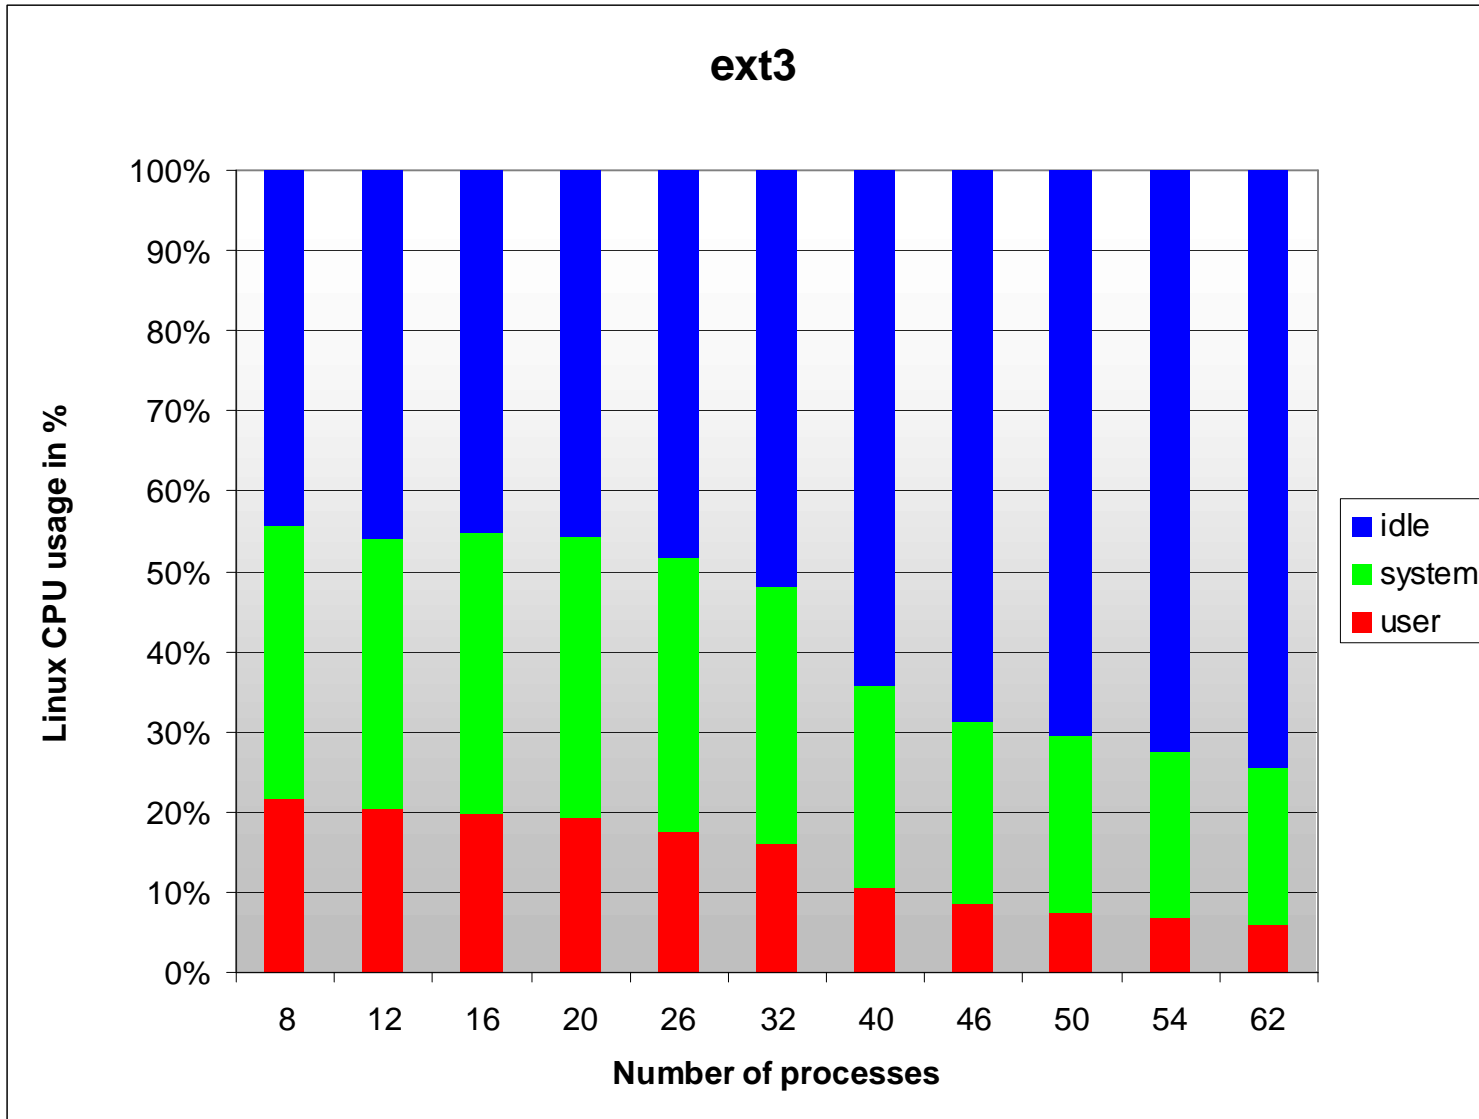
Throughput – SCSI, LVM, 2 GB








Throughput – SCSI, LVM, 2 CPU



CPU Load – SCSI, LVM, 4 CPU, 256 MB




CPU Load – SCSI, LVM, 4 CPU, 2 GB




z/VM Overhead


LVM CPU consumption






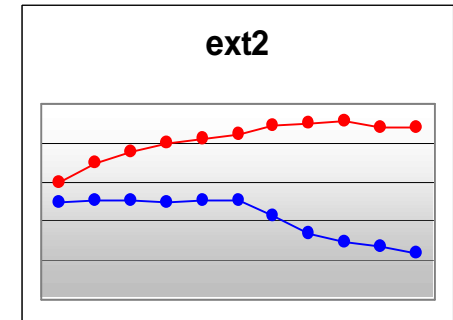

CPU Load – Comparison 265 MB




CPU Load – Comparison 2 GB




CPU Load – SCSI, LVM, 4 CPU




Throughput Per CPU Usage %



Throughput Per CPU Usage %



CPU Usage % In Relation To Throughput



Summary



- Journaling file systems increase data integrity significantly.
- Journaling file systems dramatically reduce system outage times.
- Performance cost is at least 30%.
- Journaling file systems profit from LVM.
- 2.6 brings more improvements (increased throughput, reduced CPU load).

Linux For zSeries Journaling File Systems



Questions ?

Trademarks

- **The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**
 - AIX, e-business logo, on-demand logo, IBM, IBM logo, OS/390, PR/SM, z900, z990, z800, z890, zSeries, S/390, z/OS, z/VM, FICON, ESCON

- **The following are trademarks or registered trademarks of other companies.**
 - LINUX is a registered trademark of Linus Torvalds
 - Penguin (Tux) complements of Larry Ewing
 - Tivoli is a trademark of Tivoli Systems Inc.
 - Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
 - UNIX is a registered trademark of The Open Group in the United States and other countries.
 - SMB, Microsoft, Windows are registered trademarks of Microsoft Corporation.

- * All other products may be trademarks or registered trademarks of their respective companies.

- Notes:
 - Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
 - IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
 - All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved.
 - Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
 - This publication was produced in Germany. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
 - All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
 - Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

