



Share Technical Conference
San Francisco 2/99
Share Session 2591



Carl Parris

**OS/390 Design and Performance Analysis
Mid-Hudson Valley Programming Lab
914-435-6287
parris@us.ibm.com**

Your Mileage May Vary Pt.1



The information contained in this presentation has not been submitted to any formal IBM test and is distributed on an "as is basis" without any warranty either expressed or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The information contained in this document represents the current views of IBM on the issues discussed as of the date of publication. IBM cannot guarantee the accuracy of any of the information presented after the date of publication.

Any performance data contained in this presentation was determined in a controlled environment; and therefore, the results which may be obtained in other operating environments may vary significantly. Some measurements quoted in this presentation may have been made on development-level systems. There is no guarantee that these measurements will be the same on generally available systems (continued on next page)

Your Mileage May Vary Pt.2



Some of the measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Limited. Other company, product and service names may be trademarks or service marks of others.

Trademarks or Registered trademarks of International Business Machines Inc.

ACF/VTAM tm

CICS tm

DB2 tm

DFSMS tm

ESCON tm

IBM tm

IMS tm

MVS tm

OpenEdition tm

OS/390 tm

RACF tm

RMF tm

S/390 tm

Your Mileage May Vary Pt.3



Unless otherwise noted all performance results contained in this presentation are from IBM internal workload measurements.

Agenda



- 'Why OS/390 UNIX' Overview
- Technical View of Performance
 - ▶ TCP/IP
 - ▶ File System
 - ▶ NFS
 - ▶ C/C++ Compiler
 - ▶ Kernel/RTL
 - ▶ DB2
 - ▶ Web
 - ▶ Net.commerce

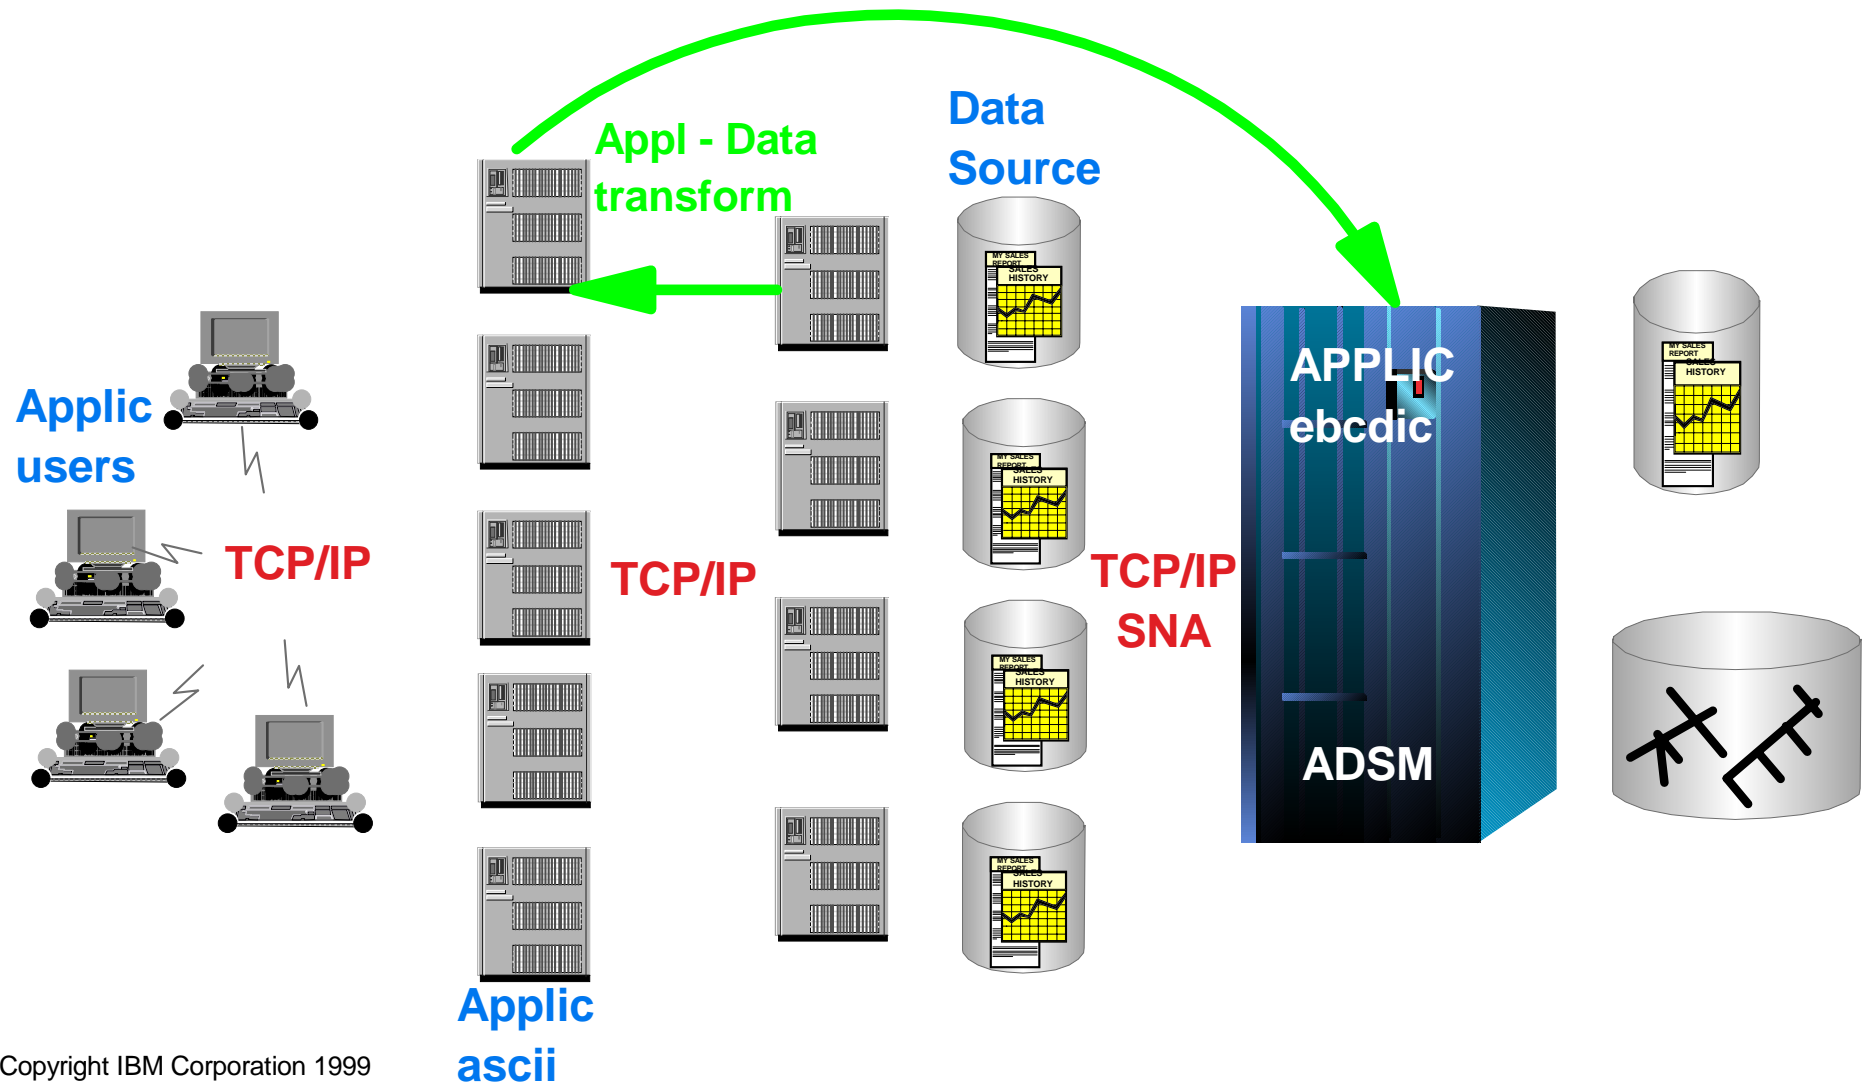
Overview

Today's Distributed Environment

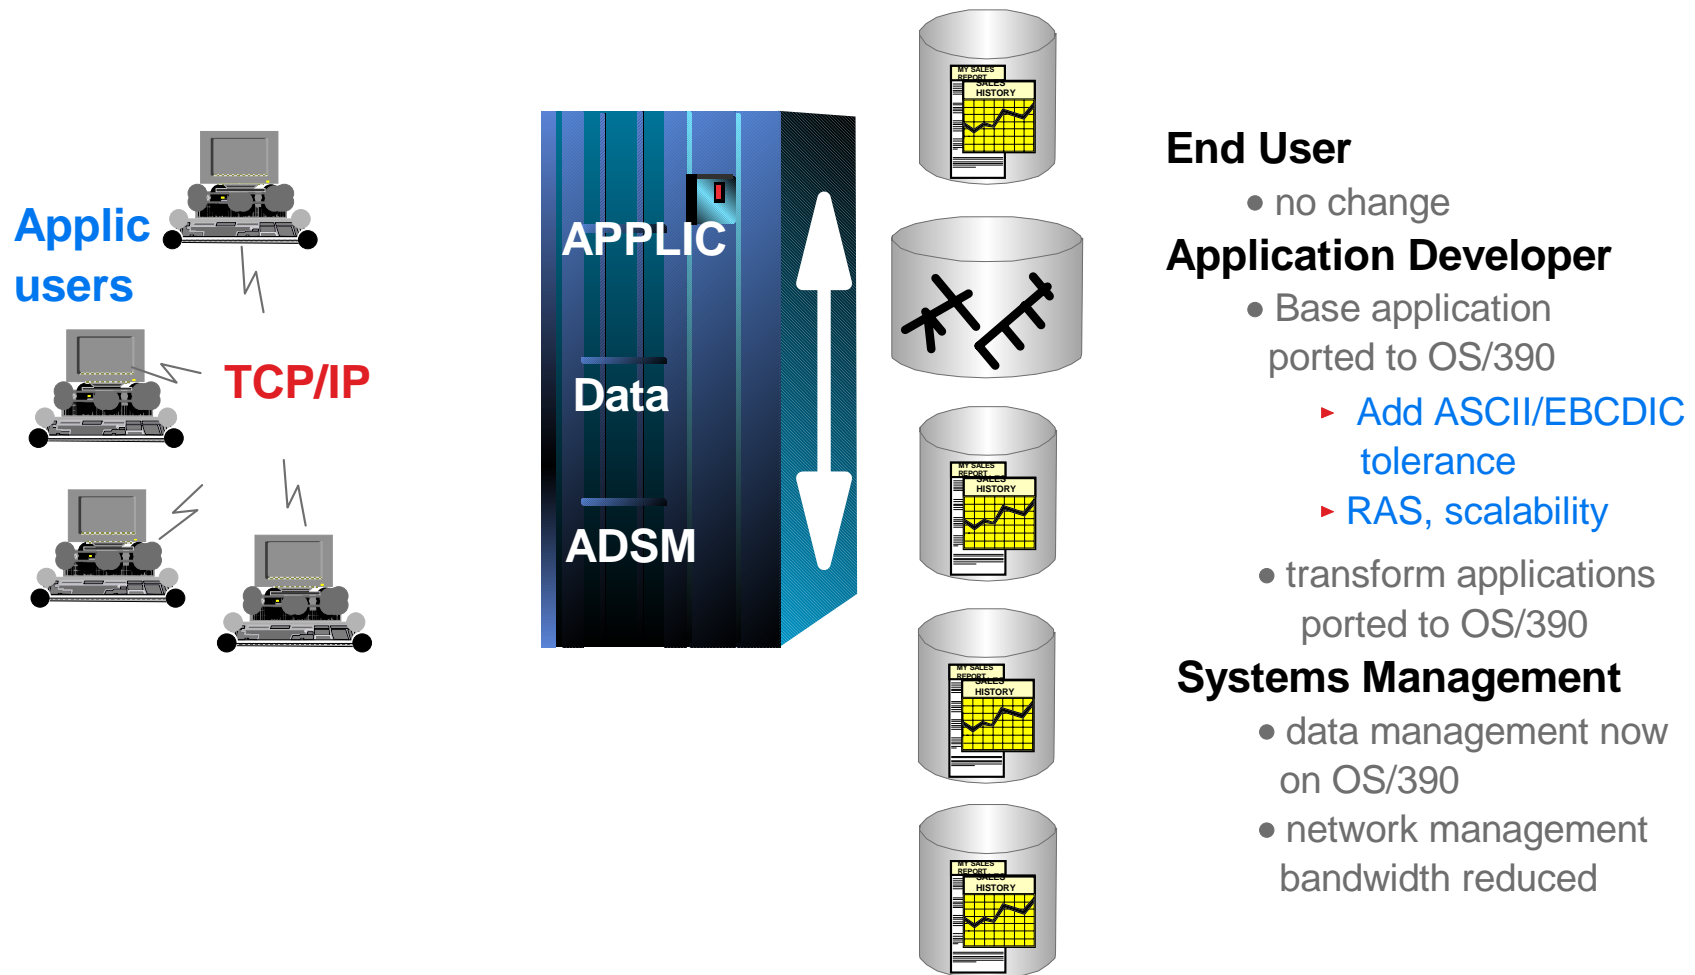


Major Data Moves

- End users access application (and data) via distributed process
- OS/390 and other servers transform data when both process it
- ADSM / MVS backup data from LAN file/database servers



Application Consolidation - End Step



End User

- no change

Application Developer

- Base application ported to OS/390
 - ▶ Add ASCII/EBCDIC tolerance
 - ▶ RAS, scalability

- transform applications ported to OS/390

Systems Management

- data management now on OS/390
- network management bandwidth reduced

Major Data Moves

- End users access applications (and data) via distributed protocols

Data Moves Avoided

- ADSM / MVS backups from LAN file/database server
- OS/390 and other servers transform data when both process it
- Application to Data move now occurs within OS/390

Application and Data Consolidation

Application migration using OS/390 to eliminate tiers of servers to Reduce:

- movement of data
- protocol handshakes
- data transforms
- points of failure
- complexity
- system management costs

and provide Improved:

- large scale deployment
- price/performance
- availability
- security
- resource balancing and management
- configuration management

Technical View

- OS/390 UNIX implementation is highly integrated in base operating system
- UNIX 95 Branding is functionally robust
- Performance Improvements coming regularly
- Development organization view performance as 'key' focus area
- Significant performance gains in '97, '98, '99



CS for OS/390 V2R6 for TCP/IP

- Protocol stack re-written for performance
- Improved SMP scalability
- RAS improvements

CS OS/390 V2R5 & V2R6 performance compared to MVS TCP/IP V3R2 on S/390 MP 3-way and 10-way configurations:

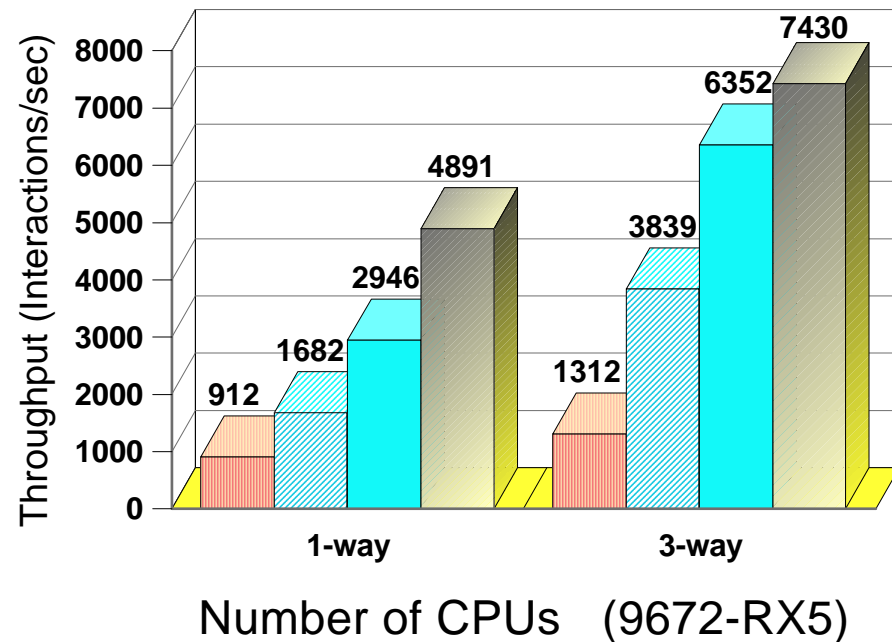
9672-RX53-Way:

Workload	CS OS/390 V2R5	CS OS/390 V2R6
Interactive	4.84X	5.66X
Web	4.16X	7.40X
File Tfr	1.64X	2.35X

9672-RY5 10-Way:

Workload	OS/390 V2R5	OS/390 V2R6
Interactive	11.51X	14.97X
Web	9.72X	14.22X
File Tfr	2.89X	5.31X

Interactive Workload



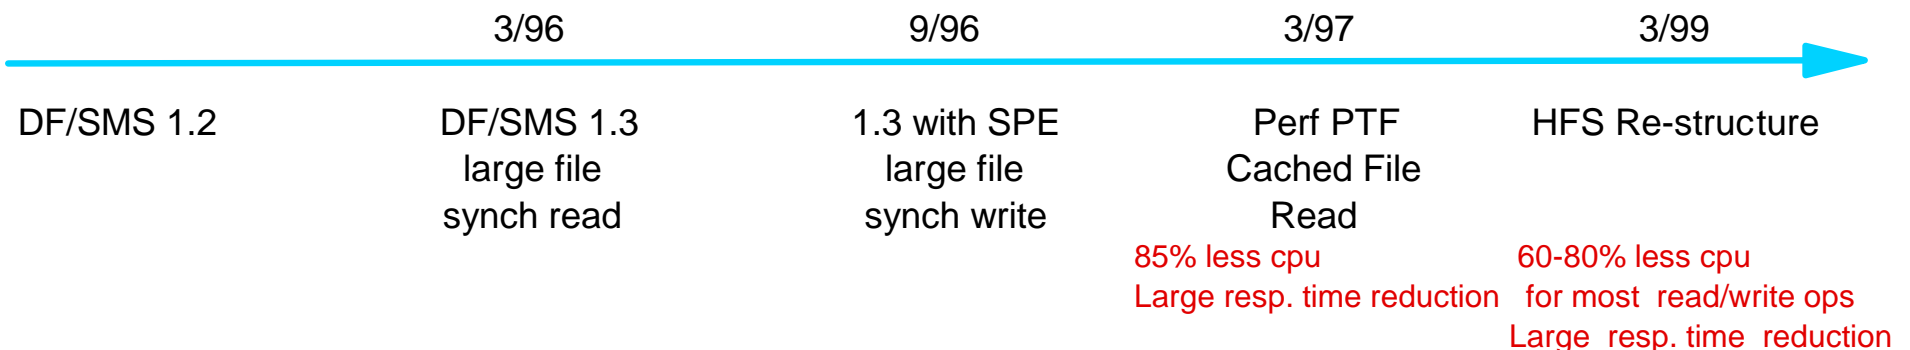
- TCP/IP V3R2*
- CS OS/390 V2R4
- CS OS/390 V2R5
- CS OS/390 V2R6

*Note: TCP/IP V3R2 tests were with performance enhancement PTFs applied. The performance difference would have been even greater without them.

HFS Performance



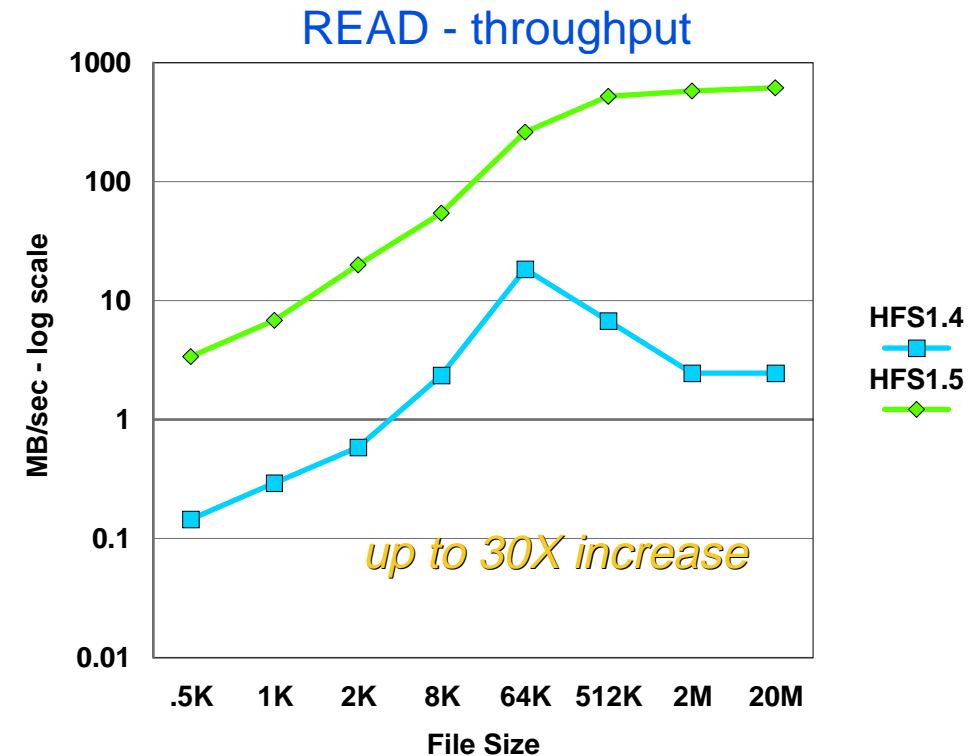

- UNIX Byte File Expectations
 - ▶ Large numbers of small files opened and closed frequently
 - ▶ Metadata I/O written on every file access
 - ▶ Synch daemon for deferred/asynchronous write
- Original HFS Implementation
 - ▶ Hardened all File Writes
 - ▶ File I/O synchronous
- New HFS Implementation
 - ▶ Deferred/asynchronous I/O metadata and file data
 - ▶ Index manager cache for index pages
 - ▶ Large fixed file data buffers
 - ▶ Pathlength reductions



HFS Performance Restructure

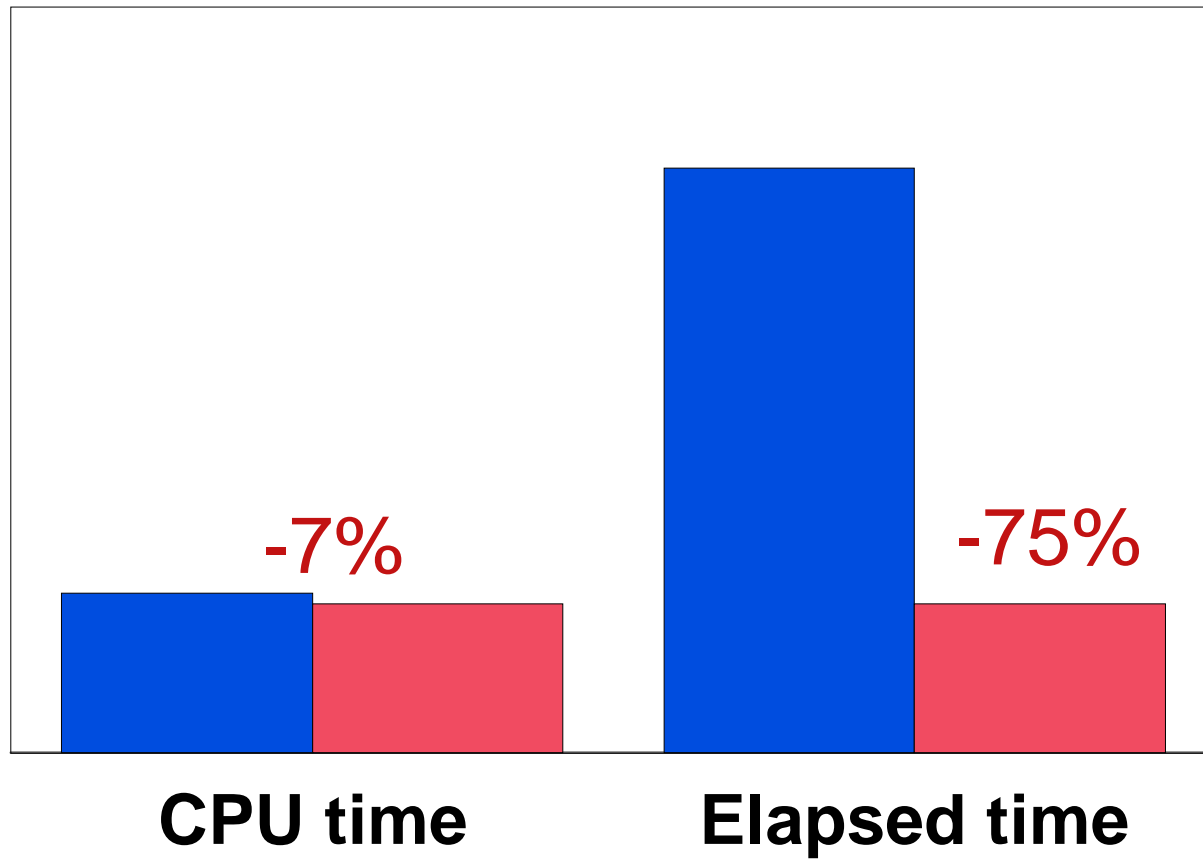


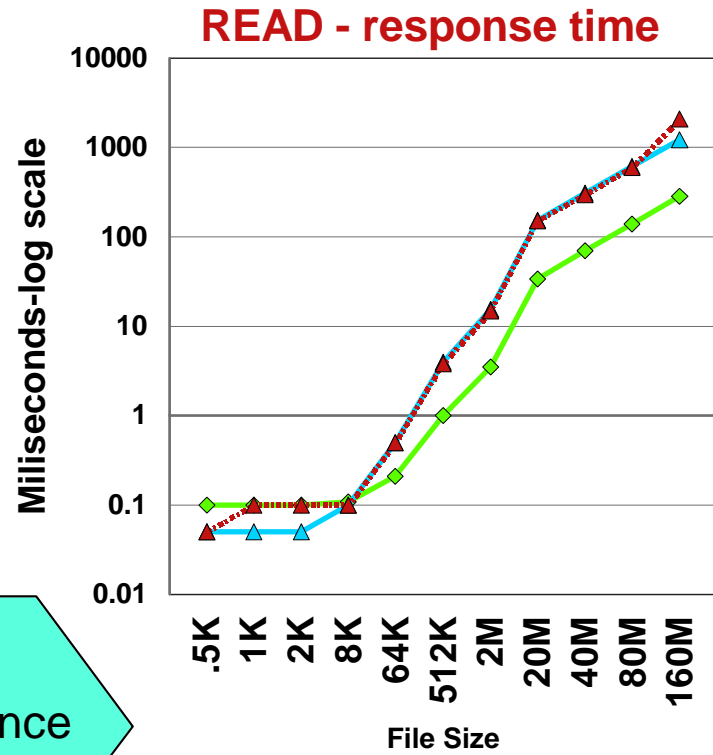
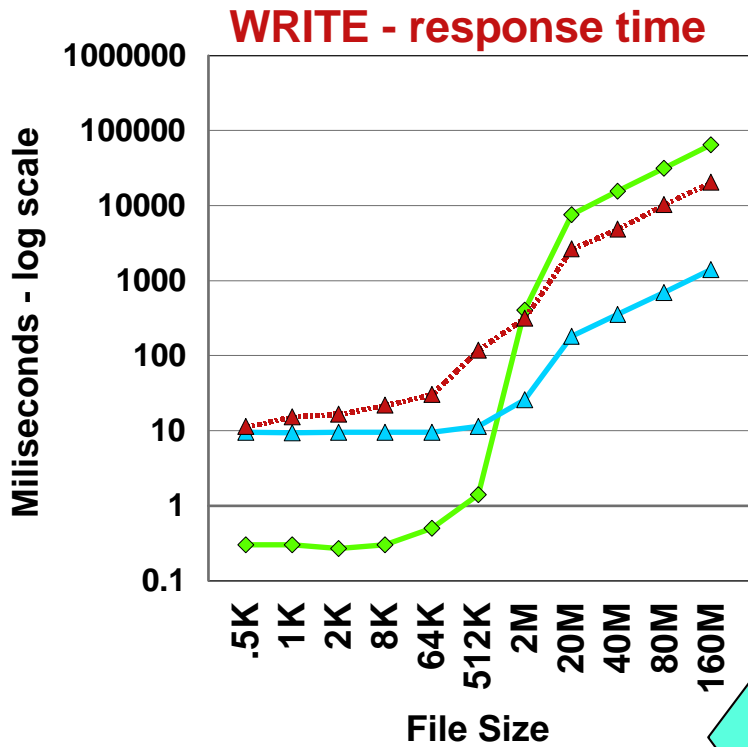
DFSMS1.5 GA- 3/99
Runs on OS/390 R5,R6,R7




- Up to 85% reduction in CPU time
- Order of magnitude improvement in throughput for files < 512KB
- Optimized for multi-user access
- Make time improved by 75%

Make test







HFS
Performance
Restructure



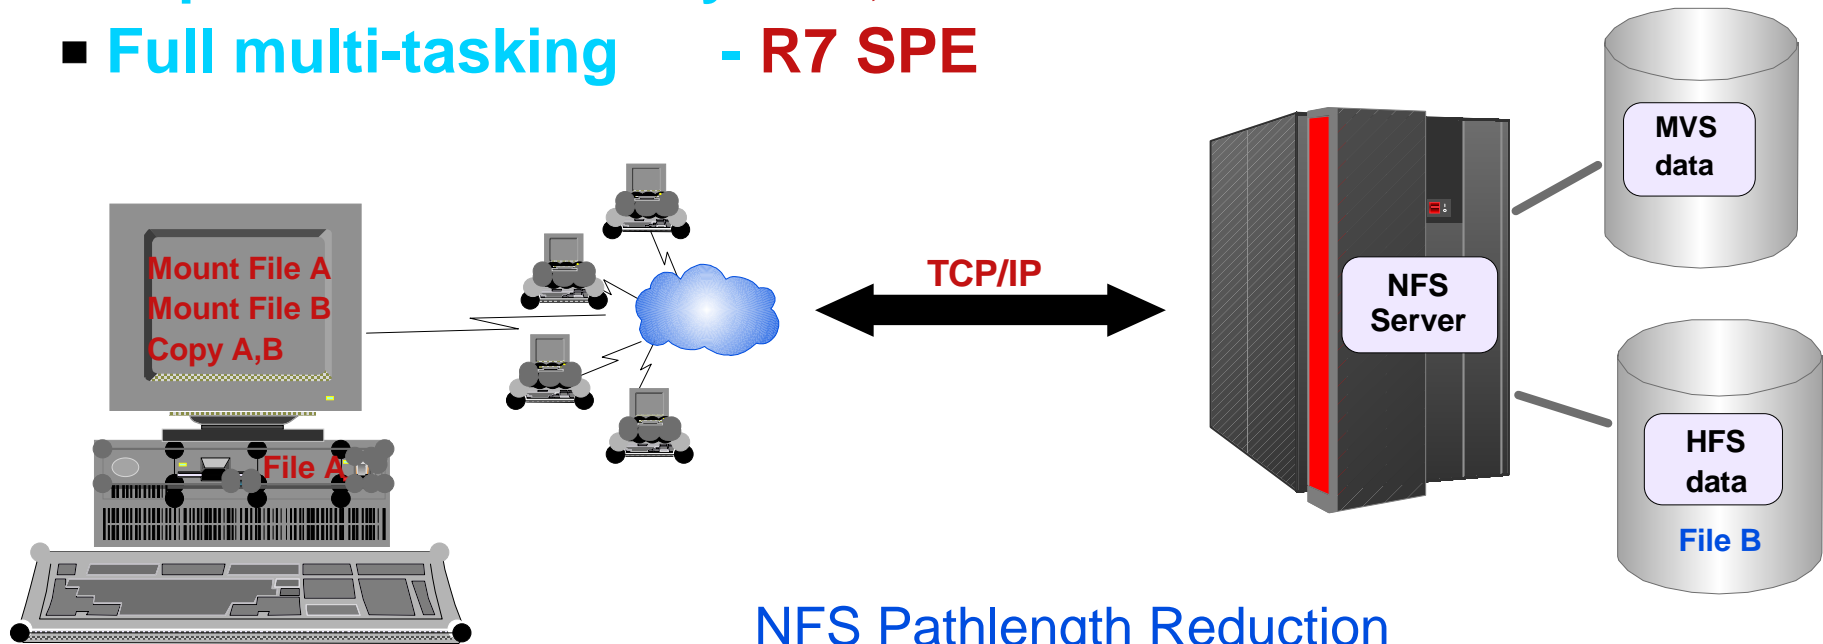
Better



NFS Performance Improvements



- Pathlength reduction - R6,R7
- Improved scalability - R6,R7
- Full multi-tasking - R7 SPE



NFS Pathlength Reduction

OS/390 Rel	R3	R6	R7(SPE)
Pathlength	Base	80%	90%

NFS Throughput Gains - R6 only

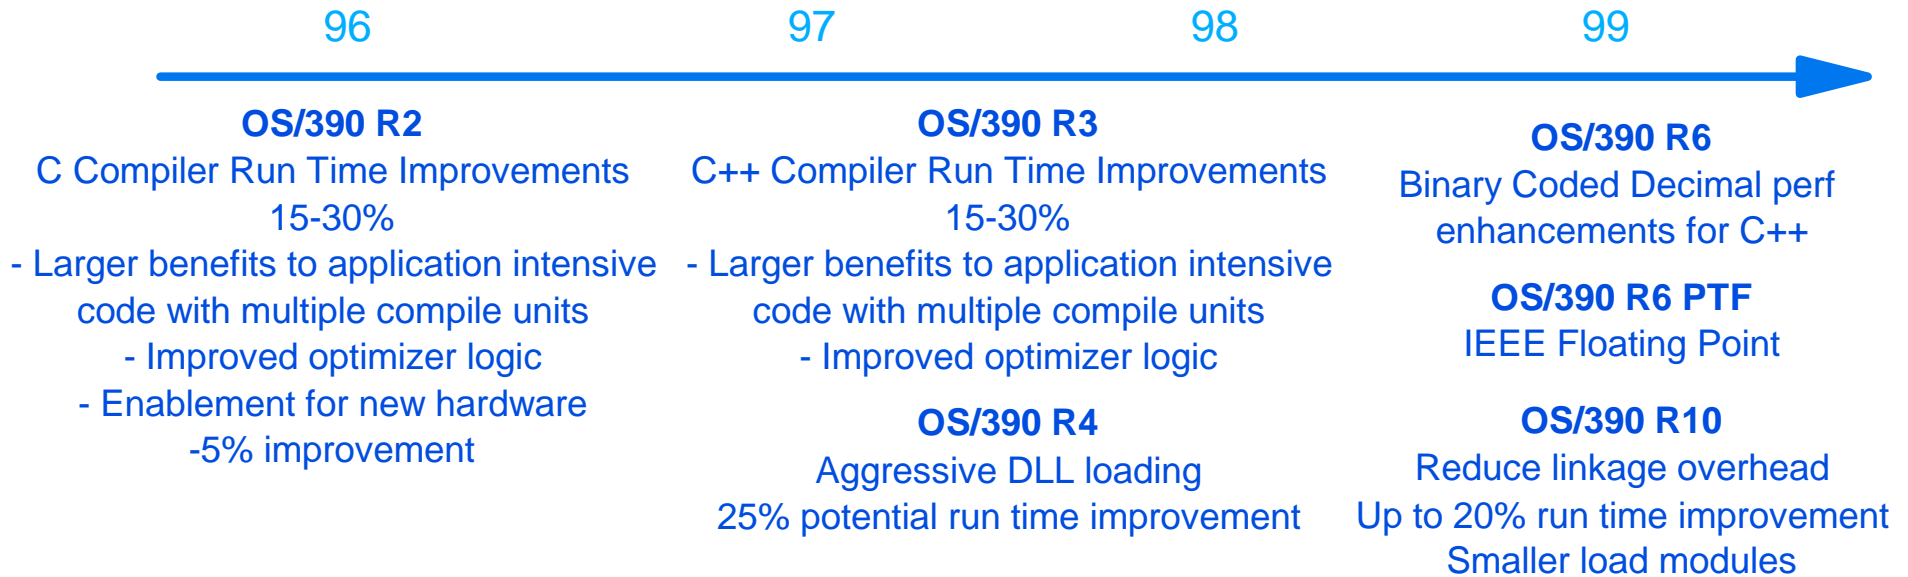
File Size	48K	256k	20MB	100MB
Single Client Read	55%	70%	85%	30%
Single Client Write	70%	10%	5%	30%
Multi-client Read	*	*	180%	*
Multi-client Write	*	*	50%	*

* not measured


C/C++ Compiler Improvements



Run Time

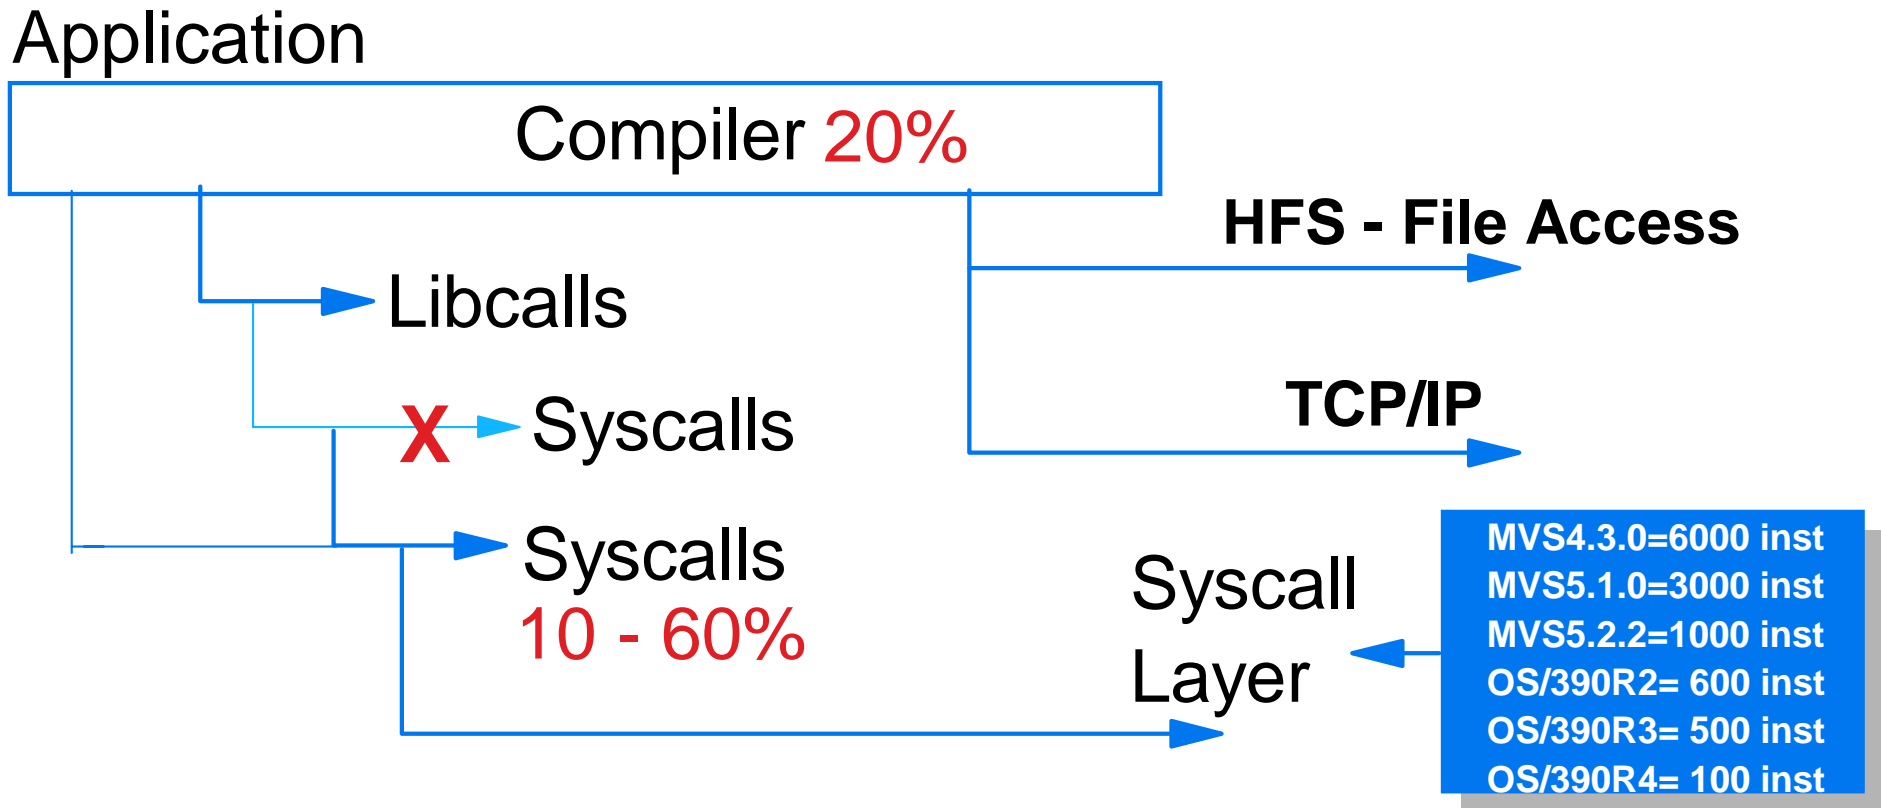


Compile Time




New Processor architecture on CMOS for C compiler
State of the Art Optimizer Technology Now!

Kernel/RTL Improvements



Focus on Pathlength Reduction for High Frequency Services

Other Kernel Performance Improvements



- Multiproc/Multiuser Support
 - Improves performance for DB2 Stored Procedures and Web
 - R6 with support back to R4
- Binary Semaphores using the PLO instruction
 - improved locking performance where GRS lock contention existed
 - R6 with support back to R4
- File system bundled syscalls - reduces pathlength
 - `_open_stat()`
 - `_accept_and_recv()` - benefits the web
 - `send_file()` - benefits the web and ftp type servers
 - R6
- Improved shared memory processing
 - greatly reduced ESQA and LSQA memory usage
 - R6
- Signaling performance enhancements
 - R6
- Optimized security checking in perf. sensitive syscalls
 - Message queues, semaphores, File system(stat,open, lookups), ps cmd
 - R7
- Fast CGI Support for Web
 - R7

Other Kernel Perf Improvements cont...



- Asynch I/O for sockets
 - multiple socket read/writes w/o waiting for I/O complete
 - R7
- Shell and Utilities
 - recompiled with OPT(2) and IPA - highest levels of optimization
 - performance tune up: vi, find, pax, tar, cpio, compress, uncompress, zcat, grep, ar, ls, du
 - R7
- Message Queues using the PLO instruction
 - potential performance benefit depending on GRS latch contention rate
 - R8
- Additional asynch I/O support
 - avoids the inefficiency of SELECT - benefits Lotus
 - R8
- Major improvement in DBX performance
 - use of local spawn instead of fork/exec
 - R8

For more detail see Don Ault's presentation:

OS/390 Unix System Services

Update for

OS/390 V2R6, V2R7, and V2R8

<http://www.s390.ibm.com/unix/>

Focus on DB2 Connectivity



■ Stored Procedures:

- Decreases costs with less network traffic
- May be used by local or remote applications
- Protects sensitive portions of your applications from unauthorized use
- Offers more programming and application design choices
- Allows access to non-SQL sources
- Provides greater flexibility for change

■ Elimination of Gateways

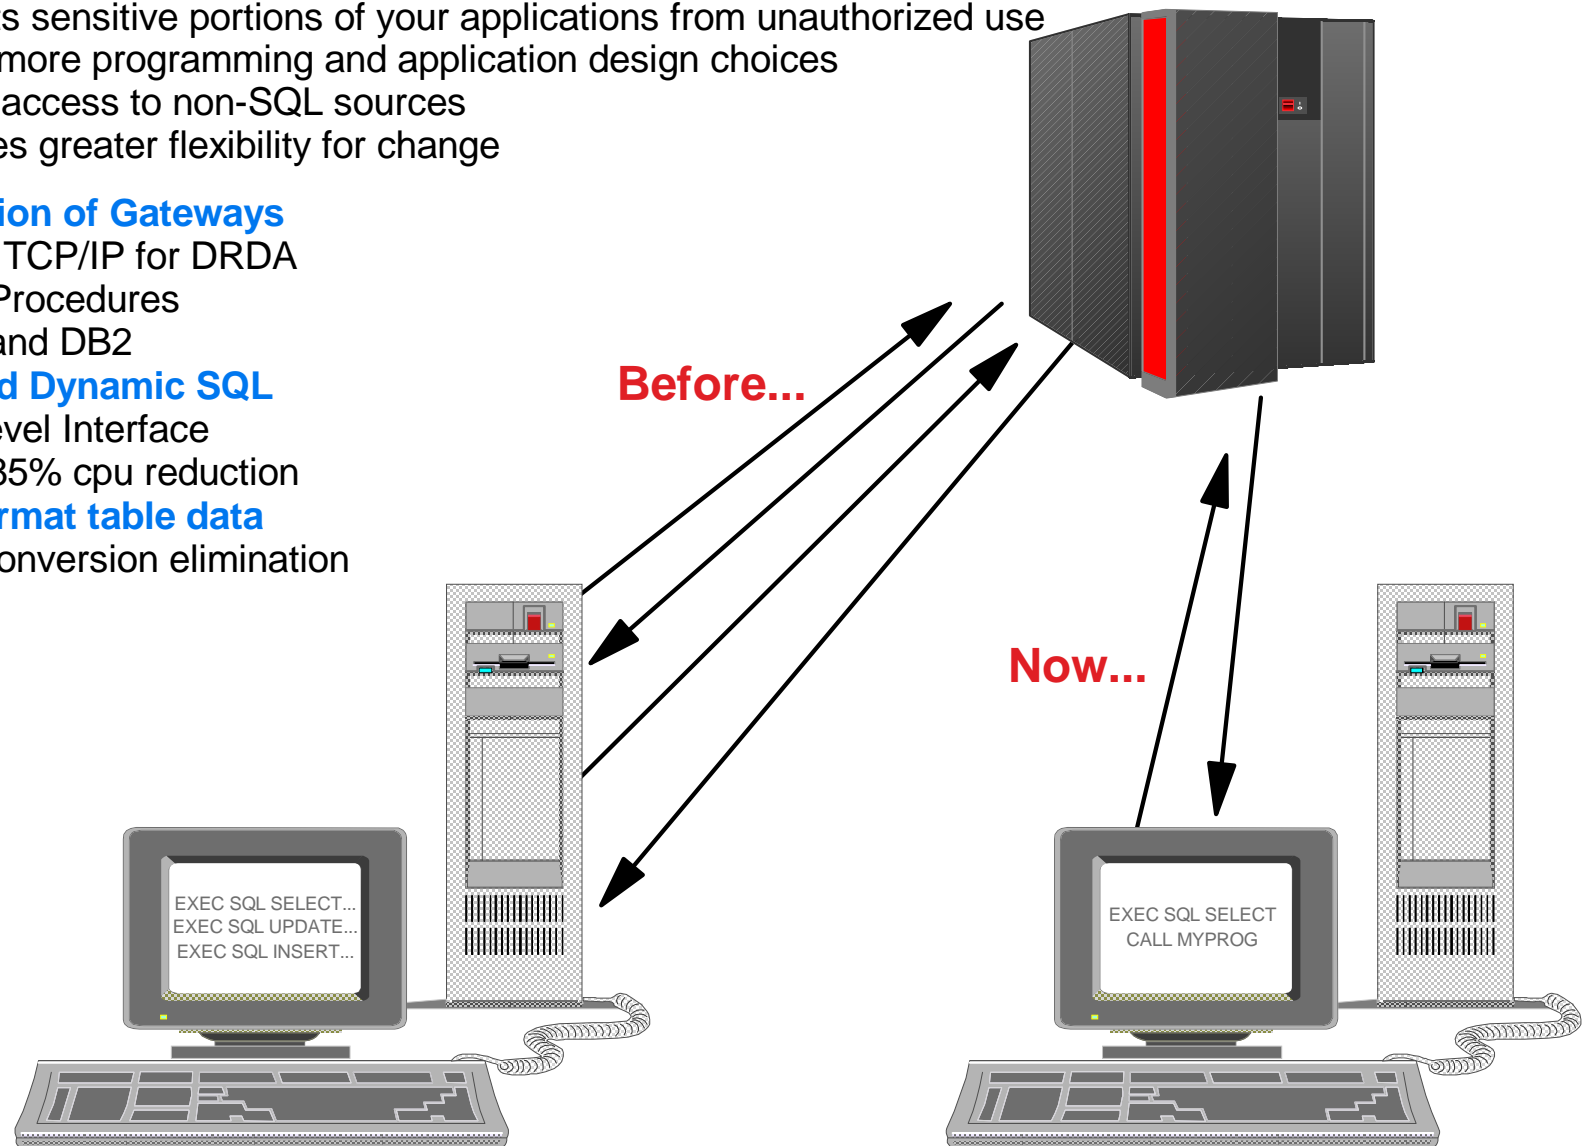
- Native TCP/IP for DRDA
- Stored Procedures
- WLM and DB2

■ Improved Dynamic SQL

- Call Level Interface
- Up to 85% cpu reduction

■ ASCII format table data

- Data conversion elimination

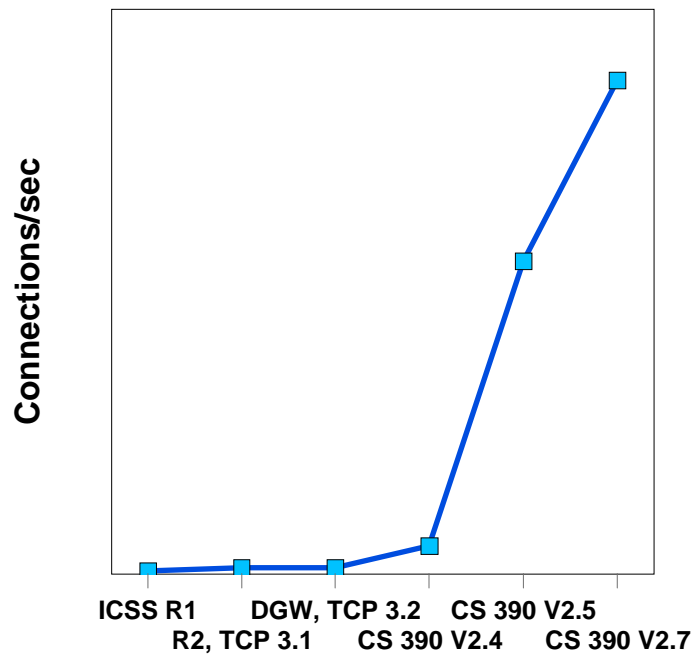


Web Serving Performance




- Websphere Application Server, TCP/IP, Unix/kernel services, and Crypto HW improvements combine to provide large reductions in WEB serving pathlength

Webstone Throughput



SSL Throughput



- 99% reduction in pathlength
- 170X increase in connections/sec

- 80% reduction in pathlength
- 6X increase in connections/sec

SPECWeb96 Published Results - 2/16/99

Rank	Company	System	Result	HTTP Version	# CPU
1	IBM Corporation	S/390 9672-YX6 G5	21591	IBM HTTP Server V5.1 for OS/390	10
2	Compaq	Compaq AlphaServer GS140 6/575	14263	Zeus 1.3.0	10
3	Hewlett-Packard	HP 9000 V2250	13811	Zeus 1.3.3	16
4	IBM	RS/6000 S70	12031	Zeus 1.3.3	12
5	Siemens	Primergy 870-40	10212	SWS 2.1 and SNCA	4
6	Siemens	Primergy 870-40	9199	SWS 2.1 and SNCA	4
7	Sun Microsystems Inc	Sun Enterprise 450	9115	SWS 2.1 and SNCA	4
8	IBM	RS/6000 S70 RS642	9081	Zeus 1.3.0	12
9	NCR Corporation	4400	9046	SWS 2.1 and SNCA	4
10	IBM	RS/6000 S70 RS642	8301	Zeus 1.3.0	12
11	IBM	Netfinity 7000 M10	7818	IBM HTTP Server 1.3.3	4
12	NCR Corporation	4400-4000-8090	7800	SWS 2.1 and SNCA	4
13	Hewlett-Packard	NetServer LXr 8000	7531	IIS 4.0 and SWS 1.1	4
14	Silicon Graphics	Origin 2000	7214	FastTrack 3.02	8
15	IBM	Netfinity 7000 M10	7180	IBM HTTP Server 1.3.3	4
16	IBM Austin	RS/6000 Model S70 RS642 Upgrade	7013	Zeus 1.3.0	12
17	Hewlett-Packard	NetServer LXr 8000	7006	Microsoft SWS 1.1 and IIS 4.0	4
18	Hewlett-Packard	9000/K580	6265	Zeus 1.3.3	4
19	Silicon Graphics	Origin 2000	6200	Zeus 1.3	8
20	IBM	Netfinity 7000 M10	5996	IBM HTTP Server 1.3.3	2

S/390 Results 50% better than previous BoB

Results from IDEAS International: www.ideasinternational.com

also check the SPEC Corp: www.specbench.org

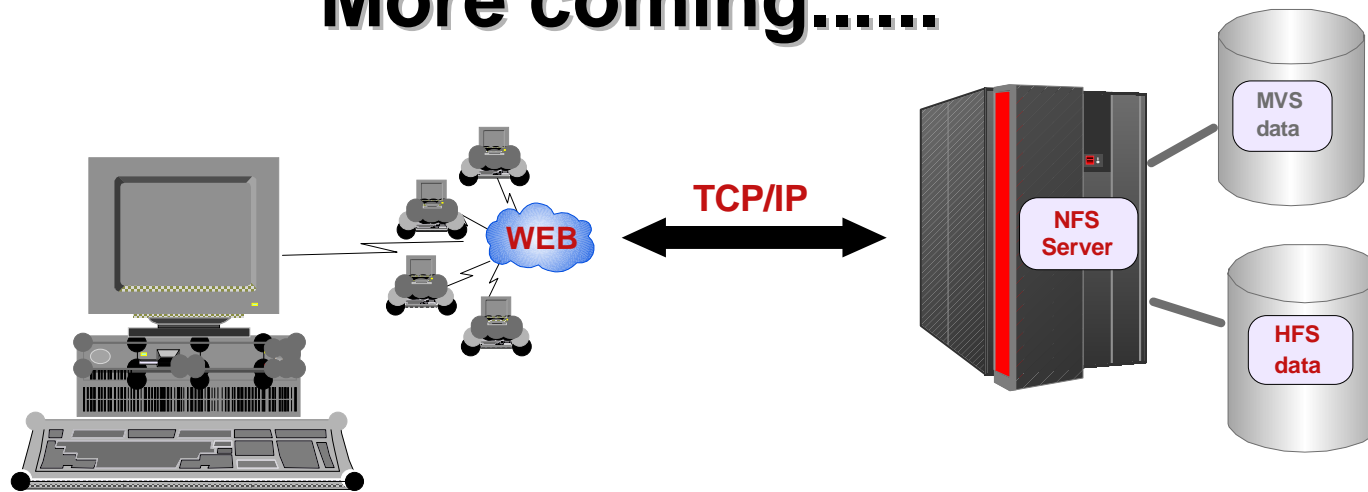
- **JDK 1.1.4**
 - ▶ Single thread performance - 85% improvement
 - ▶ Multithread performance - 10X improvement
 - available 5/98
- **JDK 1.1.6**
 - ▶ 2X increase in throughput for many applications vs JDK1.1.4
 - ▶ 100X increase in throughput for floating point intensive applications when using G5 IEEE FP hardware
 - available 12/98
- **JDK 1.1.6p**
 - ▶ significant additional performance improvements
 - available 3/99
- **JDK 1.1.8**
 - ▶ projected 2X performance improvement vs JDK1.1.6
 - available summer '99

Watch this space !!!

OS/390 Unix Performance



**Dramatic Improvements
Competitive Performance
More coming.....**



-
- ▶ **TCP/IP** 2.3X-15X transactions /sec

 - ▶ **NFS** up to 2.8X multiple client file processing

 - ▶ **HFS** up to 40X file throughput processing

 - ▶ **Web serving** up to 170X connections/sec
BoB SPECWeb96 results

 - ▶ **JAVA** up to 100X IEEE F.Pt. operations
