Benchmarketing

Stephen Turner
ITSO Poughkeepsie

IBM Technical Support

Agenda

- Overview/Introduction
- Computer Architectures and Performance Issues
- Industry Standard Benchmarks
- Strengths and Weaknesses of Industry-Standard benchmarks
- S/390 Strengths not measured by benchmarks

IBM Technical Support

Introduction/Overview

Customer Demand for Capacity

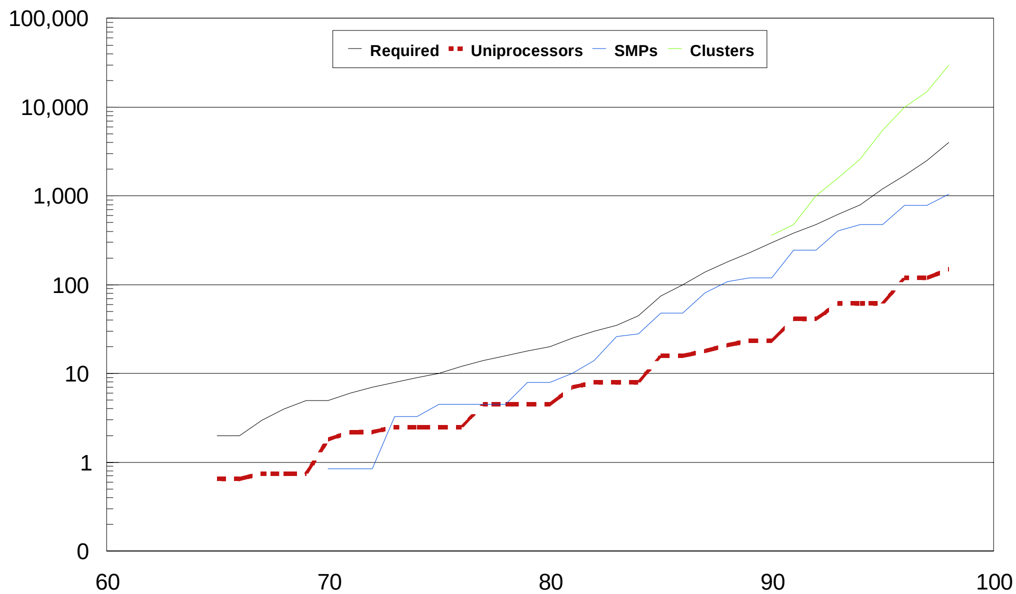

Which Car is Best?

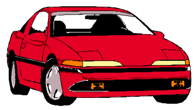

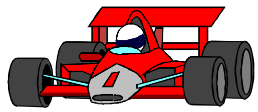

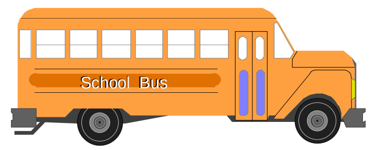

IBM Technical Support

What is a Benchmark?

Basic types of benchmarks:

- Measurement of all system performance
- Measurement of specific performance
- Measurement for marketing usage

IBM Technical Support

Measurement Criteria

- Throughput
- External throughput
- Internal throughput
- Instruction execution rate
- Response time

Key items for judging the system

- Performance
  - Throughput
  - Response time
- Cost of the system

Computer Architectures and Performance Issues

Computer System Elements

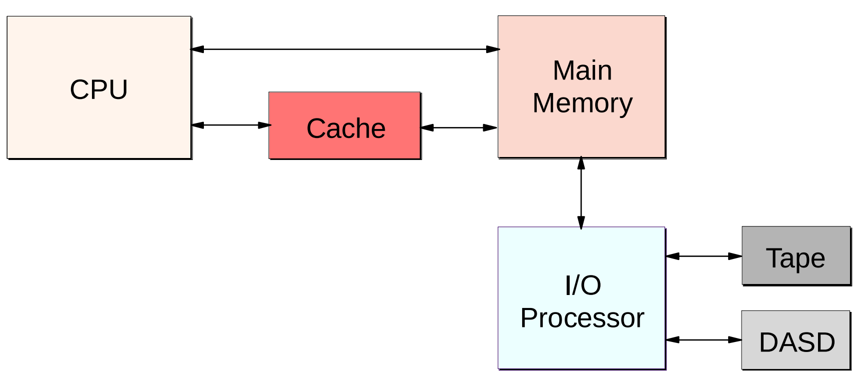

SMP (General)

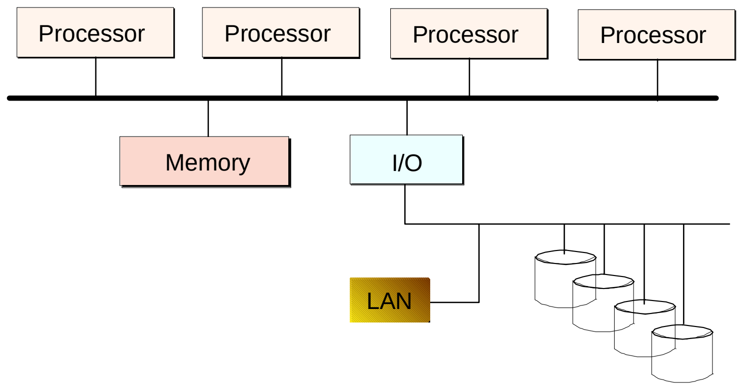

Non - SMPs

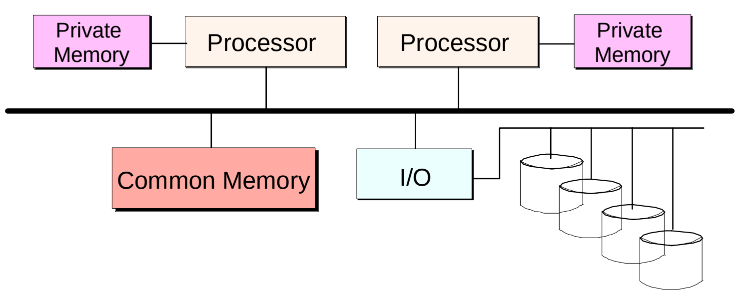

Multiple memories,
not just processors.
(also, not symmetric)

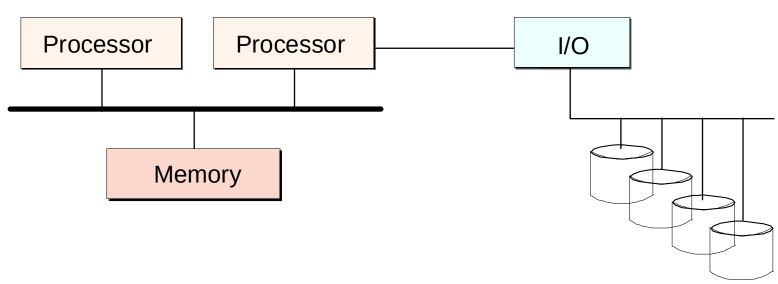

Not symmetric:
Only one processor
can do I/O

Message versus Shared

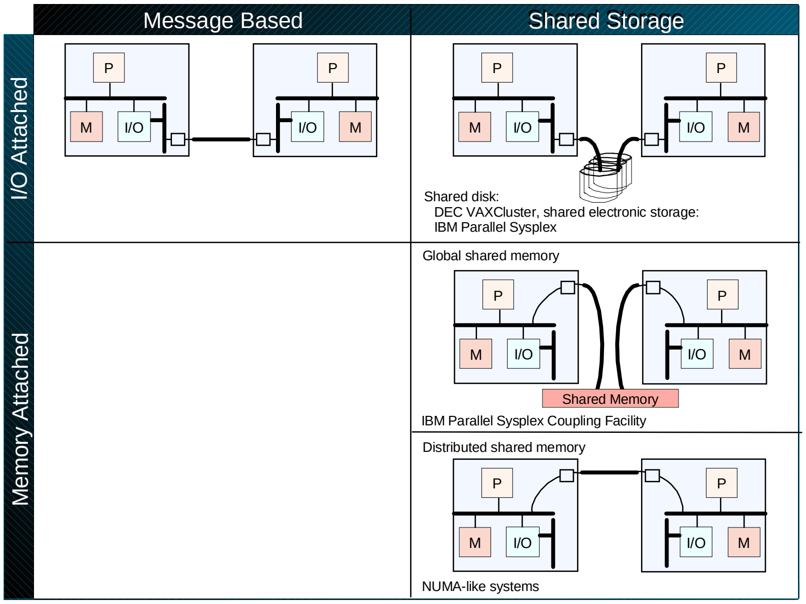

IBM Technical Support

Multiprocessor Classifications

Multiprocessors are classified by the way the memory is organized:

- Common Shared Memory
  - Shared memory / Tightly coupled multiprocessor
- Distributed Memory
  - Loosely coupled multiprocessor

IBM Technical Support

Multiprocessor System

Common Shared Memory

- has cache memory with each CPU
- has a global shared memory for all CPU
  (information shared)
- higher degree of interaction

Shared Memory System

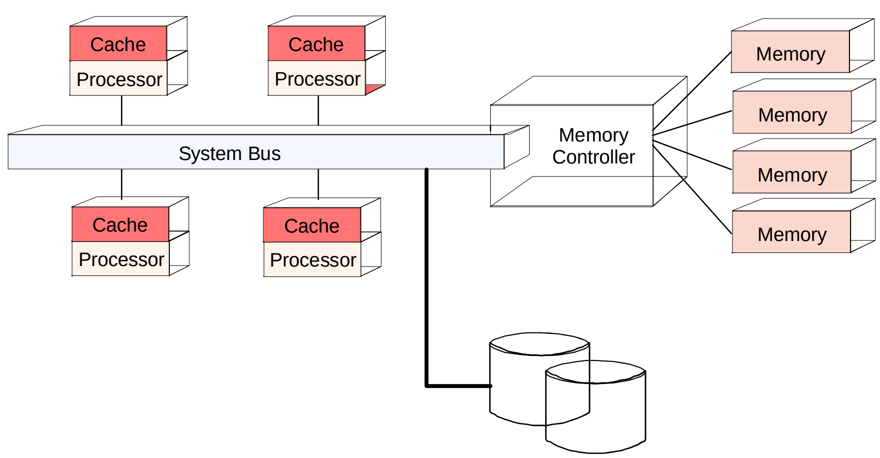

Microprocessor Efficiency

Multiprocessor efficiency depends on:

- Delay due to inter-processor communication
- Overhead in synchronizing the work of one processor with another (overhead in serial mode)
- Processing cost for controlling the system and scheduling processing unit
- Lost efficiency
  - ran out of tasks
  - wasted effort

Amdahl's Law

$$\text{Total execution time} = \frac{\text{parallel part}}{N} + \text{serial part}$$

IBM Technical Support
© Copyright IBM Corporation, 1999

Multiprocessor Systems

Clustered Multi-Processor type of design

1. Advantages
   - Straightforward
   - Highly scalable
   - Very inexpensive
2. Drawbacks
   - Absence of the shared memory may result in:
     - Difficulty of sharing data between programs
     - Lack of software support
     - Confining CMP-based machines to the scientific environment

Distributed Memory

Distributed Memory

- has a private local memory
- tied together by a switching scheme
- route information among processors via message passing scheme
- relay data via packet
  - Advantage
    Interaction between tasks is minimal

Distributed Memory System

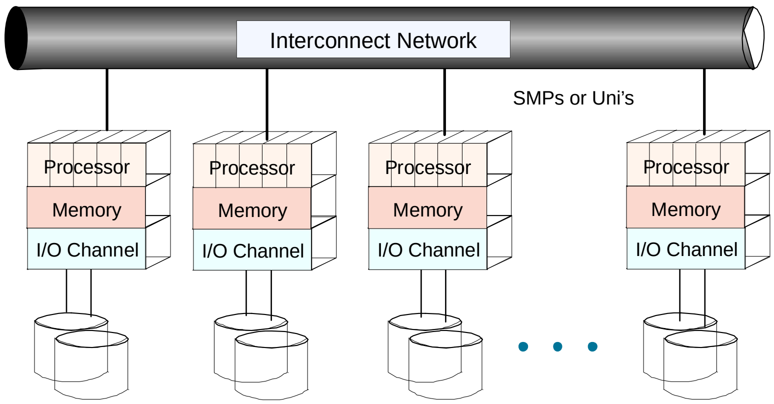

IBM Technical Support

Distributed Systems?

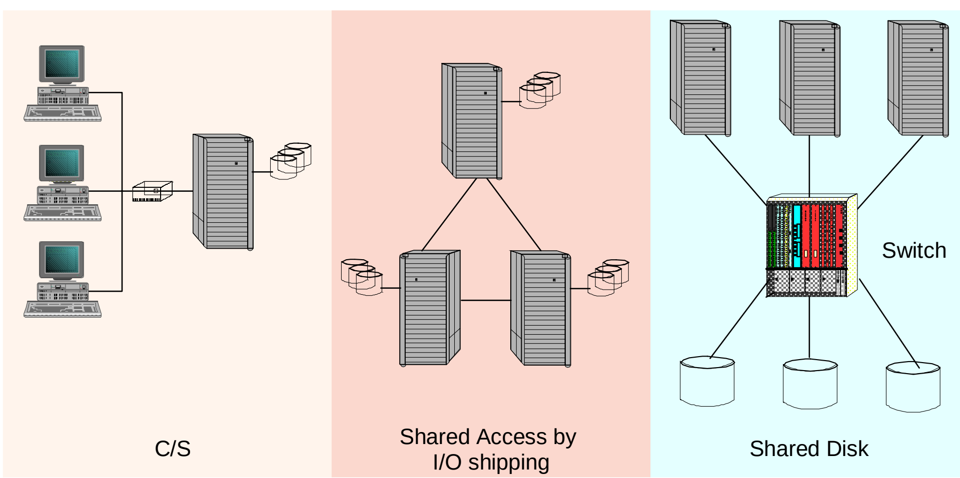

IBM Technical Support

Memory Design

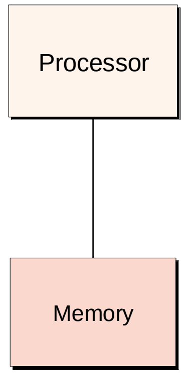

- CPU cycle
- Memory cycle (<2.5 X large than CPU cycle time)

Cache Memory

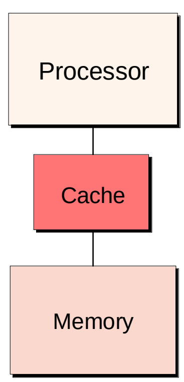

- CPU - too fast for memory
- Cache - memory run at processor speed
- Memory - too slow for processor
- Cache hit/miss ratio

TCMP Two-Level Caching

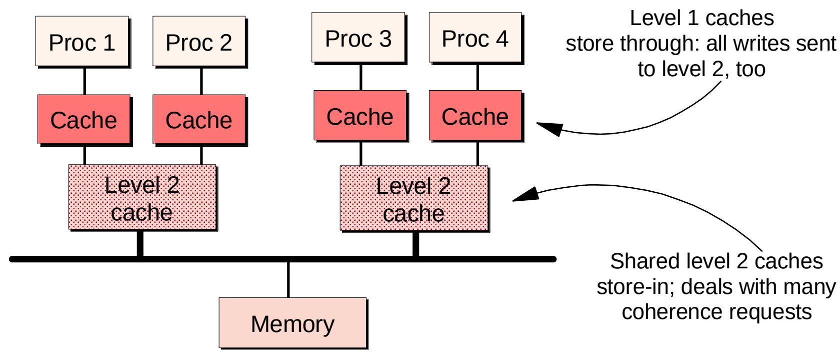

Cache Operations

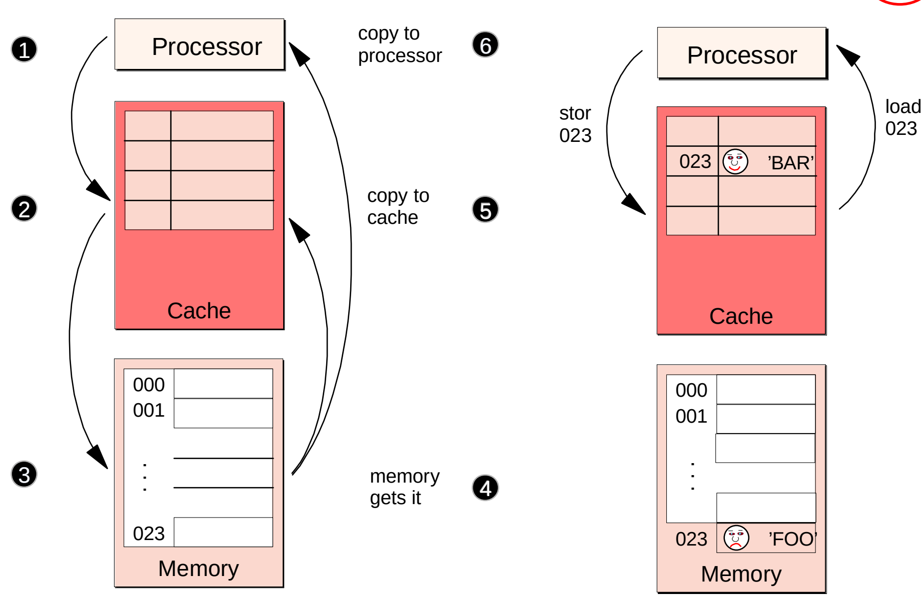

Level 2 cache

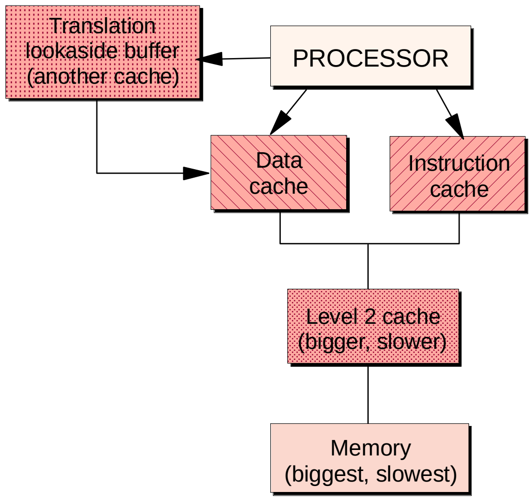

Characteristics		
Element	Size	Speed (nanoseconds)
Processor		10 ns
Level 1 Caches	10s-100s of KB	10 ns
Level 2 Caches	100s of KB to MBs	40-60 ns
Main Memory	MBs to GBs	100s of ns

IBM Technical Support

DEC System

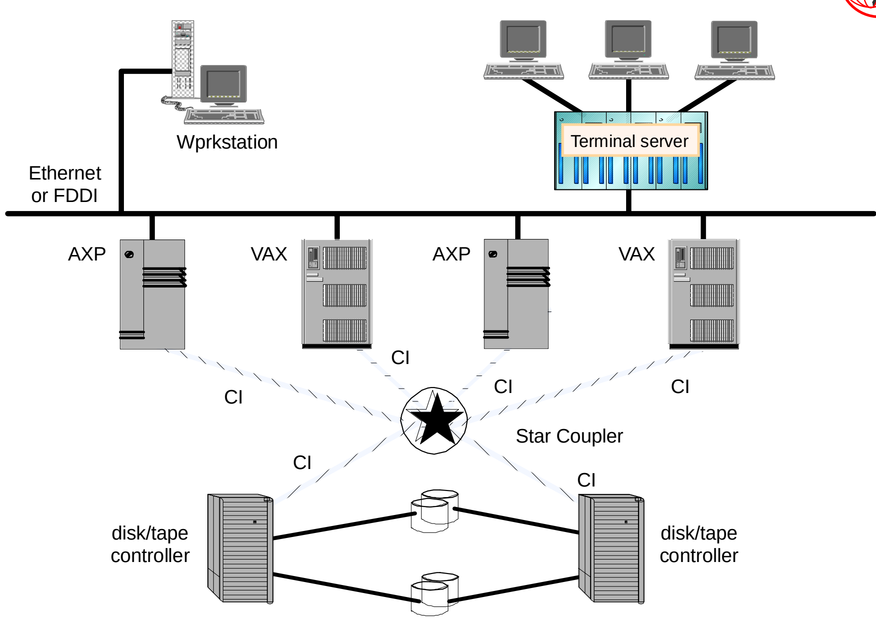

IBM Technical Support

Tandem System

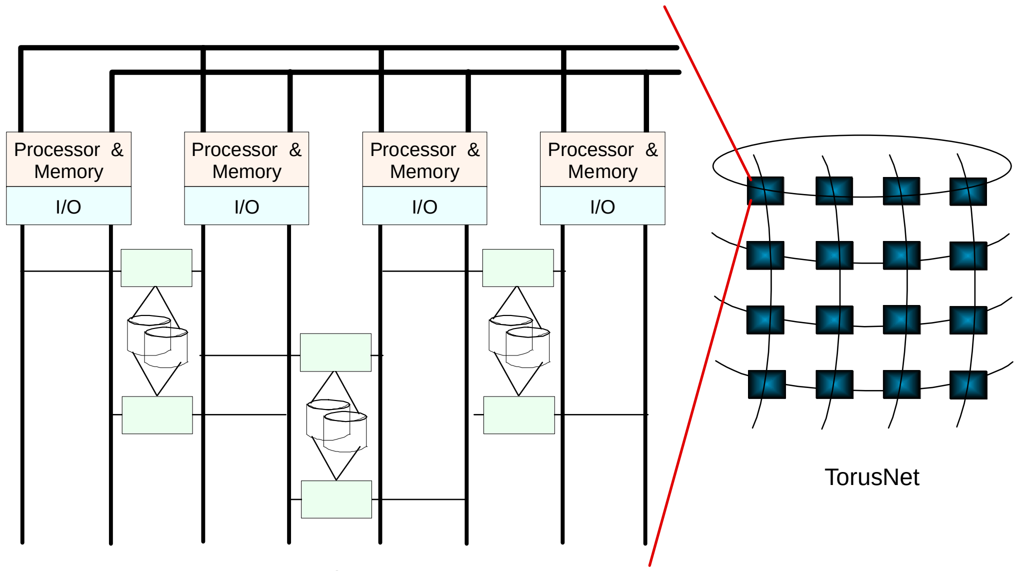

IBM Technical Support

Sun Starfire System

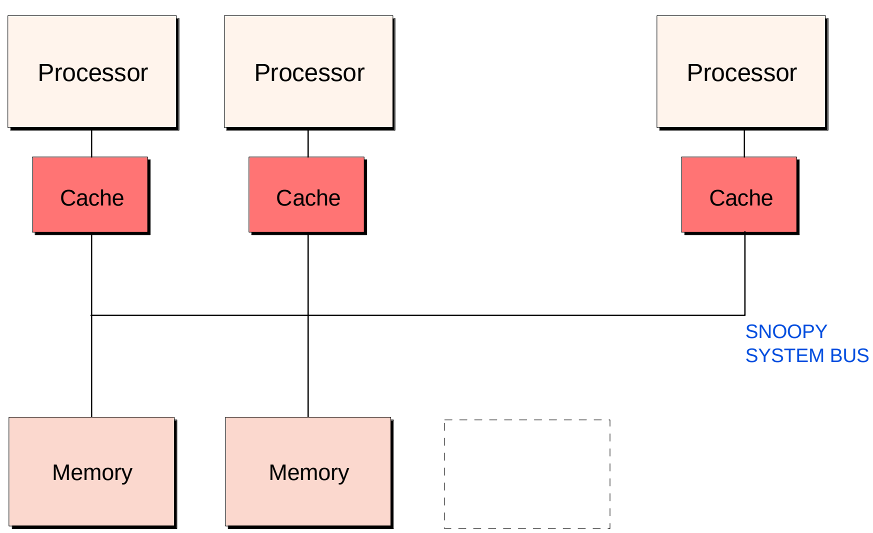

IBM Technical Support

Sun Interconnect

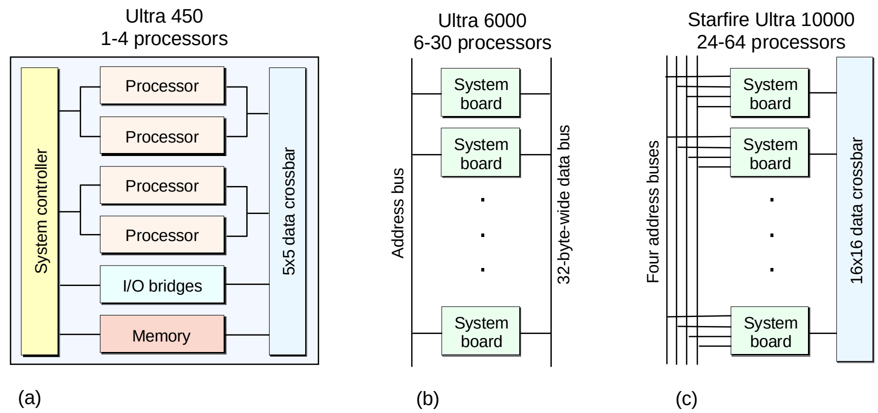

IBM Technical Support

IBM S/390 Parallel Sysplex

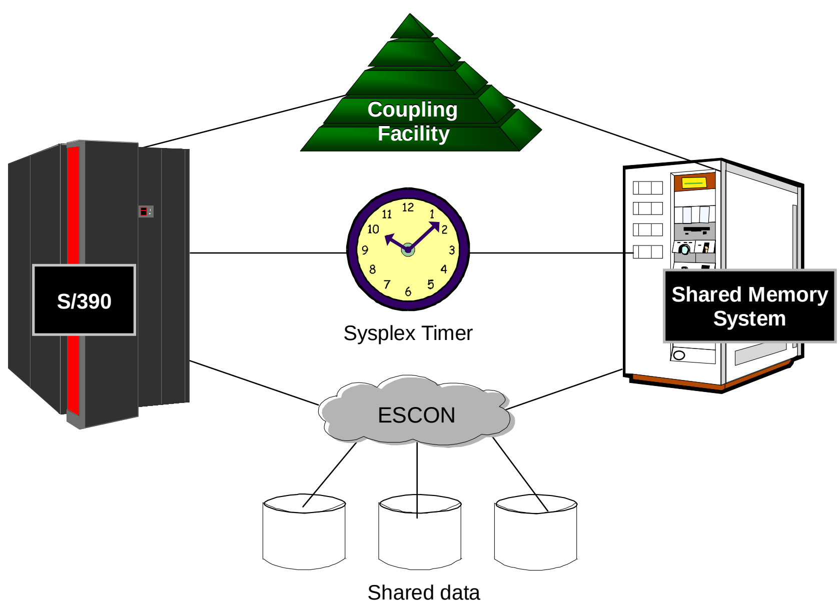

IBM Technical Support

RS/6000 SP Cluster

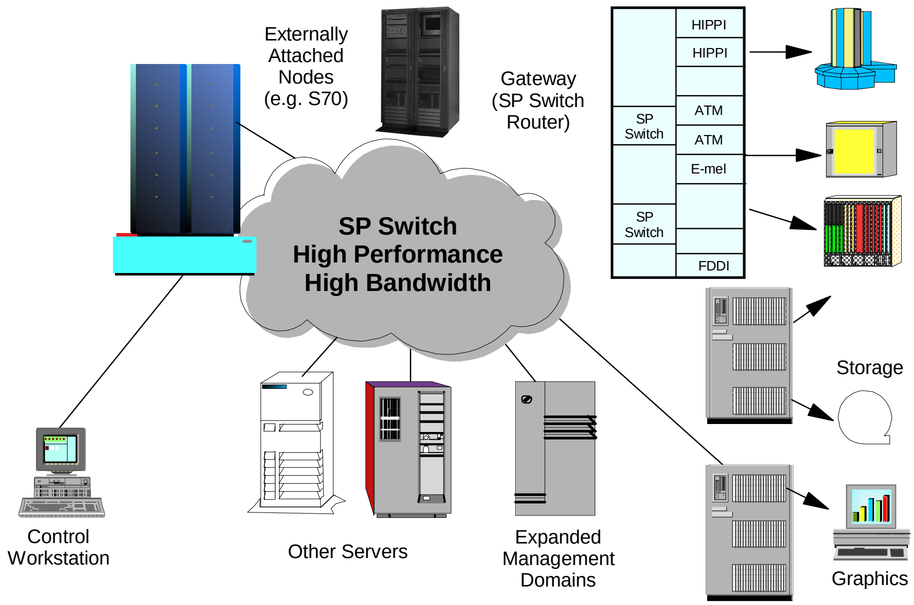

IBM Technical Support

Cluster Products

Vendor	Products
Amdahl	EnVista PS/R Cluster "Wolfpack"
Compaq	Recovery Server Option Kit
	Microsoft "Wolfpack"
DEC	TruCluster Availability Server
	TruCluster Production Server
	Digital Clusters for Windows NT
HP	HP High Availability Enterprise Clusters, MC/Service Guard,
	NetServer System Cluster; runs "Wolfpack"
Sequent	Symmetry 5000 ptx/CLUSTERS
Sun	Ultra Enterprise Cluster PDB Server,
	Ultra Enterprise Cluster HA Server

IBM Technical Support

IBM

Vendor	Products
IBM	AS/400 Opticonnect/400
	HACMP/6000 (High Availability Cluster Multiprocessor)
	PC Server Cluster, runs "Wolfpack" or Vinca
	RS/6000 Scalable Parallel Systems

NUMA

NUMA - Is "numa" Inevitable?

IBM Technical Support
© Copyright IBM Corporation, 1999

Numa?

UMA: Uniform Memory Access

- Every processor has equal access, using normal loads and stores, to all of memory. This is the way memory appears in the standard SMP programming model.

NUMA: Non-Uniform Memory Access

- Every processor has access to all of memory using normal loads and stores. However, there is a noticeably different, often very noticeably different, delay depending on what parts of memory are accessed; hence "non-uniform."

NORMA: No Remote Memory Access.

- Processors cannot access other processors' memories by normal loads and stores, but instead must communicate by other means. For all practical purposes, this is another term for message-passing systems.

CC - NUMA

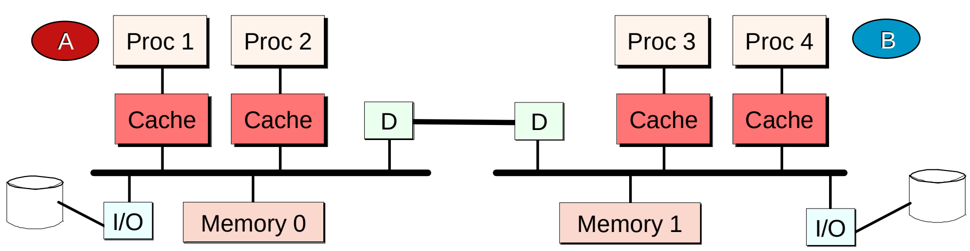

D = Attached memory bus

CC - NUMA (Remote access)

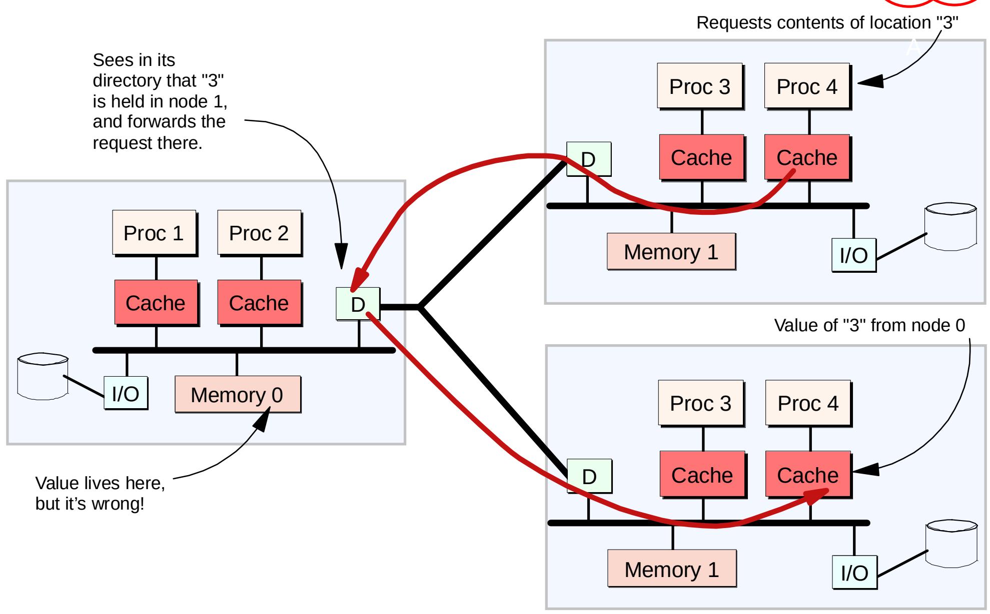

CC - NUMA (SCI)

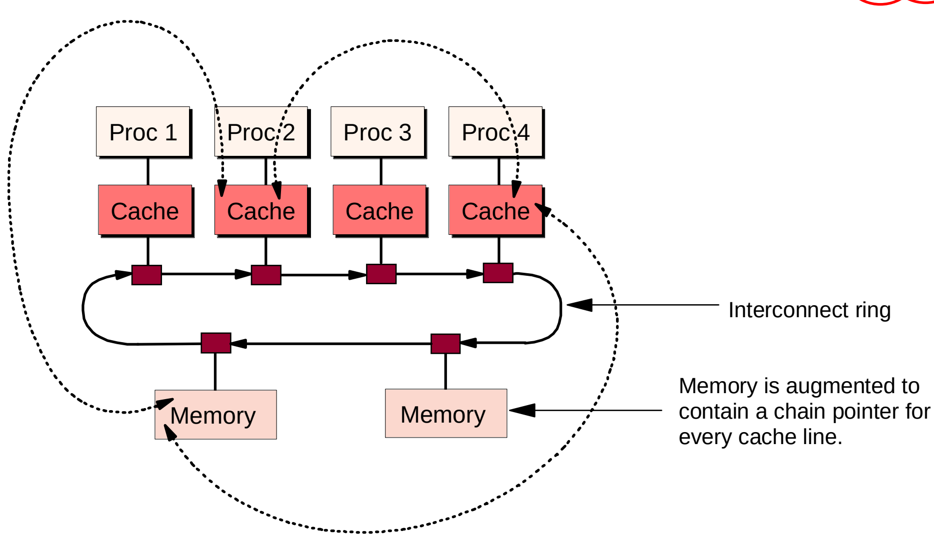

Industry Standard Benchmarks

IBM Technical Support

Popular Industry Benchmarks

Benchmark	Sponsor	Type	Range of Applicability
TPC-C	TPC	Moderately complex OLTP	Commercial OLTP
TPC-D	TPC	Decision support	Business Intelligence
TPC-W	TPC	OLTP web serving	e.Commerce
SPECint95	SPEC	Integer processing	Relative processor speed
SPECweb96	SPEC	Web serving	Web serving
Linpack	Jack Dongarra, University of Tennessee	Linear algebra	Technical computing
Baan	Baan	ERP	Baan ERP
SAP R/3 SD	SAP AG	ERP	SAP R/3 ERP
Notesbench	NotesBench Consortium	Groupware	Groupware

IBM Technical Support

Compute-Intensive Benchmarks

- LINPACK:
  - FORTRAN program for solution of a set of dense linear equations
  - Distributed by Jack Dongarra, University of Tennessee
- Livermore Loops:
  - Set of 24 FORTRAN DO-loops from Lawrence Laboratories
- Whetstones:
  - Synthetic Floating Point benchmark developed in 1976
  - Very sensitive to the size of cache
- Drystones:
  - This benchmarks spends a significant amount of time on string functions
  - Designed for measuring integer performance of small machines with simple architectures and instruction sets
  - Programs run in cache

Standard Performance Evaluation Corporation (SPEC)

- Open Systems Group
  - SPEC CPU95 (processor)
  - SPECjvm98 (Java)
  - SPEC SFS 2.0 (NFS)
  - SPECweb96 (Web serving)
- High-Performance Group
  - High performance benchmarks SPEChpc96
- Graphics Performance Characterization Group

SPECint Benchmarks

	SPEED	THROUGHPUT
Aggressive		
	SPECint95	SPECint_rate95
	SPECfp95	SPECfp_rate95
Conservative		
	SPECint_base95	SPECint_rate_base95
	SPECfp_base95	SPECfp_rate_base95

SPECint Benchmarks

- SPECint95:
  The geometric mean of eight SPECratios (one for each integer benchmark) when compiled with aggressive optimizations for each benchmark.
- SPECint_base95:
  The geometric mean of eight SPECratios (one for each integer benchmark) when compiled with the conservative optimizations for each benchmark.
- SPECfp95:
  The geometric mean of 10 normalized ratios (one for each floating point benchmark) when compiled with aggressive optimization for each benchmark.
- SPECfp_base95:
  The geometric mean of 10 normalized ratios (one for each floating point benchmark) when compiled with the conservative optimizations for each benchmark.

SPECint Ratios

- SPECint_rate95:
  - The geometric mean of eight normalized SPECrates (one for each integer benchmark) when compiled with aggressive optimization for each benchmark.
- SPECint_rate_base95:
  - The geometric mean of eight SPECrates (one for each integer benchmark) when compiled with conservative optimizations for each benchmark.
- SPECfp_rate95:
  - The geometric mean of 10 SPECrates (one for each floating point benchmark) when compiled with aggressive optimization for each benchmark.
- SPECfp_rate_base95:
  - The geometric mean of 10 SPECrates (one for each floating point benchmark) when compiled with conservative optimizations for each benchmark.

IBM Technical Support

SPECint_rate95 - Top Ten Results

Rank	Company	System	#CPUs	Processor	Result	Baseline
1	Compaq	AlphaServer GS140 6/625 Cluster	80	525 MHz 21264	15,332	12.425
2	SGI	Origin2000	128	250 MHz MIPS R10000R 3.3	14,703	14,611
3	Compaq	AlphaServer GS140 6/625 Cluster	64	612 MHz 21164	12,793	10,505
4	SGI	Origin2000	128	195 MHz MIPS R10000R 2.6	11,335	11,051
5	Compaq	AlphaServer 8400 5/625 Cluster	96	612 MHz 21164	11,330	9,546
6	Compaq	AlphaServer GS140 5/625 Cluster	48	525 MHz 21264	9,959	8,258
7	Sun Microsystems	Ultra Enterprise 10000	64	400 MHz UltraSPARC	9,181	7,095
8	Sun Microsystems	Ultra Enterprise 10000	128	250MHz UltraSPARC	8,381	7,339
9	Compaq	AlphaServer 8400 5/625 Cluster	64	612 MHz 21164	8,274	7,270
10	SGI	Origin2000	64	250 MHz MIPS R10000R 3.3	8,021	7,752

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

IBM Technical Support

SPECWeb Bechmarks

- SPECweb96 is a standardized benchmark for measuring basic Web Server performance.
- SPECweb96 features:
  - Standardized workload, agreed to by major players in WWW market
  - Full disclosures available available on this web site
  - Stable implementation with no incomparable versions

SPECWeb - Top Ten Results

Rank	Company	System	Result	HTTP Version	# CPUs
1	HP	9000/N4000	**24,139**	Zeus 1.3.3	8
2	IBM	9672-YX6	**21,591**	HTTP Server 5.1 for OS/390	10
3	IBM	RS/6000 S7A	**20,200**	HTTP Server 1.3.4	12
4	IBM	RS/6000 S7A	**19,264**	HTTP Server 1.3.4	12
5	Compaq	GS140 6/575	**14,263**	Zeus 1.3.0	10
6	HP	9000/V2250	**13,811**	Zeus 1.3.3	4
7	HP	9000/N4000	**13,051**	Zeus 1.3.3	4
8	HP	NetServer LXr 8000/500	**12,969**	Zeus 1.3.3	4
9	Siemens	Primergy 870-40	**12,126**	SISP 2.0	4
10	IBM	RS/6000 S70	**12,031**	Zeus 1.3.3	12

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

Webstones

- Works in simple fashion by generating traffic on a www HTTP protocol to create stressful conditions on the server
- Requires both the HTTP server under test and a UNIX workstation as the driver
- Workstation oriented

IBM Technical Support

NotesBench Benchmarks

- The NotesBench Consortium is an independent, non-profit organization dedicated to:
  - Providing Domino and Notes performance information to customers
  - Developing a base of performance information
- The consortium's primary focus is to:
  - Reduce customers' expenses associated with internal benchmark activities
  - Serve as an industry conduit for specifying future benchmarks, and
  - Ensure more rapid and optimized deployment of Domino and Notes.
- Lotus NotesBench for Lotus Notes R4 is a collection of benchmarks and documentation for evaluating theperformance of Notes R4 servers. The benchmarks (usually called workloads) model the behavior of Notes workstation-to-server or server-to-server operations.
- They return measurements to evaluate server performance in relation to the server system's cost of ownership.

IBM Technical Support

NotesBench Metrics

11 Measurements are defined:

- Idle Usage
- Web Buyer
- Web Walker
- Calendar and Scheduling
- Shared Discussion
- Groupware A
- Mail Only
- Mail and Shared Database
- Mail Routing Hub
- Replication Hub
- Cluster Mail and Shared Database

Common NotesBench Metrics

Idle Usage

This workload establishes an upper bound on the number of sessions,doing nothing, that a Notes server can support.

Mail Only

This workload simulates a server for mail users workload that models sites that rely only on mail for communication.

Mail and Shared Database

This workload simulates a server for active users who are performing only mail and simple shared database operations.

Lotus NotesBench Mail Only

Rank	Title	Notes Mark	Max_Users	Price_Per_Notes Mark	Price_Per_User	CPUs	Test_Date	Ver
1	IBM - RS/6000 S70 Server	**40,0757**	28,800	15.32	21.32	12	04-Jan-99	4.63
2	IBM - AS/400e 9406-S40	**35,979**	27,030	43.20	57.51	12	30-Oct-99	4.62
3	IBM - RS/6000 H70 Server	**15,372**	11,000	14.06	19.65	4	15-Apr-99	4.6
4	IBM - AS/400e 9406-S40	**13,857**	10,400	102.76	136.92	12	23-Jan-98	4.6a
5	IBM - Netfinity 5500 M20	**10,957**	8,250	4.15	5.51	2	24-Mar-99	5.0

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

Lotus NotesBench MailDB - Performance

Rank	Title	Notes Mark	Max_Users	Price_Per_Notes Mark	Price_Per_User	CPUS	Test_Date	Ver
1	Digital - Digital Server 7300/R	**10,864**	4,700	$8.50	$19.62	4	25-Jan-98	4.53
2	Digital - AlphaServer 4100 5/533	**10,864**	4,700	$8.50	$19.62	4	25-Jan-98	4.53
3	Digital - Digital Server 5300	**7,266**	3,100	$8.09	$18.97	2	25-Jan-98	4.52B
4	Digital - AlphaServer 1200	**7,266**	3,100	$8.09	$18.97	2	25-Jan-98	4.52B
5	IBM - Netfinity 7000	**6,294**	2,900	$11.38	$24.70	4	19-Sep-97	4.51

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

IBM Technical Support

© Copyright IBM Corporation, 1999

SAP R/3 Three Tier Architecture

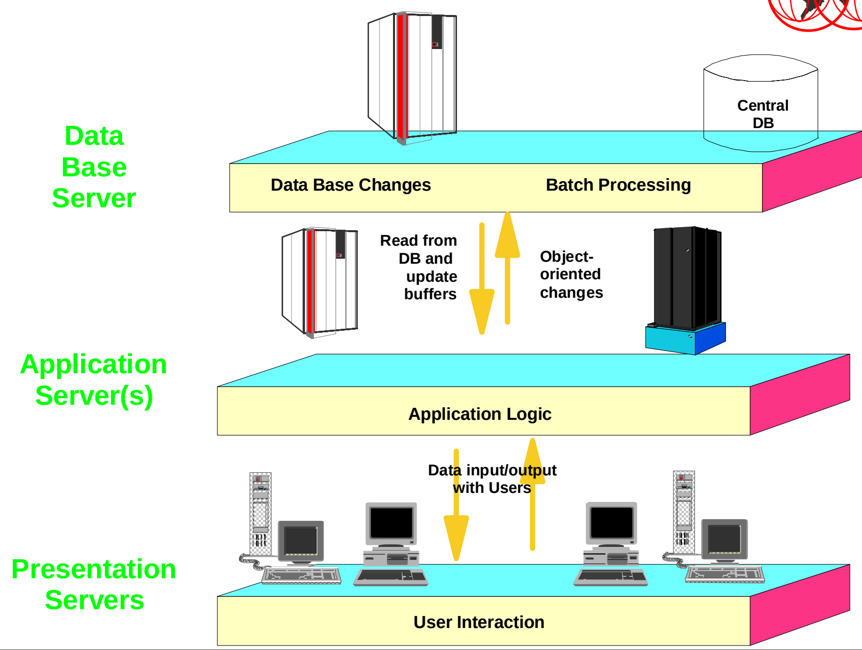

SAP R/3 DB Server on S/390

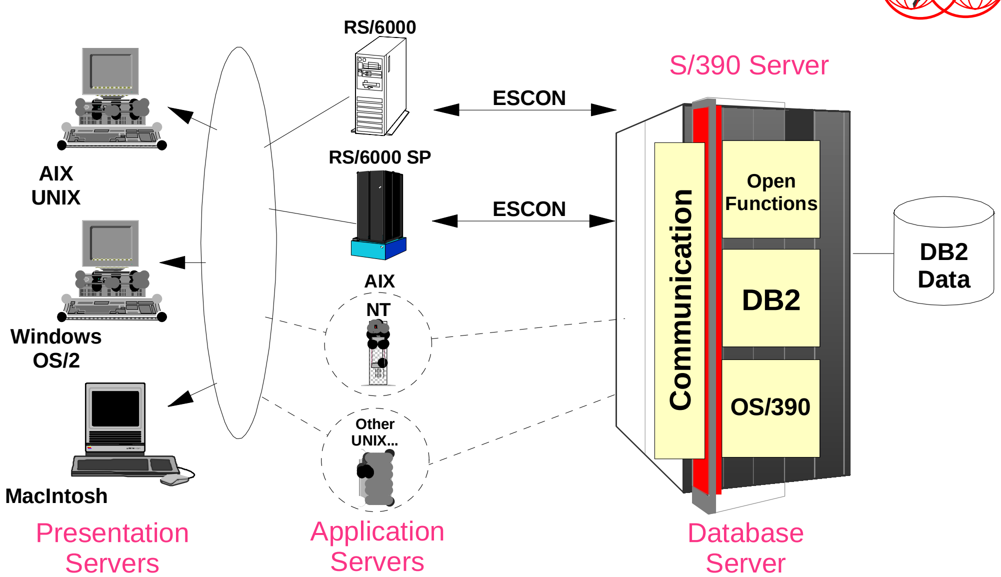

SAP R/3 Benchmarks

- SAP Standard Application Benchmark
  - Available since the R/3 software Release 1.1H (April 1993).
- Currently, there are seven different SAP application benchmark modules available
  - FI, MM, SD, PP, PS, WM and HR
  - The first six are dialog benchmarks,
  - Human Resources module is a batch benchmark
  - SD is the most common benchmark
- Results cannot be compared between two different modules.

SAP R/3 Metrics

- Number of benchmark users
- Average dialog response time (in seconds) - must be less than 2 seconds for the result to be valid.
- Configuration throughput - measured in SAPS (the number of dialog steps per hour)
- R/3 module type (e.g. SD, MM, etc.)

SAP R/3 Benchmark Types

- Central Configurations (Two-tier Client/Server)
  - Presentation and application layers are combined,
  - Database and application are on the same system.
- Three-tier Client ServerConfiguration
  - Architecture of the R/3 System is divided into the database,application and presentation layers.
  - Application and the database layers are on separate systems

SAP R/3 - Three Tier Top 10

Rank	Title	Module	Steps_Per_Hour	Users	Response_Time	Database	Utilization	Version	CPUs
1	Sun - Enterprise 10000 Server	SD	4,725,000	**14,400**	0.97	Oracle 8.0.4	82	3.1H	64
2	HP - HP 9000 Model V2250	SD	2,034,000	**6,750**	1.95	Oracle 8.0.4	99	3.1H	16
3	IBM - AS/400e 9406-S4 0	SD	2,016,000	**6,651**	1.88	DB2/400 V4R3	99	3.0F	12
4	HP - HP 9000 Model V2250	SD	1,958,000	**6,200**	1.4	Informix 7.30 FC6	92	3.1H	16
5	IBM - AS/400e 9406-S40	SD	2,013,000	**6,060**	0.84	DB2/400 V4R3	99	3.0F	12
6	Sun - Enterprise 10000 Server	SD	1,820,000	**6,030**	1.93	Oracle 7.3.3.3	82	3.1G	64
7	Compaq - AlphaServer 8400 6/575	SD	1,649,000	**5,312**	1.6	Informix 7.30 FC6	95	3.1H	10
8	HP - HP 9000 Model V2200	SD	1,602,000	**5,320**	1.95	ORACLE 7.3.3	99	3.1G	16
9	IBM - RS/6000 S70 SMP Server	SD	1,518,000	**4,960**	1.76	Oracle 8.0.5	98	4.0B	12
10	Bull Escala RL 470	SD	1,504,000	**4,961**	187	Oracle 8.0.5	99	4.0B	12

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

SAP R/3 SD Parallel Benchmark

- New version - for parallel and distributed environments
- SD Parallel benchmark is closer to production usage than strictly partitioned data benchmarks
- Round Robin method:
  - All business units spread across all DB servers
  - Significant inter-system data-sharing
  - Large number of inserts
  - Probably some more overheads than for a typical customer
- IBM first to show results
- Summary:
  - Different from normal SD Benchmark (more difficult)
  - More representative of parallel environments than other benchmarks
  - A little more conservative than normal SD Benchmark

SAP R/3 - Parallel Top 10

Platform	Steps_Per_Hour	Users	Resp Time	Database	CPU Utilization	SAP Version	CPUs
IBM 3 x 9672-RY5	2,487,000	**8,000**	1.58	DB2 for OS/390 V5.1	97%	3.1H	3x10
IBM 3 x 9672-RY5	2,128,000	**6,900**	1.69	DB2 for OS/390 V5.1	98%	3.1H	3x10
IBM 4 x 9672-RX3	1,065,000	**3,400**	1.50	DB2 for OS/390 V5	94%	3.0E	4x10
IBM 3 x 9672-RX3	819,000	**2,570**	1.29	DB2 for OS/390 V5	98%	3.0E	3x10
IBM 2 x 9672-RX3	565,000	**1,750**	1.16	DB2 for OS/390 V5	98%	3.0E	2x10
IBM 5 x RS/6000 SP	573,000	**1,700**	0.69	Oracle Parallel Server 7.3.2	83%	3.0D	5x8
IBM 5 x RS/6000 SP	510,000	**1,520**	0.74	Oracle Parallel Server 7.3	89%	3.0D	5x8
IBM 3 x RS/6000 SP	354,000	**1,050**	0.67	Oracle Parallel Server 7.3	83%	3.0D	3x8
IBM 2 x RS/6000 SP	256,000	**760**	0.68	Oracle Parallel Server 7.3	94%	3.0D	3x8

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

Transaction Processing Council

- October 1988 - Invitation to computer vendors to form a council to develop standard benchmarks
- November 1988 - TPC founded
- IBM one of founding members
- Each vendor has only one (equal) vote
  - PC vendor with one product has one vote
  - IBM with 4 server groups has one vote
- There is a set of fair rules of conduct associated with the use of TPC benchmarks
  - Failure to comply could result in expulsion

IBM Technical Support

TPC Benchmarks

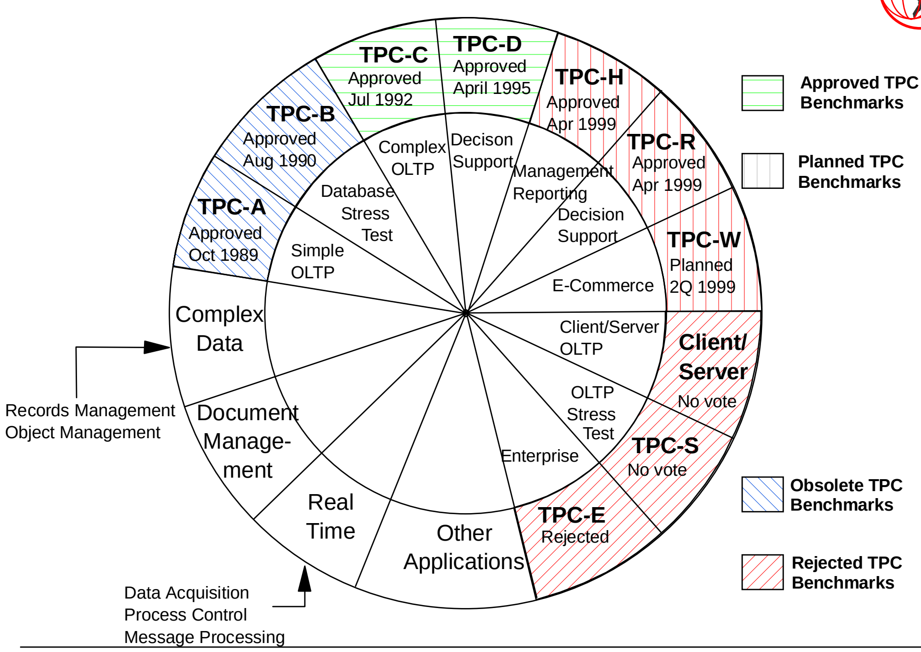

IBM Technical Support

TPC-A, -B and -C

Online transaction processing (OLTP) and database benchmarks:

TPC-A	Simple banking debit/credit - Online
TPC-B	Similar to tpc-A - Batch
TPC-C	Warehouse order entry and fulfillment (New version in 2000)

$$\$tpmC = \frac{\text{Computer system cost}}{\text{Throughput (tpmC)}}$$

TPC-D

Ad-hoc business queries:

Different scaling factors:

- 1 GB database
- 10 GB database
- 30 GB database
- 100 GB database
- 300 GB database
- 1TB database
- 3TB database

Numbers for different sized databases must not be compared.

Performance metrics:

- QppD@size
- QthD@size

Price-performance metric:

- Qphd@size

TPC-H and TPC-R

- TPC-D intended to represent ad-hoc query environment
- Vendors used knowledge of the benchmark to develop techniques for pre-building index structures that effectively reduced some queries to zero time
- Techniques are of use to customers when information required in known - useful for business reporting
- Two new benchmarks:
  - TPC-R for business reporting
  - TPC-H for ad-hoc query processing - harder than TPC-D and no pre-built indexes allowed
- TPC-R and TPC-H mandatory from June 1999
- Few results to date

TPC Reporting

1. Vendor completes benchmark
2. Hardware and software must be generally available within a defined period after the benchmark
3. Results reviewed by an independent auditor
4. Results documented in a Full Disclosure Report
5. Council members can challenge results within a 60 day period
6. Successful challenge means that vendor must withdraw result
7. Vendors often withdraw results because:
   - H/W or S/W does not meet GA date
   - Results look bad versus other vendor's results
   - Better results available

Fair Use Rules

- Must always quote Performance and Price/Performance metrics together
- Must not compare prices in different currencies (e.g. USD and JPY)
- Hardware and software must be available within a defined period (6 months)

TPC-A and TPC-C

Order entry and warehouse

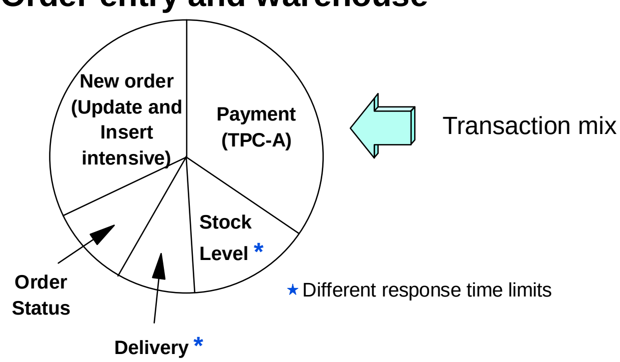

	TPC-A	TPC-C	Typical OLTP
Response Time (90%)	2 sec	5 sec	Sub second
Read/Write Ratio	1 : 2	1 : 1	7 : 1

TPC-C Options (1)

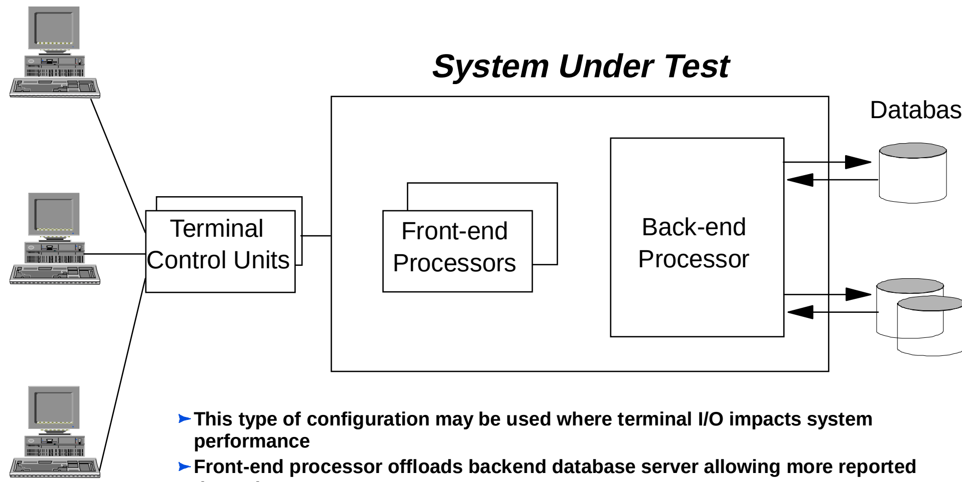

- This type of configuration may be used where terminal I/O impacts system performance
- Front-end processor offloads backend database server allowing more reported throughput
- Configuration may require more complex software and operational procedures
- "Front-end" processor (C/S) TPC implementations may have significantly greater throughput than "standalone" systems using the same back-end processor

IBM Technical Support

TPC-C Options (2)

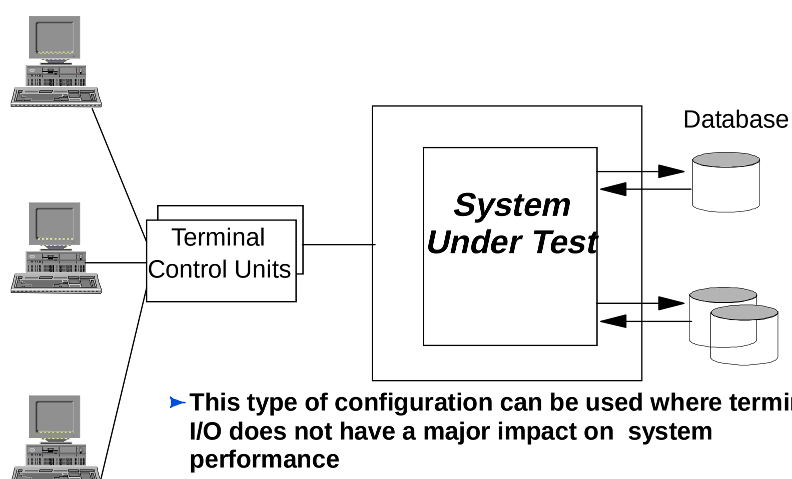

- This type of configuration can be used where terminal I/O does not have a major impact on system performance
- This implementation may be easier to configure and operate in real customer environments

IBM Technical Support

TPC-C Clusters Top Five - tpmC

Rank	Company	Config	tpmC	$/tpmC	Database
1	IBM	RS/6000 S7A (5 node x 12-way)	**110,434,10**	$122.44	Oracle8 v8.0.5
2	Compaq	AlphaServer 8400 (8 node x 12-way)	**102,541.85**	$139.49	Oracle8 v8.0.5
3	IBM	RS/6000 SP Model 309 (12 node x 8-way)	**57,053.80**	$147.40	Oracle8 Enterprise Edit'n 8.0.4
4	Sun	Ultra Enterprise 6000 c/s (2 node x 22-way)	**51,871.62**	$134.46	Oracle8 Enterprise Edit'n 8.0.3
5	NEC	Express 5800 HV8600 (4 node x 8-way)	**50,208.43**	$94.05	Oracle8 Enterprise Edit'n 8.1.6.0

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

TPC-C Non-Clusters Top Ten - tpmC

Rank	**Company**	**Config**	**tpmC**	**$/tpmC**	**Database**
1	Sun	E10000 (64-way)	**115,395.73**	$105.67	Oracle8 v8.1.5.1
2	Sequent	NUMA-Q 2000 E300 (64-way)	**93,900.85**	$131.67	Oracle8 Enterprise Edit'n 8.0.4
3	HP	HP 9000 V2500 (32-way)	**92,832.96**	$786.94	Oracle8 Enterprise Edit'n 8.1.5
4	Sun	Ultra Enterprise 6000 c/s (24-way)	**53,049.97**	$76.00	Sybase ASE 11.9.3
5	HP	HP 9000 V2250 (16-way)	**52,117.80**	$81.17	Sybase ASE 11.5 EBF 7817
6	HP	HP 9000 N4000 (8-way)	**49,308.00**	$56.67	Sybase ASE 11.9.3 EBF 8338
7	Sequent	NUMA-Q 2000 E300 (32-way)	**48,793.40**	$127.53	Oracle8 Enterprise Edit'n 8.0.4
8	IBM	AS/400e 9406-S40 (12-way)	**43,419.15**	$69.91	DB2 for AS/400 V4 Release 3
9	IBM	AS/400e 9406-S40 (12-way)	**43,169.85**	$128.91	DB2 for AS/400 V4 Release 1
10	Compaq	ProLiant 8000-500-2M (8-way)	**40,013.30**	$18.86	M'soft SQL Server 7.0 Ent'prise Edition

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

TPC-C Non-Clusters Top Ten - $ per tpm-C

Rank	Company	Config	$/tpmC	tpmC	Database
1	Compaq	ProLiant 5500-500-2M (4-way)	**14.62**	22,057.45	M'soft SQL Server 7.0 Ent Ed'tn
2	Compaq	ProLiant 5500 6/500 (4-way)	**15.11**	20,195.50	M'soft SQL Server 7.0 Ent Ed'tn
3	Compaq	ProLiant 5500R 6/500 (4-way)	**15.51**	20,190.50	M'soft SQL Server 7.0 Ent Ed'tn
4	Dell	PowerEdge 6350 (4-way)	**16.29**	23,187.90	M'soft SQL Server 7.0 Ent Ed'tn
5	Dell	PowerEdge 6300 (4-way)	**16.31**	23,187.90	M'soft SQL Server 7.0 Ent Ed'tn
6	Acer	AcerAltos 21000 (4-way)	**16.66**	23,235.57	M'soft SQL Server 7.0 Ent Ed'tn
7	Dell	PowerEdge 6350 (4-way)	**17.24**	23,460.57	M'soft SQL Server 7.0 Ent Ed'tn
8	Dell	PowerEdge 6300 (4-way)	**17.26**	23,460.57	M'soft SQL Server 7.0 Ent Ed'tn
9	Compaq	ProLiant 7000 (4-way)	**18.84**	22,478.90	M'soft SQL Server 7.0 Ent Ed'tn
10	Compaq	ProLiant 8000-500-2M (8-way)	**18.86**	40,103.30	M'soft SQL Server 7.0 Ent Ed'tn

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

TPC-D 300GB Results - QppD

Rank	TPC-D Power	Qppd	QthD	QphD	$/QphD	Database
1	NCR WorldMark 5200	**133,966.8**	13,756.2	42,928.80	440.00	Teradata V2R3.0
2	HP 9000/V2500	**56,246.7**	17,540.0	20,593.20	208.00	Oracle8 v8.1.5.1
3	Compaq AlphaServer GS140	**29,711.6**	4,868.7	12,027.40	192.05	Oracle8 v8.1.5.1
4	IBM RS/6000 SP Model 550	**10,469.6**	6,166.5	8,035.00	721.00	IBM DB2 Universal DB 5.2.0
5	Compaq AlphaServer GS140	**8,273.1**	3,487.8	5,371.70	999.00	INFORMIX IDS/AD/XP 8.21 UD2
6	HP NetServer LXr 8000	**8,124.3**	1324.7	3,280.60	162.00	Oracle8 v8.1.5.1
7	Sun Ultra Enterprise 10000	**8,113.2**	3,343.9	5,208.60	1,400.00	INFORMIX IDS/AD/XP 8.21
8	Sequent NUMA-Q 2000 (405 MHz)	**7,734.4**	3,055.4	4,861.30	740.00	INFORMIX IDS/AD/XP 8.21 UD2
9	SGI Origin 2000	**4,853.9**	1,861.9	3,006.20	973.00	Oracle8 v8.0.4
10	DG AViiON AV20000	**3,305.8**	1,277.7	2,055.20	1,320.00	Oracle8 v8.0.4

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

TPC-D 1TB Results - QppD

Rank	TPC-D Power	Qppd	QthD	QphD	$/QphD	Database
1	NCR WorldMark 5200 (44 x 4-way nodes)	**366,509.5**	24,142.0	94,065.30	273	Teradata V2R3.0
2	Sun E10000 (64 way)	**121,824.7**	10,556.3	35,878.10	283	Oracle8i v8.1.5.1.2
3	Sun E10000 (64 way)	**70,343.7**	7,658.6	23,210.70	438	Oracle8i v8.1.5.1.1
4	IBM Netfinity 7000 M10 (32 x 4-way nodes)	**36,872.0**	8,166.9	17,353.10	352	IBM DB2 UDB 5.2.0
5	Sequent NUMA-Q 2000 (64-way)	**27,441.8**	4,870.3	11,560.70	756	Oracle8i v8.1.5
6	Sun E10000 (64 way)	**27,024.6**	5,740.0	12,454.80	776	Oracle8i v8.1.5
7	IBM RS/6000 SP Model 550 (48 x 4-way nodes)	**19,137.5**	10,661.5	14,284.10	797	IBM DB2 UDB 5.2.0
8	Sun E6000 (44 x 24-way nodes)	**12,931.9**	5,850.3	8,698.00	1,353	INFORMIX IDS/ AD/XP 8.21

Ideas International (www.ideasinternational.com/benchmarks) as at 27th July 1999

Strengths and Weaknesses of Industry Standard Benchmarks

Strengths of Industry Standard Benchmarks

- They are standard tests
- They can compare different architectures
- They can compare different database management systems and versions (if sufficient data points exist)
- Tuning by vendors can flow through to useful performance for customers
- Multiple types of workloads can be used to compare to customer applications:
  - OLTP
  - ERP (Baan, SAP R/3, Peoplesoft)
  - Web Serving
  - Business Intelligence
  - Technical computing
- At the system level, they can reveal what is needed for an application in production

IBM Technical Support

Weaknesses of Industry Standard Benchmarks

- It is difficult to compare:
  - Different database management systems
  - Different versions of DBMSs
- High utilization in benchmarks is generally not achievable in production
  - Difficult to compare with systems that can achieve high utilizations in production
- The extent of tuning required for benchmarks is generally not practical for production systems
- Benchmarks use single workloads
  - Customers are increasingly requiring multiple concurrent workloads
- Scalability with the number of SMP engines is usually not visible because:
  - Cost of benchmarks to vendors, favour one measurement on their largest system
  - Even if there are multiple data points, each is a tuned environment and does not show production scaling

IBM Technical Support

Weaknesses of Industry Standard Benchmarks

- They run at the system level, and not at the processor level
- They do not test security capabilities
- They do not test data integrity functions
- They do not test backup/recovery or disaster recovery capabilities
- They do not show operational capabilities such as:
  - Ease of operations
  - Resilience to failure situations
  - Mixed workload management according to business priorities
  - On-line database re-organization capabilities

S/390 Strengths NOT Measured in Benchmarks

Strengths Not Measured by TPC

S/390 Function	Significance	Availability on UNIX
Memory management	S/390 system require less memory than UNIX databases. Benchmarks do not stress the paging system.	UNIX paging systems are typically not as good as S/390
I/O bandwidth	TPC benchmarks do not stress I/O systems.	UNIX I/O systems do not have as high a bandwidth as S/390
Data compression	Hardware data compression can reduce disk space requirements significantly, without the overhead found in software compression	Software compression available, but not widely used
DB2 Sort Assist, and Sorting instructions	Limited sorting in TPC benchmarks.	Not available on UNIX systems
Multiple concurrent workloads	Limited workload management needed in TPC benchmarks. Commercial systems demand workload management	Limited workload management available on most UNIX systems

IBM Technical Support

Strengths Not Measured by TPC

S/390 Function	Significance	Availability on UNIX
Fault tolerance	Not measured by any benchmark. Loss of a processor engine in UNIX usually causes the loss of the entire system	Limited fault tolerance on most UNIX systems, unless sepcial fault-tolerant designs
High availability	Not measured by any benchmark	Single S/390 systems more available than UNIX Clusters
Peripheral device performance	Benchmarks focus on total, highly tuned, system performance. Extra I/O devices often used to maximize throughput	S/390 I/O devices have superior performance and availability than most UNIX devices
Cluster performance, scalability and availability	Most benchmarks do not measure the capabilities of IBM Parallel Sysplex	Parallel Sysplex capabilities not available on any UNIX system
Systems management	Systems management functions not assessed by any benchmark	S/390 system mangement is superior to most UNIX systems

IBM Technical Support

Strengths Not Measured by TPC

S/390 Function	Significance	Availability on UNIX
Backup and recovery	Not measured in benchmarks. S/390 times are usually significantly faster than UNIX due to better I/O handling, better system design, and parallel handling of I/O	S/390 backup and recovery are generally acknowledged to be superior to UNIX systems
Data Integrity	Not measured in any benchmark. S/390 has protection mechanisms to prevent deliberate or accidental loss of data integrity	S/390 data integrity is superior to most UNIX systems
Security	Not measured in benchmarks. S/390 and OS/390 have earned the highest security ratings of any general purpose commercial operating system	S/390 security is superior to UNIX systems
Cryptography	Not measured in benchmarks. Growing importance for e-business. Software encryption can cripple UNIX systems	Not available in UNIX

Dynamic I/O System

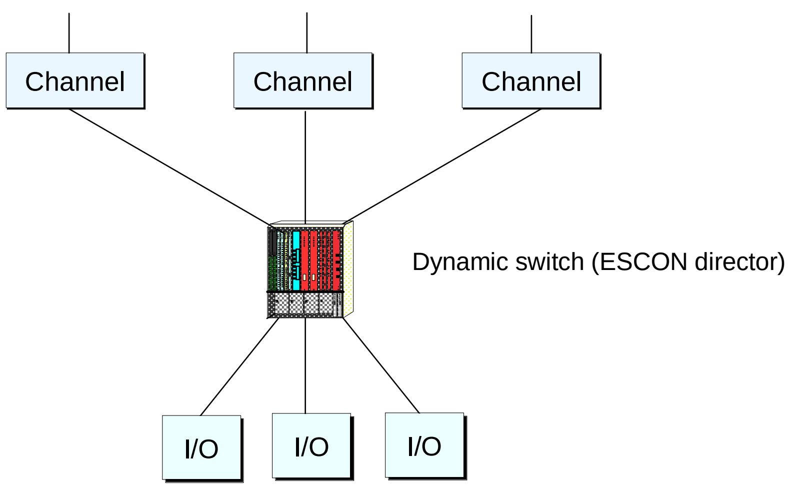

End of Presentation

IBM Technical Support