RMS 9 Replenishment Benchmark on IBM S80

Results and Implications for Future Sizings

Ralf Schmidt-Dannert,
IBM
ERP Advanced Technical Support

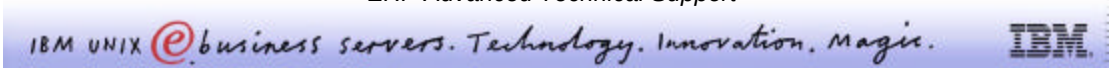

Agenda

Introduction

- IBM and Retek Relationship
- IBM S80 - an Overview

The Replenishment Benchmark

- RMS 9 - Overview
- What was benchmarked
- HW and SW configuration
- General Results
- The Details
- What did we learn

Contacts

Special Notices

This presentation was produced in the United States. IBM may not offer the products, programs, services or features discussed herein in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products, programs, services, and features available in your area. Any reference to an IBM product, program, service or feature is not intended to state or imply that only IBM's product, program, service or feature may be used. Any functionally equivalent product, program, service or feature that does not infringe on any of IBM's intellectual property rights may be used instead of the IBM product, program, service or feature.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. Sources for non-IBM list prices and performance numbers are taken from publicly available information including D.H. Brown, vendor announcements, vendor WWW Home Pages, SPEC Home Page, GPC (Graphics Processing Council) Home Page and TPC (Transaction Processing Performance Council) Home Page. IBM has not tested these products and cannot confirm the accuracy of 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.

IBM may have patents or pending patent applications covering subject matter in this presentation. The furnishing of this presentation does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of a specific Statement of General Direction.

The information contained in this presentation has not been submitted to any formal IBM test and is distributed "AS IS". While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. The use of this information or the implementation of any techniques described herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. Customers attempting to adapt these techniques to their own environments do so at their own risk.

IBM is not responsible for printing errors in this presentation that result in pricing or information inaccuracies.

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

All prices shown are IBM's suggested list prices; dealer prices may vary.

IBM products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Special Notices (Cont.)

Information provided in this document and information contained on IBM's past and present Year 2000 Internet Web site pages regarding products and services offered by IBM and its subsidiaries are "Year 2000 Readiness Disclosures" under the Year 2000 Information and Readiness Disclosure Act of 1998, a U.S statute enacted on October 19, 1998. IBM's Year 2000 Internet Web site pages have been and will continue to be our primary mechanism for communicating year 2000 information. Please see the "legal" icon on IBM's Year 2000 Web site (www.ibm.com/year2000) for further information regarding this statute and its applicability to IBM.

Any performance data contained in this presentation was determined in a controlled environment. Therefore, the results 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 these measurements will be the same on generally-available systems. Some measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this presentation should verify the applicable data for their specific environment.

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: ADSTAR, AIX, AIX/6000, AIXwindows, AS/400, C Set++, CICS, CICS/6000, DB2, ESCON, IBM, Information Warehouse, Intellistation, LANStreamer, LoadLeveler, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, OS/2, OS/400, OS/390, Parallel Sysplex, POWERparallel, RS/6000, S/390, Service Director, System/390, ThinkPad, TURBOWAYS, VisualAge. The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: AIX PVMe, AS/400e, DB2 OLAP Server, DB2 Universal Database, DEEP BLUE, e-business (logo), eNetwork, GigaProcessor, HACMP/6000, Intelligent Miner, Network Station, POWER2 Architecture, PowerPC 604, SmoothStart, SP, Videocharger, Visualization Data Explorer, WebSphere. A full list of U.S. trademarks owned by IBM may be found at http://iplswww.nas.ibm.com/wpts/trademarks/trademar.htm.

Microsoft, Windows, Windows NT and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States, other countries, or both. UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group. Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Lotus, and Lotus Notes are trademarks or registered trademarks of Lotus Development Corporation. Tivoli, TME, TME 10 and TME 10 Global Enterprise Manager are trademarks of Tivoli Systems, Inc. Other company, product and service names may be trademarks or service marks of others.

Notes on Benchmarks and Values

The benchmarks and values shown herein were derived using particular, well configured, development-level computer systems. Unless otherwise indicated for a system, the values were derived using 32-bit applications and external cache, if external cache is supported on the system. All benchmark values are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the following on the Web:

TPC	http://www.tpc.org	LINPACK	http://www.netlib.no/netlib/benchmark/performance.ps	Pro/E	http://www.proe.com
SPEC	http://www.spec.org	GPC	http://www.spec.org/gpc	NotesBench Mail	http://www.notesbench.org
VolanoMark	http://www.volano.com				

Unless otherwise indicated for a system, the performance benchmarks were conducted using AIX V4.2.1 or 4.3, IBM C Set++ for AIX/6000 V4.1.0.1, and AIX XL FORTRAN V5.1.0.0 with optimization where the compilers were used in the benchmark tests. The preprocessors used in the benchmark tests include KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors.

Notes on Performance Estimates

ROLTP

Relative OLTP (ROLTP) is an estimate of commercial processing performance derived from an IBM analytical model. The model simulates some of the system's operations such as CPU, cache, and memory. However, the model does not simulate disk or network I/O operations. Although the model uses general database and operating system parameters, the model does not reflect specific databases or AIX version or releases. Unless otherwise indicated for a system, the model assumes the use of 32-bit applications.

Unless otherwise indicated for a system, ROLTP is estimated only at the time the system is introduced. An IBM RS/6000 Model 250 is the baseline reference system and has a value of 1.0. Although ROLTP may be used to compare estimated RS/6000 commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying. For additional information about ROLTP, system performance, and benchmarks, contact your local IBM office or IBM authorized reseller or access the following Web sites:

SPEC	http://www.spec.org	Pro/E	http://www.proe.com
TPC	http://www.tpc.org	GPC	http://www.spec.org/gpc
Linpack	http://www.netlib.no/netlib/benchmark/performance.ps	NotesBench Mail	http://www.notesbench.org
VolanoMark	http://www.volano.com		

Background

IBM and Retek

The Retail Solution Requires A Wide Range of Hardware, Software, Services and Training

Platforms - Wide range of reliable, scalable, open servers supporting the Retek Retail Solution

People - Skills are required for domain expertise, project management, implementation, deployment, training

Process - A well-defined implementation methodology for new implementations and ongoing enhancement for the retailer

Proof - Experience always helps

Partnership - Strong technology, business and industry alignment

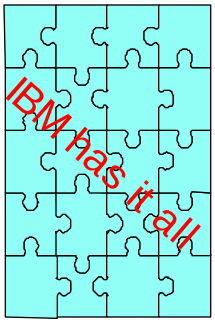

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- Retek and IBM have formed a unique partnership that will benefit customers immensely through better integration between their retail management systems, hardware and other software installations.
- The advancements are clear -- the best platforms for Retek are from IBM.
- Proof - The experience and knowledge available within IBM assure that proposed solutions are viable. IBM goes one step further in doing benchmarks not only to get high water numbers, but also to get better information about the key parameters of the Retek core applications.

IBM and Retek Have Recently Forged A Strategic Partnership

IBM-Retek Strategic Alliance Targeted to Generate $1 Billion Through Joint Solutions for Worldwide Retail Market

September 7, 2000 -- IBM (NYSE: IBM) and retail applications leader Retek (Nasdaq: RETK) today announced an expansion of their strategic alliance that will produce significant benefits for the worldwide retail industry, while generating revenues of more than $1 billion by 2003 for the two companies.

The new elements of the alliance call for both companies to jointly market, sell, and service a comprehensive retail e-business solution consisting of Retek applications and IBM software and hardware technologies, including DB2* Universal Database. IBM will promote Retek's products worldwide as the foundation for its retail e-market solution.

Retek has selected IBM Global Services as its leading consulting and implementation partner. In return, IBM will substantially increase its Retek-dedicated staff from 60 to more than 300 professionals, and will lead with Retek's applications in IBM's retail solution sales that are serviced by more than 200 dedicated account executives worldwide ....

Finding the Right Servers for Our Customer

Our first challenges are to characterize the performance of Retek core functions on Web Server platforms and develop sizing guidelines for our customers.

Finished

- RMS - POS Upload
- Retek Demand Forecasting
- RMS - Replenishment

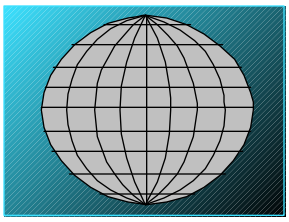

Planned

- Retek Distribution Management
- Retek Data Warehouse

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- The RS/6000 platform offers the performance, scalability, reliability and availability retailers need for their Retek installations.
- The RS/6000 line includes the world's most powerful single server - the RS/6000 S80, which has produced industry leading results on critical Sales Upload and Forecasting tests using Retek software solutions for the retail industry.
- Its unprecedented data transaction performance establishes the S80 as a preeminent platform for the retail sector.

The Server(s) for Retek Applications

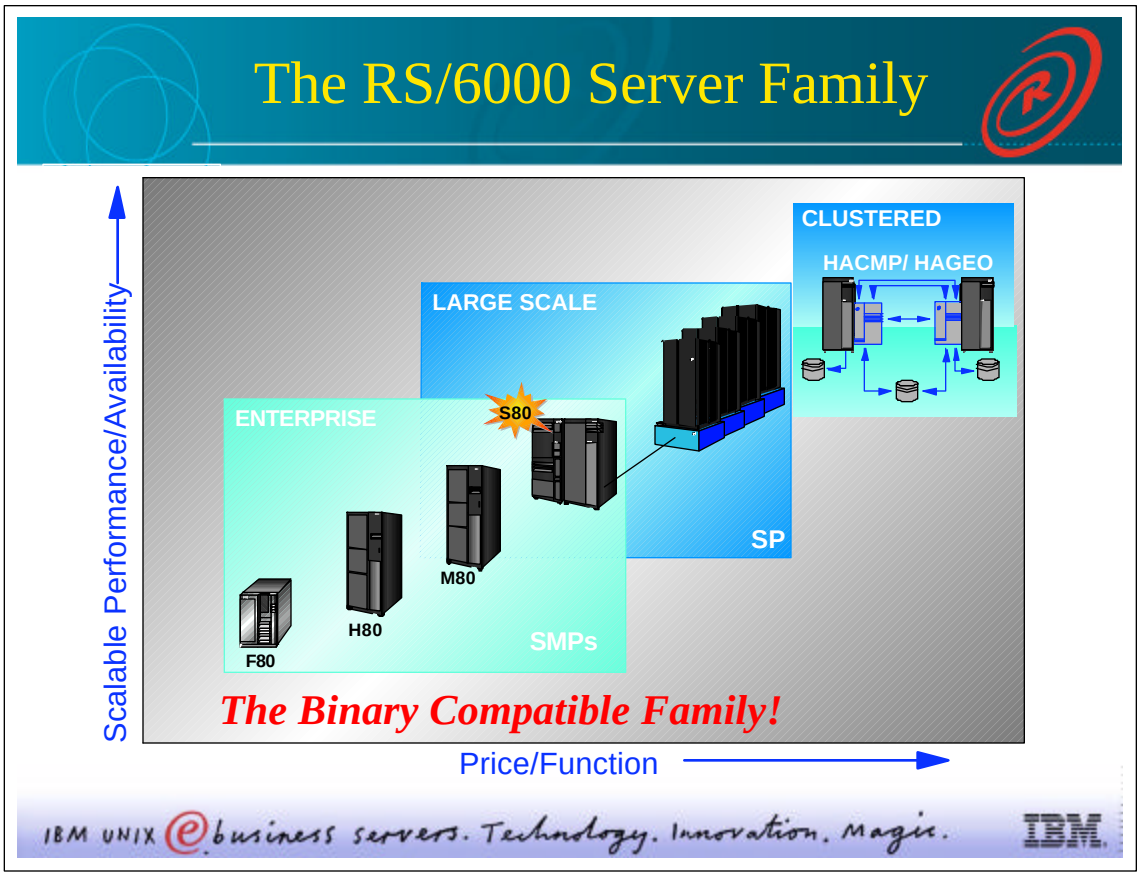

- Industry leading price/performance and scalability options
- Superior high availability features
- Key platforms for deploying multi-channel retailing applications
- Uniquely effective systems management technology

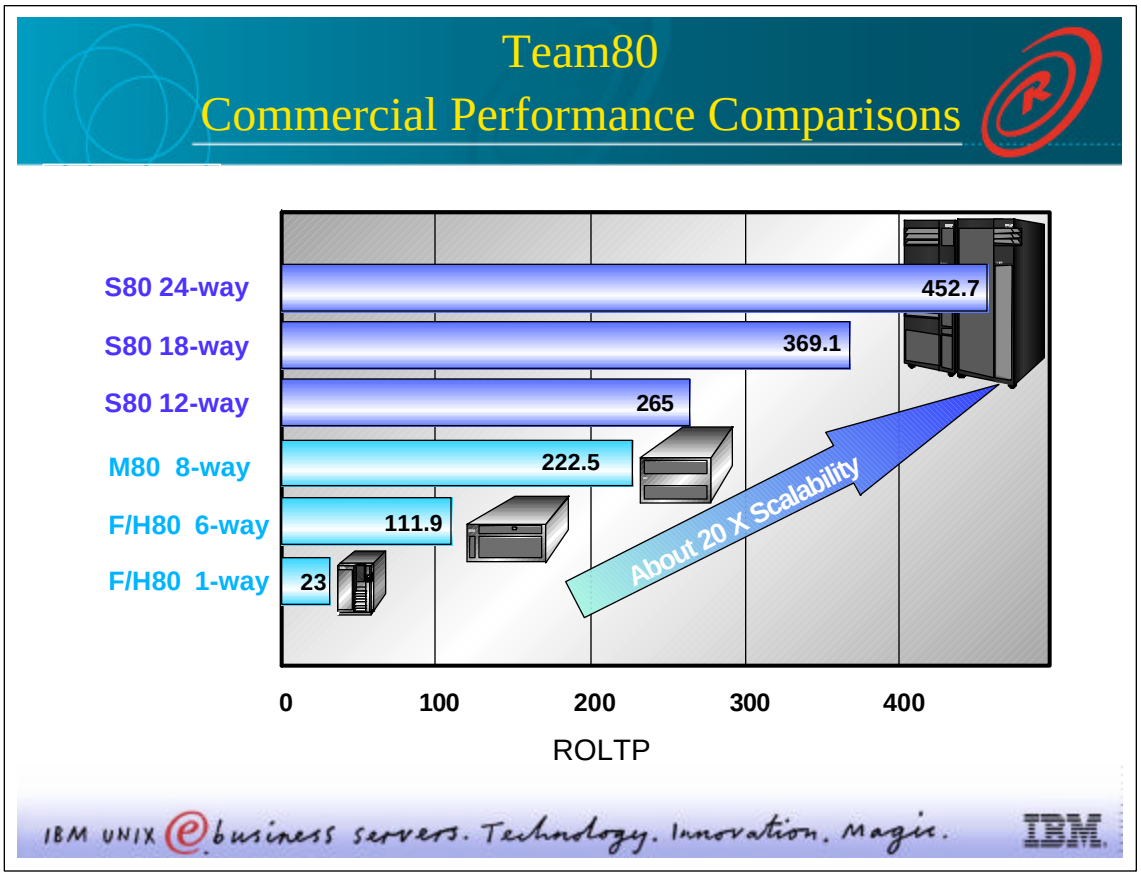

- Relative OLTP (ROLTP) is an estimate of commercial processing performance derived from an IBM analytical model
- An IBM RS/6000 Model 250 is the baseline reference system and has a value of 1.0

S80 Product Specifications

- **PowerPC RS III 450 MHz**
  - Copper technology
- **6, 12, 18, 24-way processors**
- **128KB/128KB L1 cache**
- **8MB L2 cache / CPU**
  - On-chip L2 cache controller
- **2GB - 64GB memory**
- **SMPs and I/O : 24GB/sec**
  - Dual buses (processors & I/O)
  - 2.4GB/sec/bus
- **Memory bus: 19.2GB/sec**
  - 4x 64-byte wide memory paths
  - 4.8GB/sec/bus
- **43.2 GB/sec total aggregate data switch bandwidth**

- **14-56 PCI slots**
- **Ultra SCSI 6-packs**
- **High availability I/O drawer**
- **Max disk bays : 48**
- **Max I/O drawer hot-swap disk: 873.6GB**
- **Max media bays : 8**
- **CEC upgrade from S70 Advanced**
- **AIX 4.3.3 required**
- **About 4X ROLTP over the S70 Advanced**
- **High Availability Solutions HACMP/HAGEO and GeoRM**
- **Dynamic CPU Deallocation**

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

S80 - UNIX Performance Leader

Single-server TPC-C (vs Sun & HP)

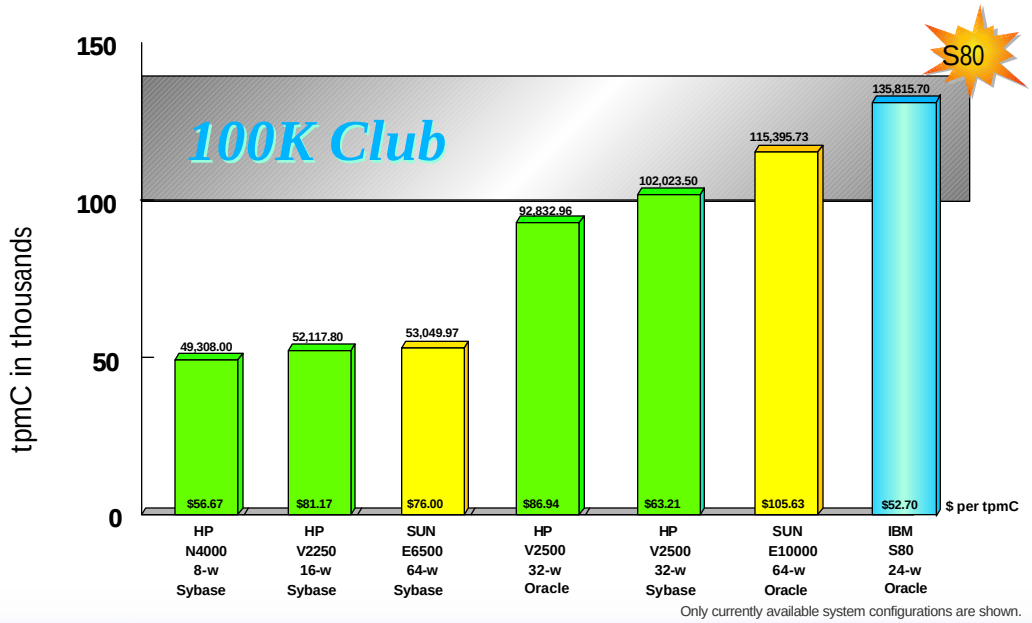

Only currently available system configurations are shown.

IBM UNIX @business servers. Technology. Innovation. Magic.

The Replenishment Benchmark

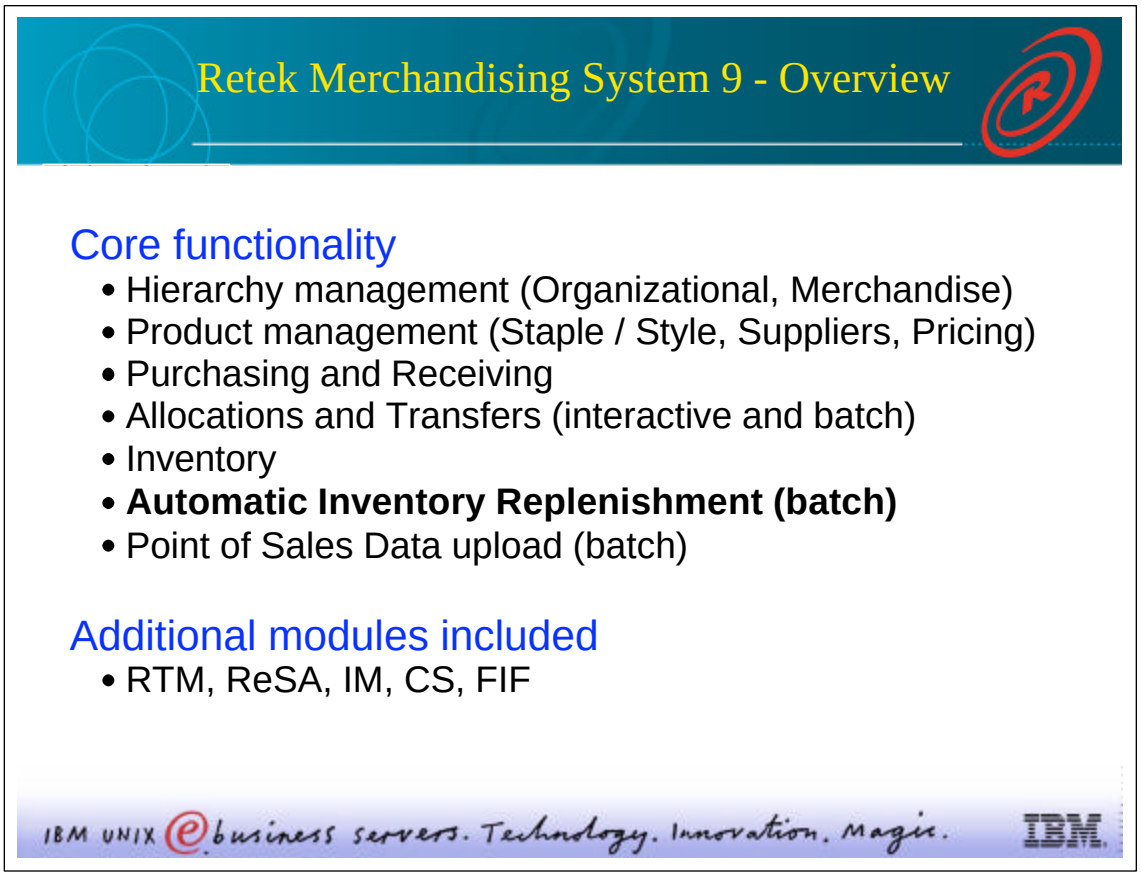

- Retek Trade Management (RTM)
- Retek Sales Audit (ReSA)
- Invoice Matching (IM)
- Retek Competitive Shopping (CS)
- Retek API to Oracle Financials (FIF)

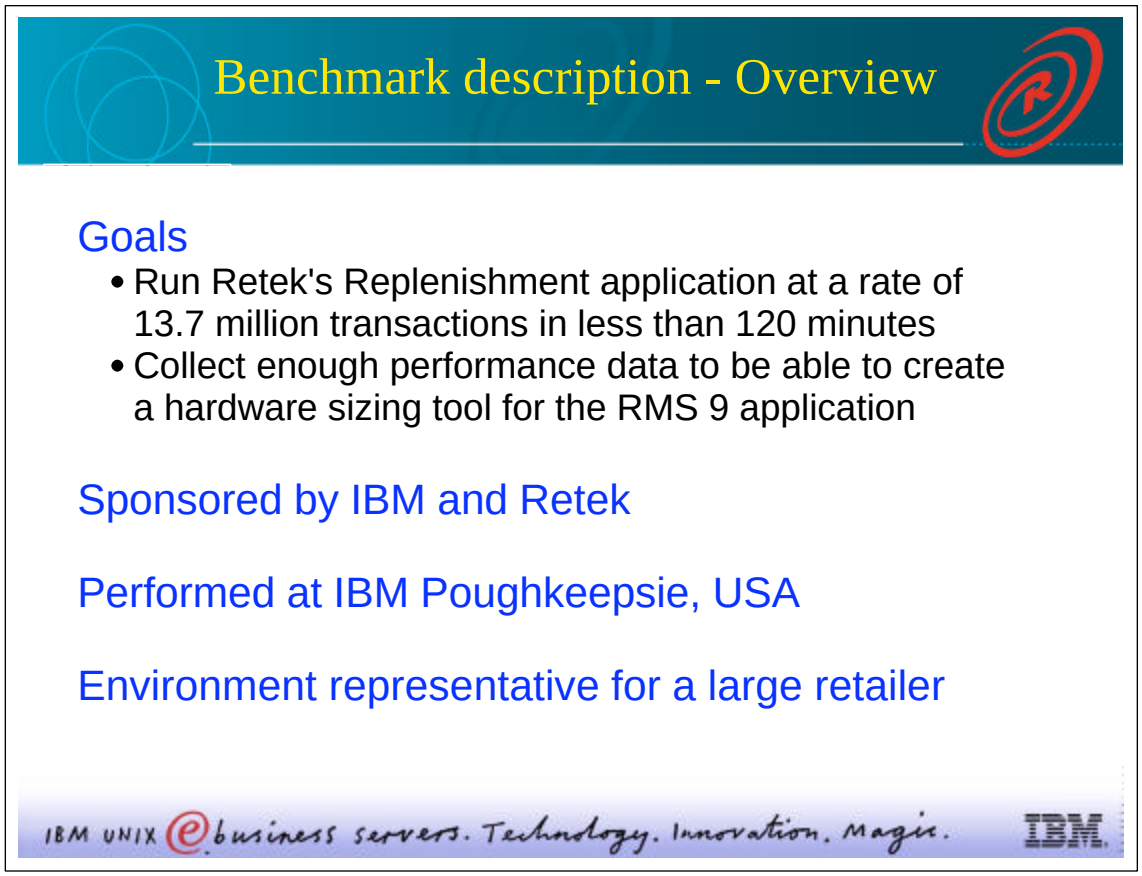

- There is no standard benchmark defined by Retek.

Benchmark description - Overview (cont.)

The following 3 major modules of the replenishment process were benchmarked

- RPLEXT
  The Replenishment Extract module maintains optimum stock levels of replenished staple/fashion stock items by determining the Recommended Order Quantity for a location (ROQ).

- RPLBLD
  The Replenishment Order Build module builds the actual purchase orders after all the store and warehouse ROQs have been determined and written to the temporary order table. The Replenishment Order Build (rplbld) module creates a new order for each supplier.

- REQEXT
  The Item Requisition Extract module creates transfers for all SKU/store records for styles/staple SKUs that are being replenished, where the SKU/store is active and its stock category is Warehouse Stocked (W).

SW configuration

- AIX 4.3.3
- ORACLE 8.1.6
  - 32bit and 64bit version used
  - table partitioning used
- RMS 9 - prerelease + patches

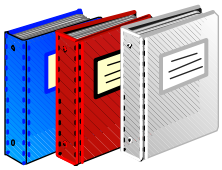

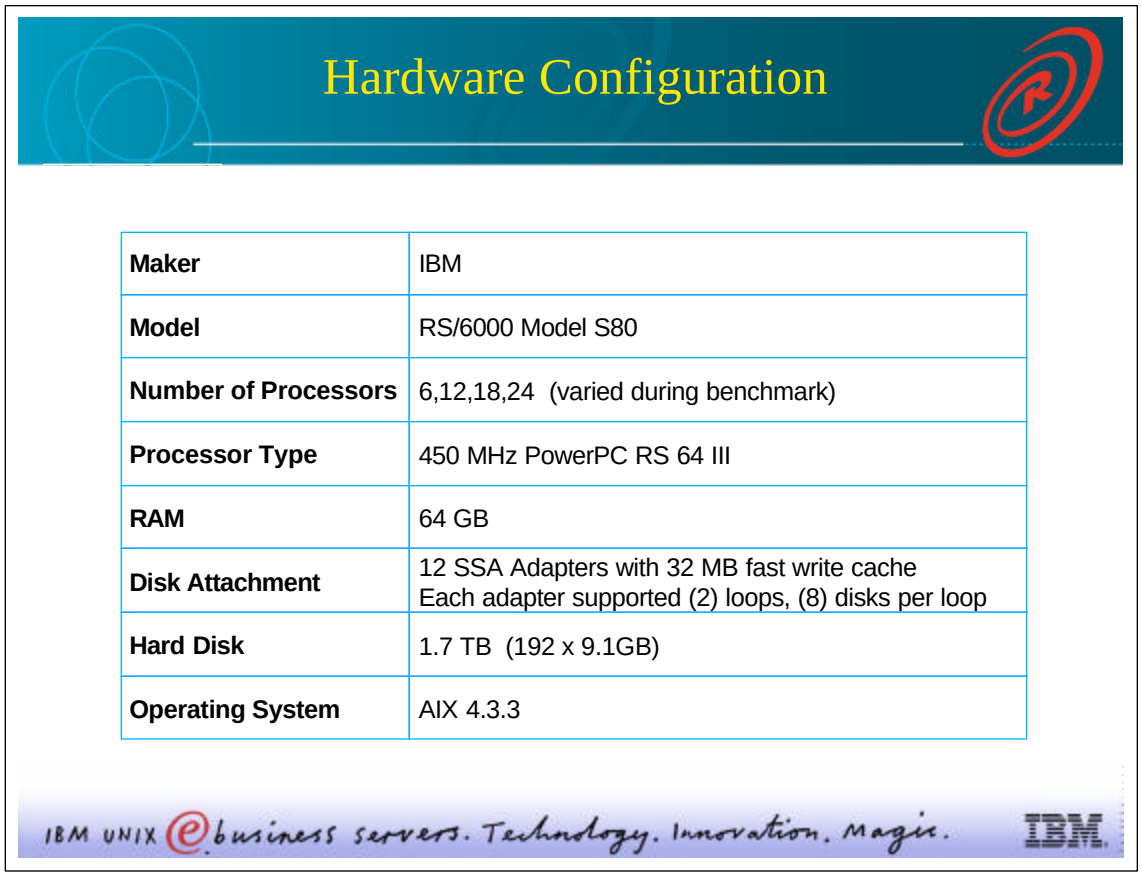

- Only a part of the available memory was realy used for the application during the benchmark - see memory chart later.

Retail Configuration

- 5 Warehouses
- 200 Stores
- Maximum of 67,200 SKUs / store
- 7 different workloads (WL1 - WL7) defined - active SKUs to be evaluated per location during one replenishment run

Replenishment Settings - Baseline

- Min/Max Method used
- Due Order Processing was not used
- Scaling was not used
- 9,280,000 SKU/Store Items set as Warehouse Replenished (100% Generating Transfers)
- 4,160,000 SKU/Store Items set as Cross-Dock Replenished (100% generating 232,000 Orders and 4,160,000 Allocations)
- 232,000 SKU/Warehouse Items for Vendor Replenishment (50% actually generated Orders)

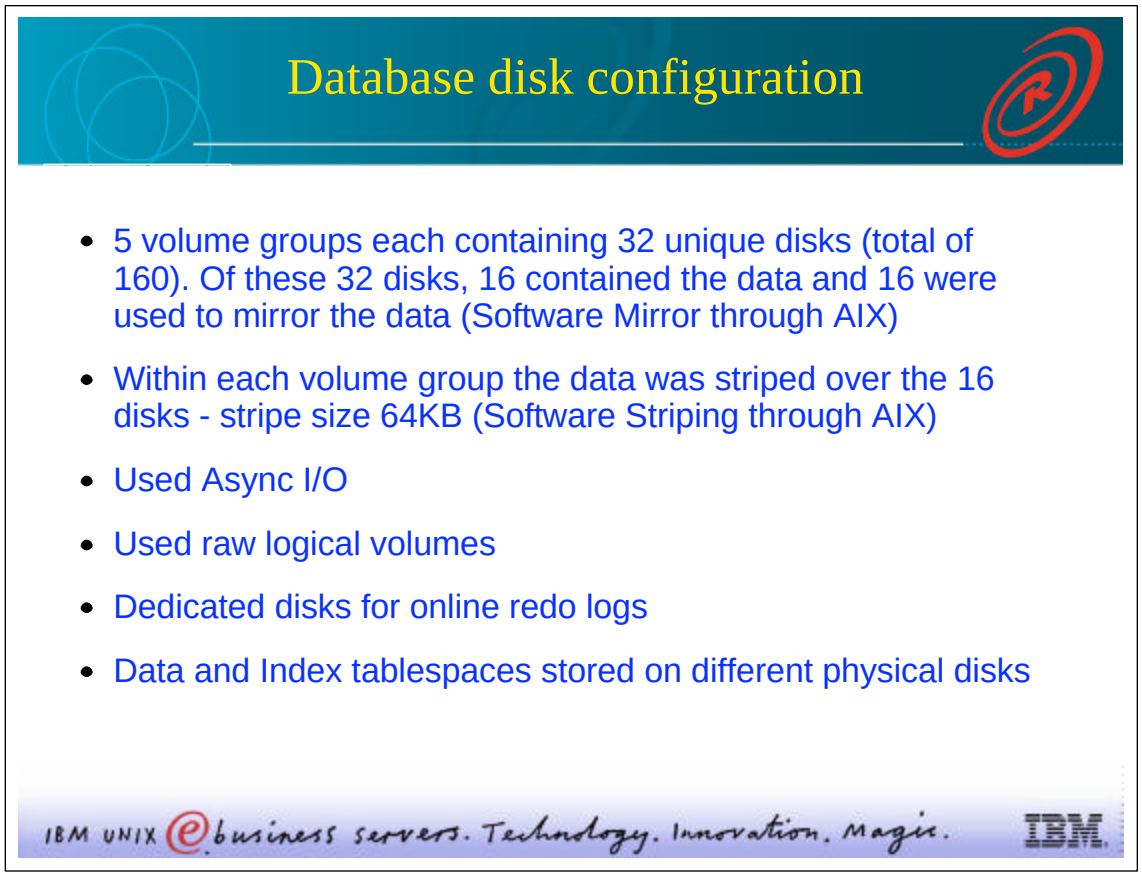

- VG #raw log. vol. contents
- retek1 10 online redo logs, roll back segments, system TBSPC
- retek2a 36 TBSPC: TMP, RETEK_DATA*
- retek2b 36 TBSPC: TMP, RETEK_DATA*
- retek3a 29 TBSPC: RETEK_INDEX*
- retek3b 29 TBSPC: RETEK_INDEX*
-
- 3 redo log groups
- size of one redo log file - 4GB
- no log archiving
- circular logging
-
- DB parameters:
- shared_pool_size: 104,857,600
- db_block_buffers: 190,000 x 8192
- log_buffer: 62914560
- ----------
- SGA: ~1.7GB

Benchmark Parameters

Number of Processors

- 6,12,18,24

Number of parallel threads

- # of processors + 1,....,+10

Number of Active SKUs / Location to be evaluated during each replenishment run

Module	WL1	WL2	WL3	WL4	WL5	WL6	WL7
rplext	20,800	18,200	15,200	12,200	9,200	6,200	3,200
reqext	46,400	40,600	33,800	27,200	20,400	13,800	7,200
total SKU x Store combinations	13.7M	12M	10M	8M	6M	4M	2M

WL stands for workload

High Water Result

Parameters	WL1
Item/Location Combinations	13,776,000
Active Item/Location	13,776,000
Percentage Active	100%
Transfers generated	9,280,000
Allocations generated	4,160,000
Orders generated	348000
Batch Process	**Time (mins)**
rplext	13:53
rplbld	7:36
reqext	38:04
Total	**59:33**

ORACLE 8.1.6, 32Bit, thread optimized

IBM UNIX @business servers. Technology. Innovation. Magic. IBM

S80 Scalability

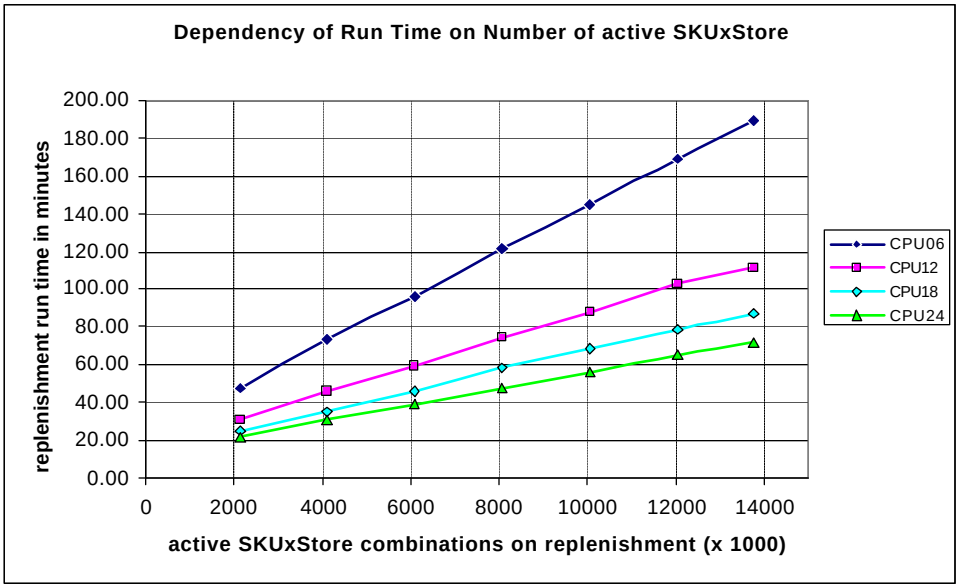

ORACLE 8.1.6, 32Bit; not thread optimized

IBM UNIX @business servers. Technology. Innovation. Magic. IBM

- This graph shows that the replenishment run time is directly linear correlated to the number of active SKU x Store combinations to be evaluated during the replenishment run.

S80 Scalability (cont.)

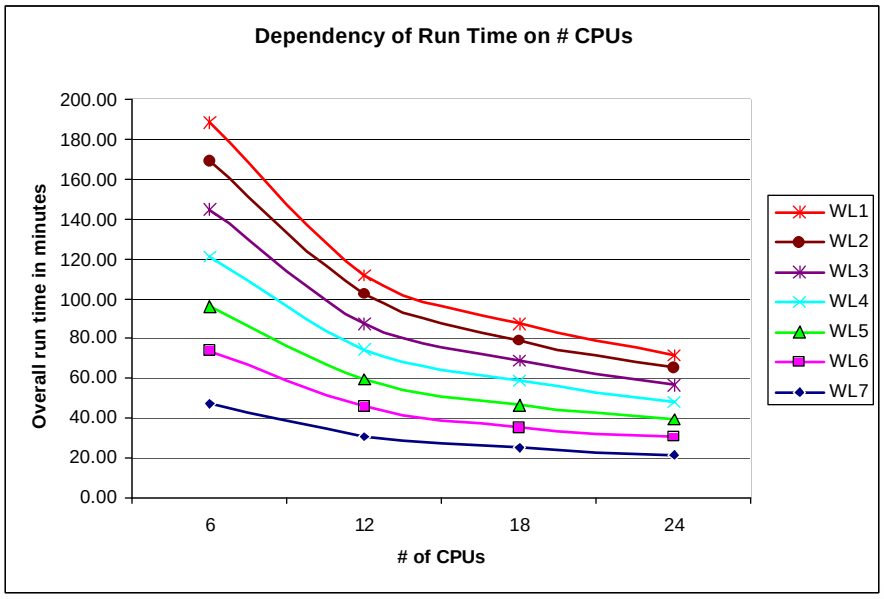

ORACLE 8.1.6, 32Bit; not thread optimized

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- Data shown is for ORACLE 8.1.6, 32 bit
- #threads: rplext = #CPU+1, rplbld = #CPU, reqext = #CPU+1
- thread optimization is more effect with higher number of CPUs e.g. 71min -> 59min for CPU24, WL1
- The small gain for low loads (WL7, WL6,...) is explained when you look at CPU utilization for that runs. 6 CPUs are about 100% busy, but for e.g. 24 CPUs have idle times from 20% to 40% for WL7.

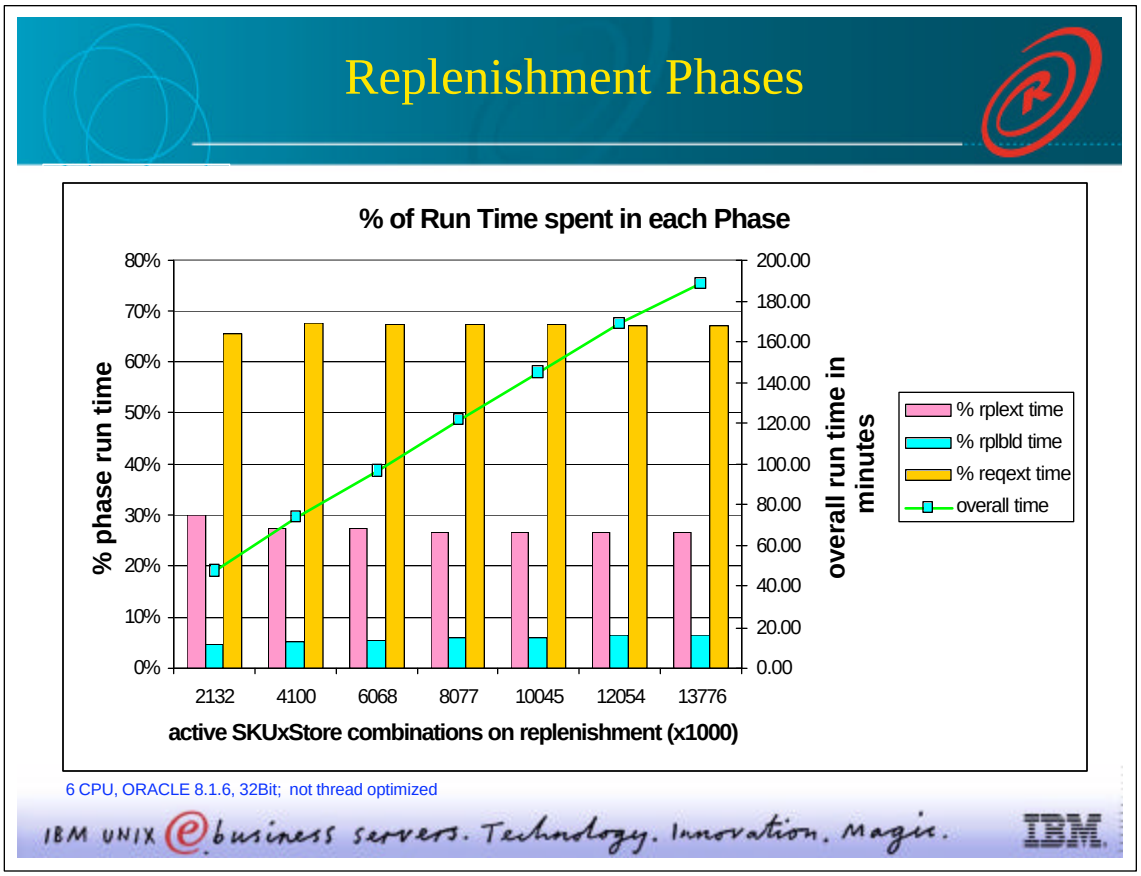

- rplext percentage of overall run time is about constant. 24CPU - 23%, 18CPU - 24%, 12CPU - 25.5%, 6CPU - 27%.
- rplbld percentage of overall run time is growing. There seems to be a non-linear correlation between the numbers of orders generated and the required run time to do so. Another point to look at is the heavy I/O load generated by this phase, which results in significant I/O wait times. 24CPU - 5.5%-10%, 18CPU - 4.8%-8.7%, 12CPU - 4.6%-7.9%, 6CPU - 4.4%-6.3%
- reqext percentage of overall run time is about constant. 24CPU - 68%,18CPU - 68%, 12CPU - 68%, 6CPU - 68%

Threads - Impact on Phases

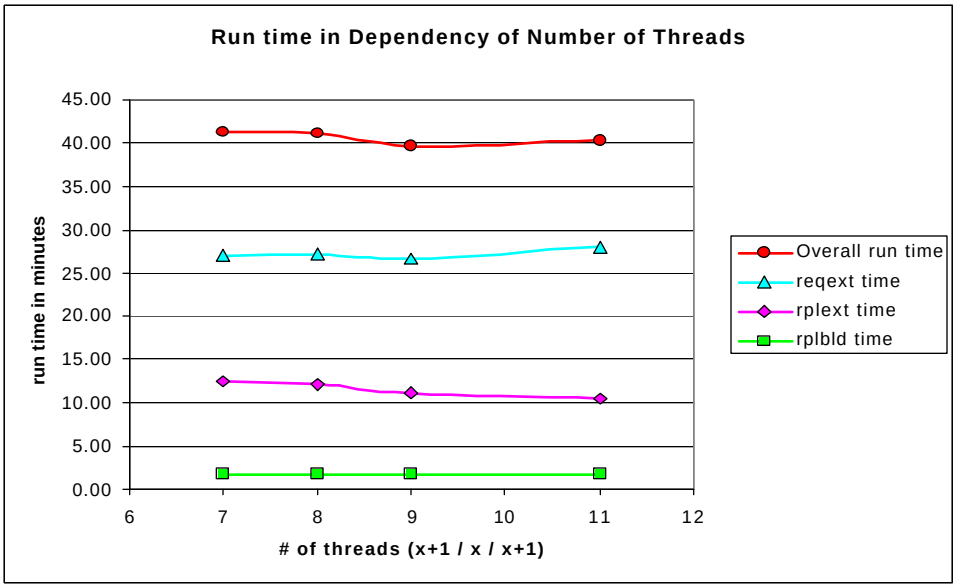

6 CPU, ORACLE 8.1.6, 32Bit, WL7

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- Run time is dependent on number of threads
- The optimum number of threads is different for each phase
- Maximum reduction in run time in this 6 CPU szenario is 5%. Would be more with higher workload.
- rplext was run with x + 1 threads (8-12)
- rplbld was run with x threads (7-11)
- reqext was run with x + 1 threads (8-12)

Threads - Impact on Run Time

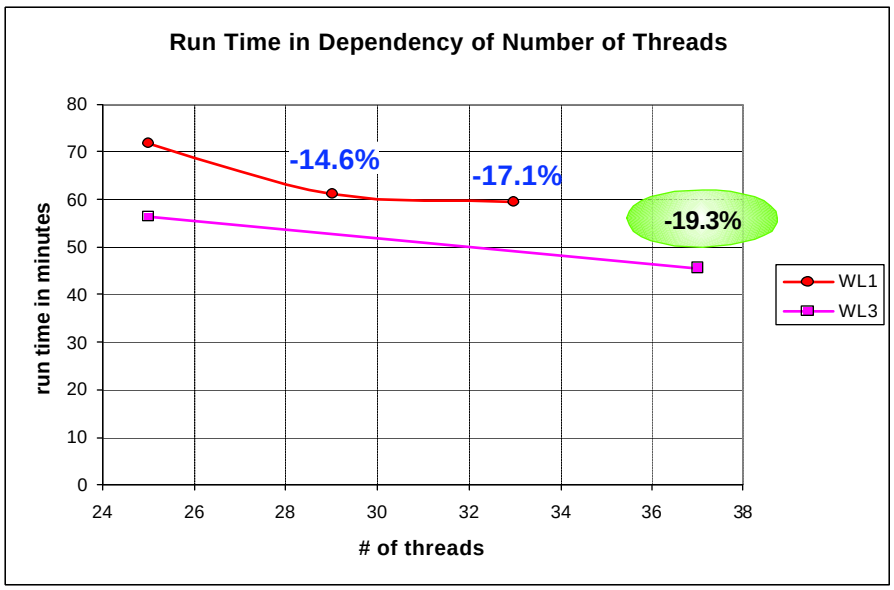

24 CPU, ORACLE 8.1.6, 32Bit

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- WL1: rplext- x threads, rplbld - x-1 threads, reqext - x threads
- WL3: Higher value: rplext- 25 threads, rplbld - 24 threads, reqext - 25 threads
- WL3: Lower value: rplext- 37 threads, rplbld - 33 threads, reqext - 33 threads
- WL1 - 13.7M, WL3 - 10M active SKUxStore combinations

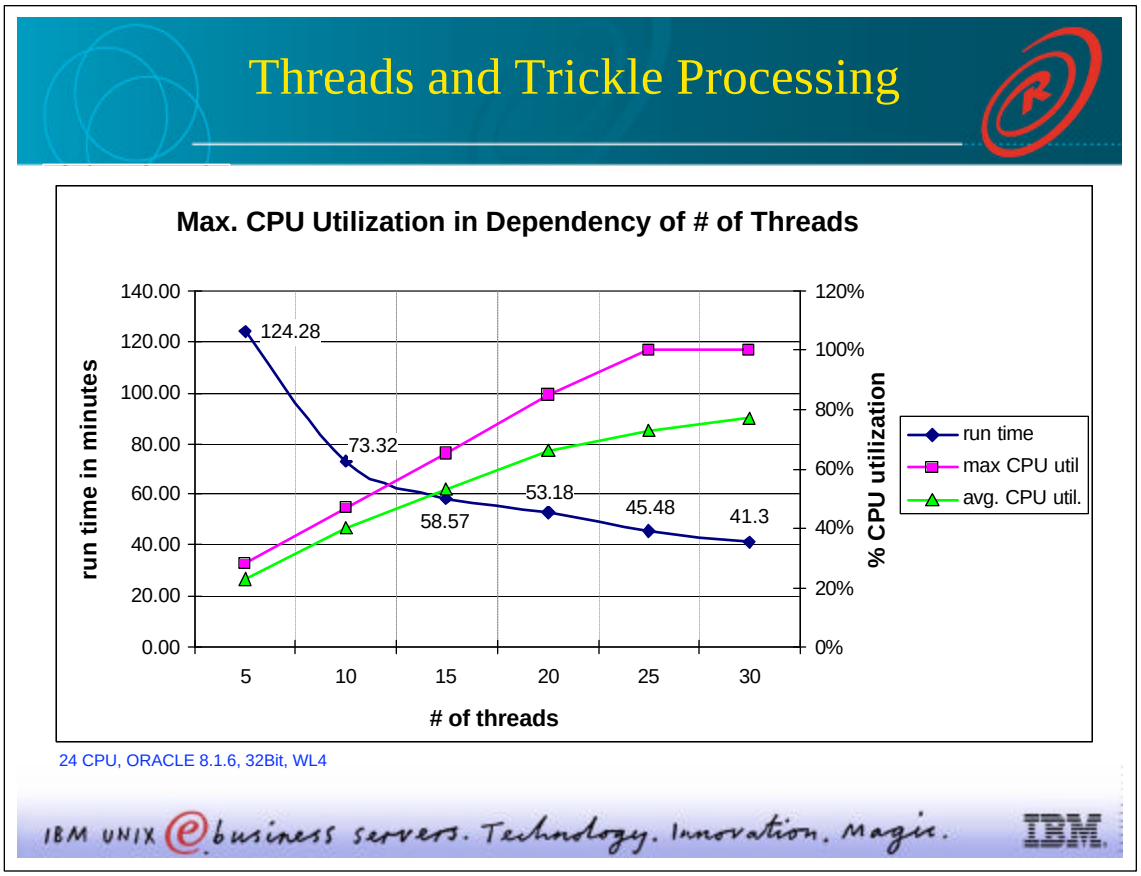

- WL4 -- 8M active SKUs
- The sharp decline in the run time in correlation to the number of threads ends at about 12 threads. The reason is, that each thread has a connection to an ORACLE thread, which satisfies the threads database requests. At about 24 = 12 + 12 threads we reach the point where we have one thread per CPU. Adding more threads still reduces the run time, because of I/O waits, but the reduction is much smaller and will eventually level off and then even reverse.
- The run time reduction is mostly seen in phase reqext and some in rplext. The run time of phase rplbld is about constant when using more than 12 threads.
- The average CPU utilization stay's well below 100% because this average also includes the transition times between the 3 phases and the ramp-up and ramp-down phases. The threads are not started all at once, but one after each other with a sleep time in between.

Required Memory

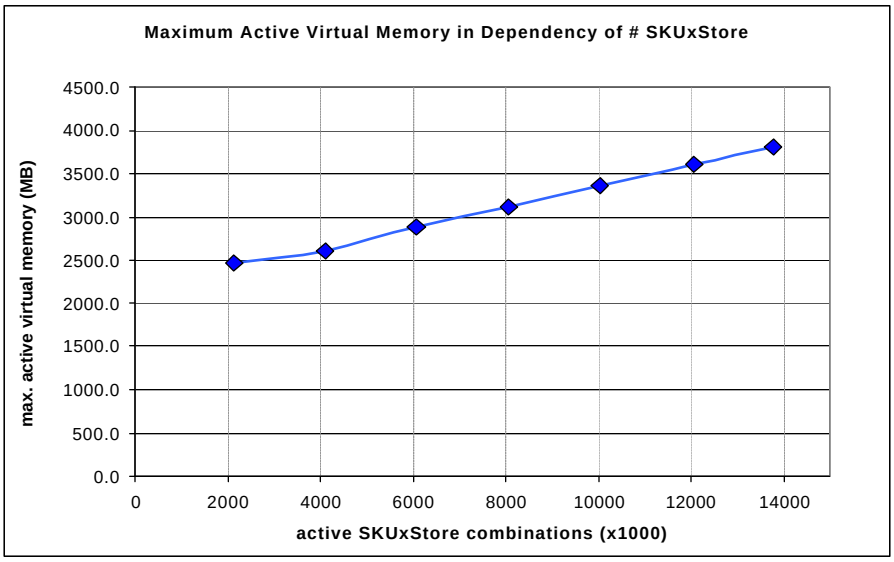

24 CPU, ORACLE 8.1.6, 32Bit

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- rplext: 25 threads, rplbld: 24 threads, reqext 25 threads

ORACLE 32bit Versus 64bit

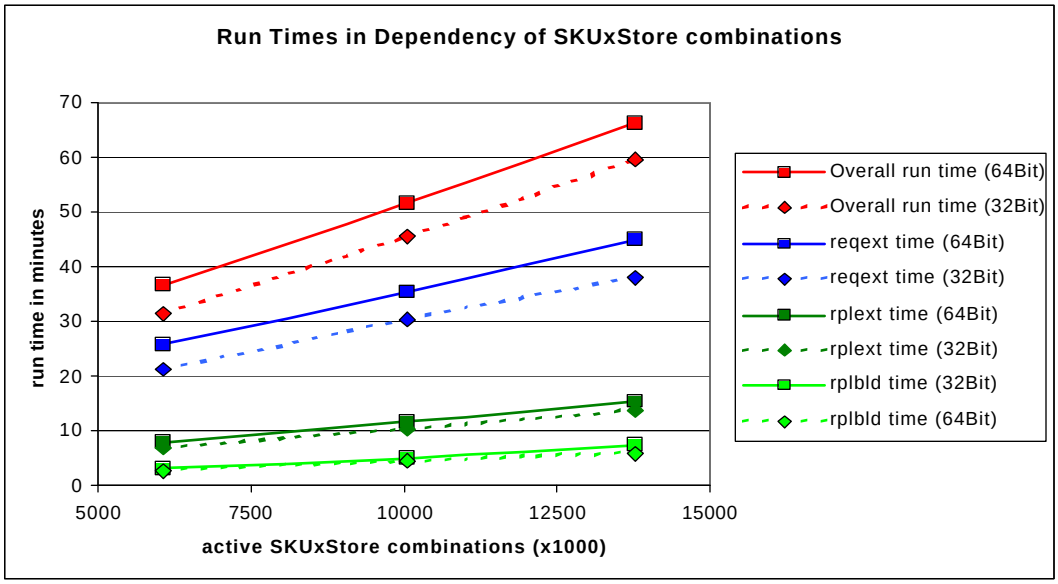

24 CPU, ORACLE 8.1.6, 32Bit, thread optimized

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- The run times when using ORACLE 64bit - all other parameters are identical - is between 10% and 14% longer.
- Memory usage also increases, when using 32bit instead of 64bit. The amount of active virtual memory used by the 64 bit version is 200MB (WL7) to 500MB (WL1) higher than the amount used by the 32bit version.

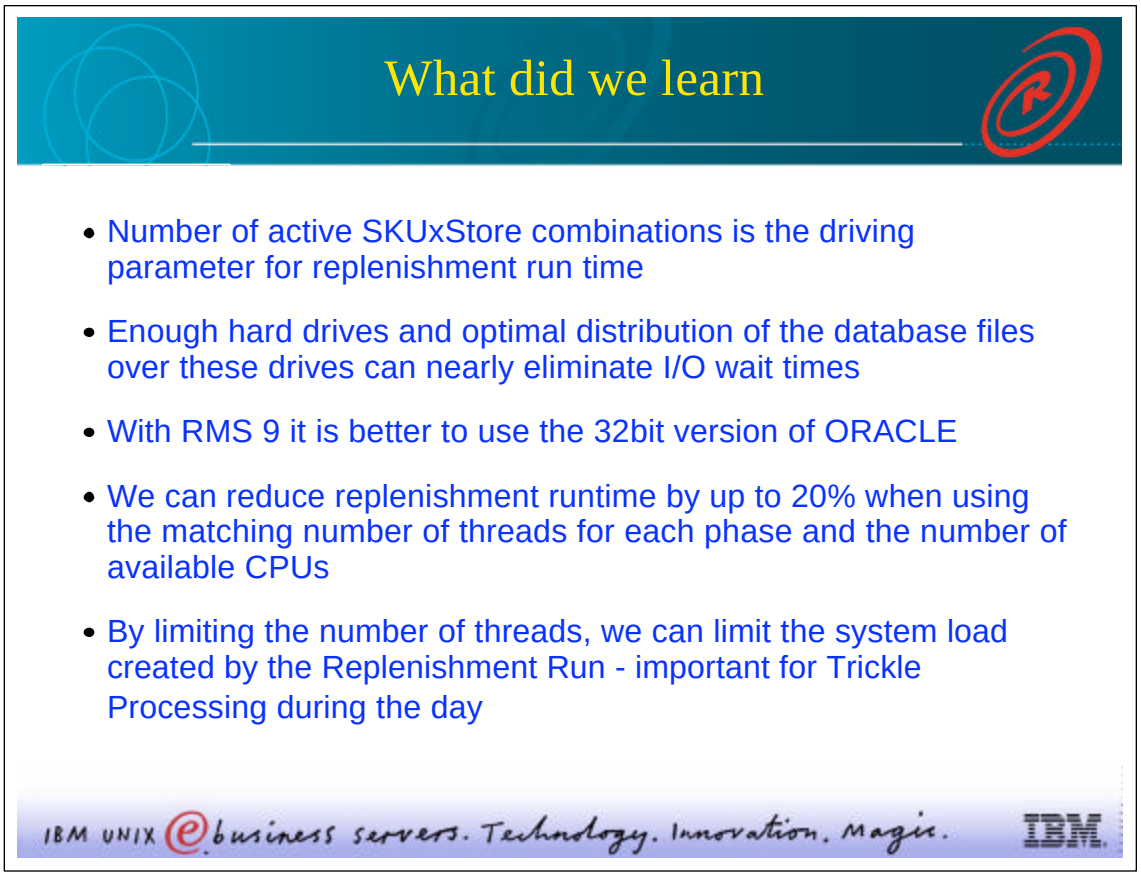

- I/O info
- Total System I/O Statistics: -- WL1, CPU24, interval 30s
- ---
- Max tps during an interval: 4,528
- Avg tps during an interval: 1,856
- Total number of Kbytes read: 20,119,640
- Total number of Kbytes written: 59,557,759
- Read/Write Ratio: 0.34
-
- Total System I/O Statistics: -- WL7, CPU24, interval 60s
- ---
- Max tps during an interval: 5,188
- Avg tps during an interval: 1,138
- Total number of Kbytes read: 7,176,210
- Total number of Kbytes written: 10,441,634
- Read/Write Ratio: 0.69
-

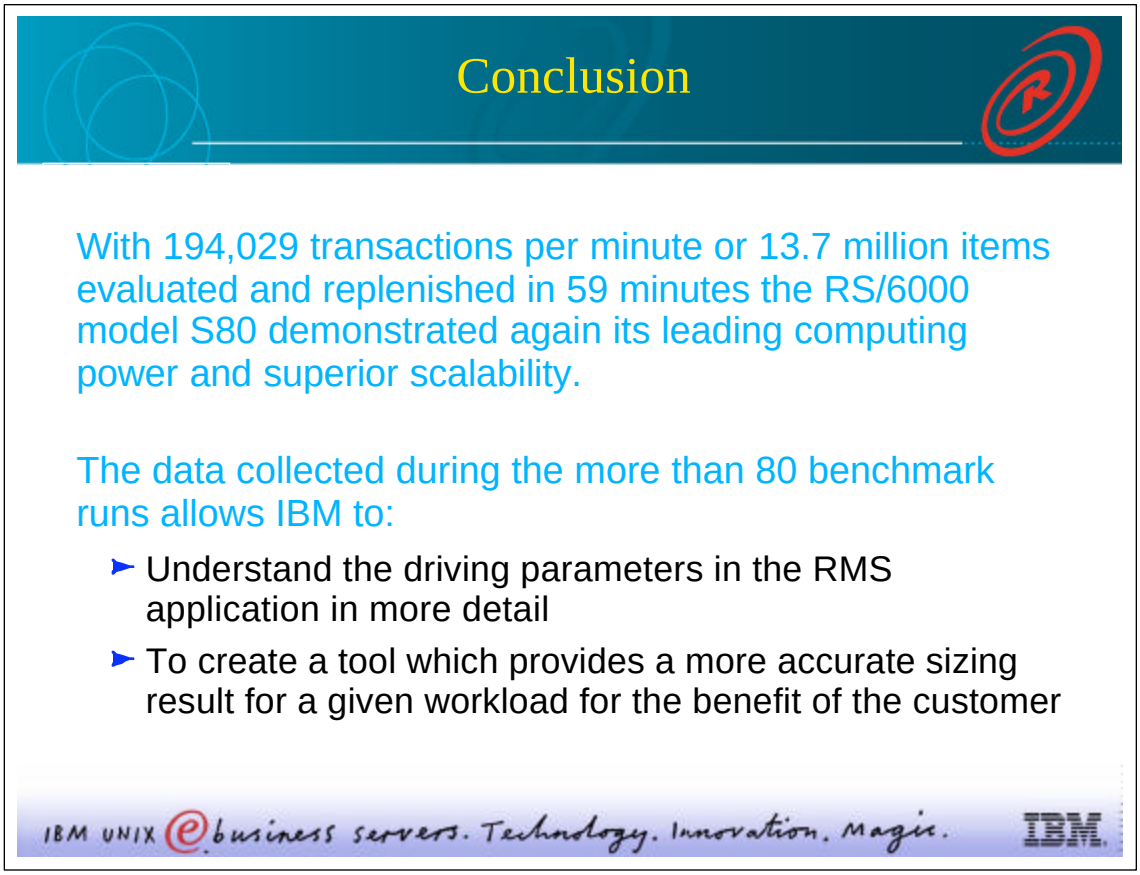

- "normal" RMS determines a configuration where the replenishment module is used down to the store level
- for larger Retailers the required memory is not determined by the replenishment run but by the number of interactive users
-

Contacts

John Pappas

Partner Account Manager
IBM Web Server Group
Phone: (770) 206-9410
E-mail: jpappas@us.ibm.com

Christa Tollefsrud

IBM / Retek Solution Sales - Enterprise Applications
Phone: (612) 397-5135
E-mail: ctollefs@us.ibm.com

For more information on IBM's RS/6000 and NUMA-Q servers, please visit www.ibm.com/servers.

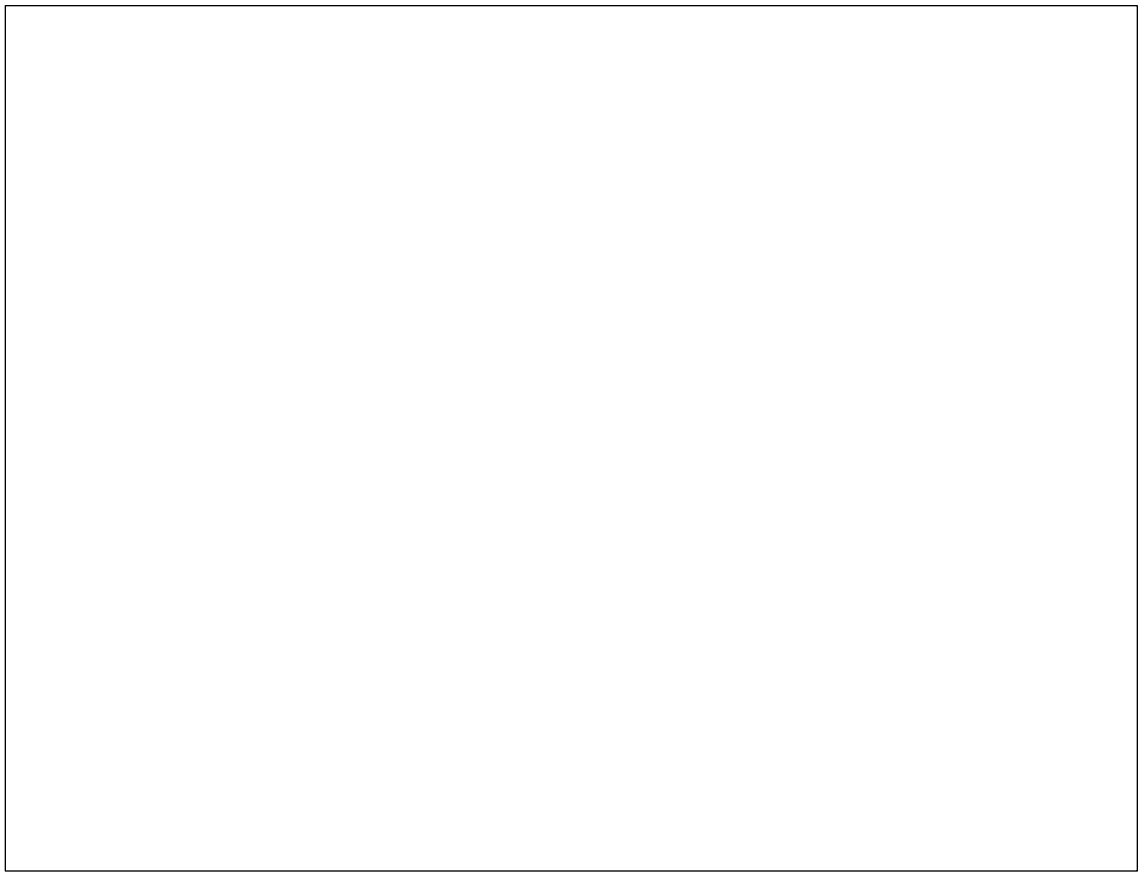

Database disk layout (cont.)

	Loop A						Loop B	
ssa0	6	7	8	9	10	12	13	20
ssa1	22	23	24	25	26	28	29	36
ssa2	38	39	40	41	42	44	45	52
ssa15	166	167	168	169	170	172	173	180
ssa4	54	55	56	57	58	60	61	68
ssa6	70	71	72	73	74	76	77	84
ssa8	86	87	88	89	90	92	93	100
ssa16	182	183	184	185	186	188	189	196
ssa9	102	103	104	105	106	108	109	116
ssa11	118	119	120	121	122	124	125	132
ssa12	134	135	136	137	138	140	149	148
ssa14	150	151	152	153	154	156	165	164

Comments

- ssa0 - ssa16: 12 x SSA adapter
- color coded disks are assigned to one of the five defined volume groups
- not all connected disks are shown (192 were connected)
- mirror disks are not shown (16 per volume group)
- 160 disks used for database
- remaining 32 disks used for database backup

High Water Result (cont.)

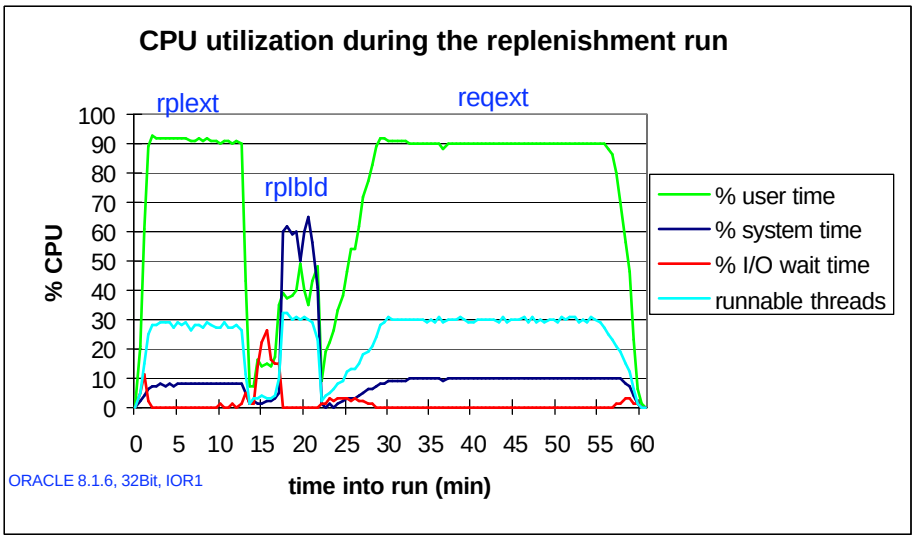

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

- rplext - 33 threads, sleep between thread starts 2s
- rplbld - 32 threads, sleep between thread starts 0s
- reqext - 33 threads, sleep between thread starts 9s
- The consistent smooth CPU profile achieved during the benchmark tests demonstrates the high User CPU utilization.
- The drop off after each module completes is rapid. This indicates that each thread is completing in approximately the same amount of time, this highlights that each thread has encountered little contention within the database.
- There is minimal I/O wait encountered throughout the entire process.
- The system CPU usage remains constant throughout each module's execution.

ORACLE 32bit <-> 64bit (cont.)

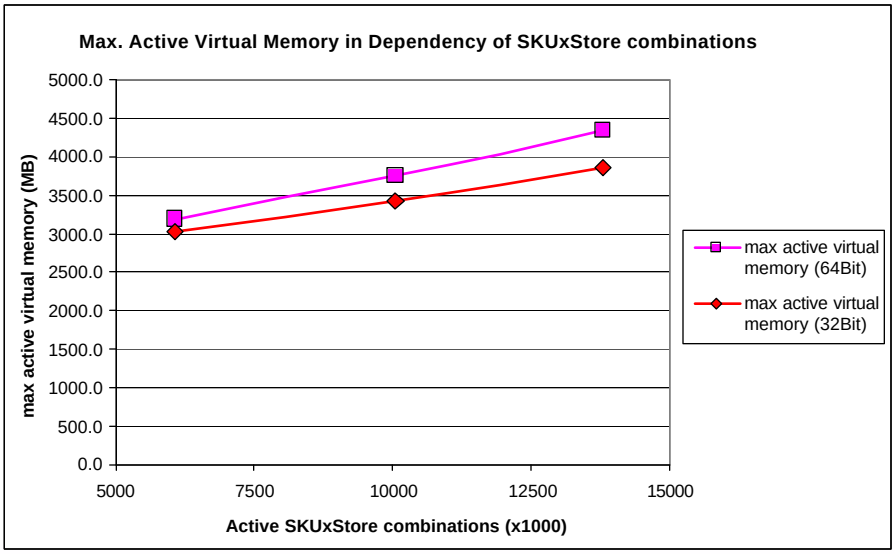

- You need between 5% (IOR3) and 10% (IOR1) more memory, when you are using ORACLE 64bit instead of ORACLE 32Bit.

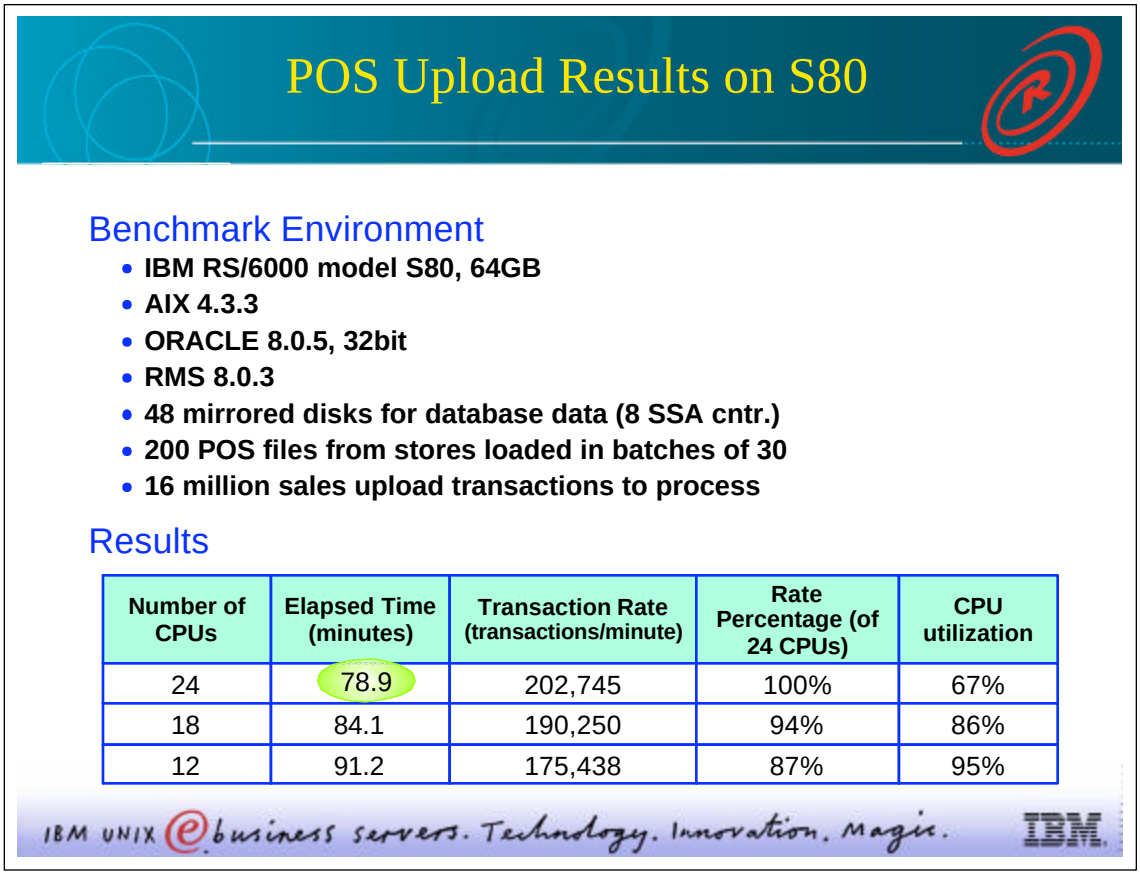

- "Transactions/Minute with100% CPU util" (assumption no I/O contention) 24 CPU - 302669, 18 CPU - 221221, 12 CPU - 184672
- Benchmark heavily I/O bound for 18 and 24 CPUs.
- Number of CPU's: This column shows the number of CPUs in the configuration.
- Elapsed Time:This column shows the number of minutes it took to run of 16,000,000 POS Upload transactions.
- Transaction Rate: The transaction rate per minute was arrived at by dividing the 16,000,000 transactions run by the number of minutes required for the run.
- Rate %: This column shows the scalability of running 16,000,000 transactions with a different number of CPUs. The 24 CPU run with a result of 202,745 transactions per minute is set at 100% for the base. Reviewing the table, notice that when the "Number of Processors" are reduced to 12 CPUs which is 50% of the 24 CPUs but the transaction rate only drops off 13% or to 87% from the 100% base.
- CPU Utilization: This data shows the utilization of the CPUs in the different CPU configurations.

RDF Benchmark Results on S80

Benchmark Environment

- **IBM RS/6000 model S80, 64GB, 24 CPU**
- **AIX 4.3.3**
- **RDF 4.1**
- **176 disks for 60 domains (8 SSA cntr.)**
- **5000 stores x 12,000 SKUs = 62 million combinations**
- **20 million forecasts need to be generated (Horizon 13 weeks)**

Results

Batch Process	Processing Time (minutes)
Loading of Hierarchies	20
Loading of Sales Data	27
Generation of Forecasts	40
Exporting of Forecasts	35
Total	**122**

IBM UNIX e business servers. Technology. Innovation. Magic. IBM

Parameter	Final Level	Source Level
Forecasting Method	Simple/CrostonAutoES*	
Source Level	SKU/Store*	N/A
Training Window	All* (110 weeks)	All* (110 weeks)
Forecast Horizon	13 Weeks	N/A
Cumulative Interval	Yes*	N/A
Export Forecasts	All*	N/A

Parameters	Szenario 1	
SKU/Store Combinations	62,680,800	
Active SKU/Stores	20,735,820	
Domains	60	(run mode {1,2,3}
Data Density	15.68%	
Forecast Horizon	13	