Optimize Performance of your Application with POWER7

IBM

Technical Forum & Executive Briefing

17 al 21 Octubre 2011

Imagine PODER Imagine CAPACIDAD

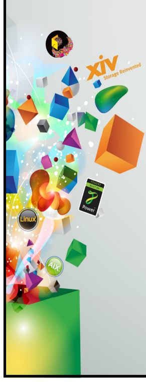

IBM Power Systems

IBM Foro and CIO Briefing - Cancun, Mexico 10/18/2011

Optimize Performance of your Application with POWER7

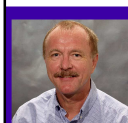

IBM

Gottfried Schimunek

Senior IT Architect
Application Design
IBM STG Software Development
Lab Services

IBM ISV Enablement

3605 Highway 52 North
Rochester, MN 55901

Tel 507-253-2367
Fax 845-491-2347

Gottfried@us.ibm.com

Acknowledgements

Many thanks to my colleagues and friends:

- Helen Olson-Williams – Performance Analyst,
  Lab Services Performance and Scalability Services

- Cindy Mestad – Senior Project Manager,
  Lab Services Performance and Scalability Services

- Anil Madan - Consulting IT Specialist,
  Lab Services Performance

- Charles Farrell – System Performance Analyst,
  Lab Services Performance and Scalability Services

- Eric Barsness – Application Consultant,
  Lab Services Performance

- Mark Funk - Software Performance Analyst,
  Systems HW Development

Agenda

- POWER7 Architecture and Performance Features
- Application Performance Characteristics
- Performance data analysis tools and methodology
- Tuning applications and system for best performance
- Practical Experiences and best practices
- Hardware configuration optimization

Comparative Anatomy … POWER6 vs. POWER7

More on Cache: http://en.wikipedia.org/wiki/CPU_cache

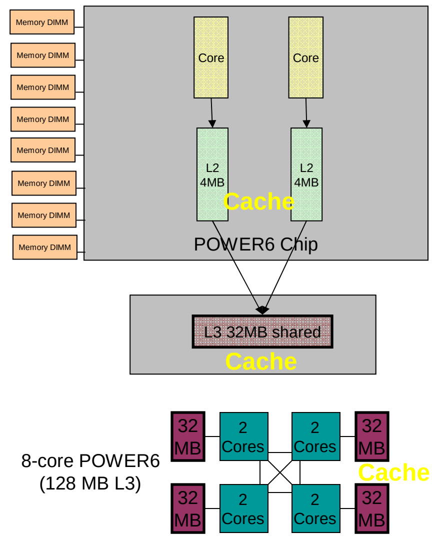

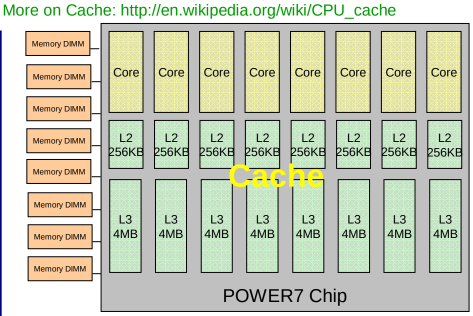

Important points:

- "Multi-core" is not new.
  - Fewer chips for same core count.
- POWER7 has 4x more cores per chip
- Cache topology is very different.
  - The latencies have improved with POWER7.
  - L3 data laterally cast out into neighbor L3.
- POWER7 has fewer memory DIMMs per core.

 Power your planet.

Systems Made From Multiple Chips

All memory and contents of all active cache accessible to all processors. (Cache-Coherent SMP)

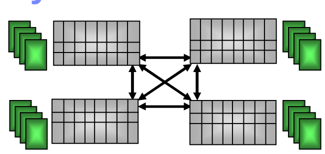

750
- 1 – 4 Chips
- 8 or 6 physical cores per chip
- 8 DIMM slots per chip

All chiplets (cores and cache) are active (i.e., no CUOD)
OS license count defines cores which execute instructions (i.e., a subset of active/physical cores).
- Single partition … Cores packaged densely.
- Multiple partitions … Each partition packaged densely if possible, partitions on separate chip's cores per memory.
- Cache Active only on OS-licensed chiplets … Energy Savings.

Active chiplet count based on CUoD … 4, 8, 16, 32, 48, 64 (and 24 in "TurboCore")
A 770/780 drawer always has 2 chips … 16 physical cores if 8-core chips.
OS licensed core count placed onto active chiplets.

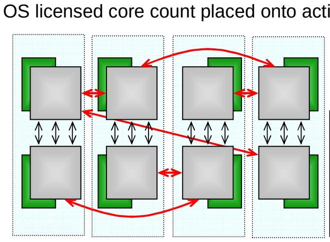

770 and 780
- 1 to 4 drawers
- 2 Chips per drawer
- 6 or 8 cores per chip
- 8 DIMM slots per chip

16-core Systems:
- POWER6 … 4 drawers and 8 chips
- POWER7 770 … 1 drawer and 2 chips.
- POWER7 780 … 2 drawers and 4 chips.

 Power your planet.

Simultaneous Multi-Threading (SMT) – POWER Systems

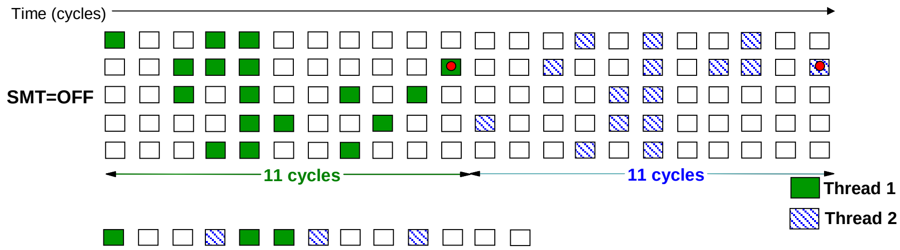

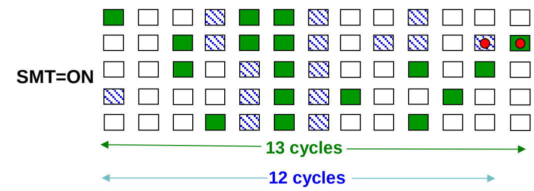

... in this trivial example -With SMT:
- Thread 1: 18% longer
- Thread 2: 9% longer
- 69% throughput increase

Amount of overlap/contention is environment dependent.

Capacity is increased but individual thread execution time in CPU is slightly more when running concurrently – because of possible contention.

4-WAY SMT (Simultaneous Multi-Threading)

POWER5/6 have 2-Way SMT

SMT4 Provides Still More Capacity.

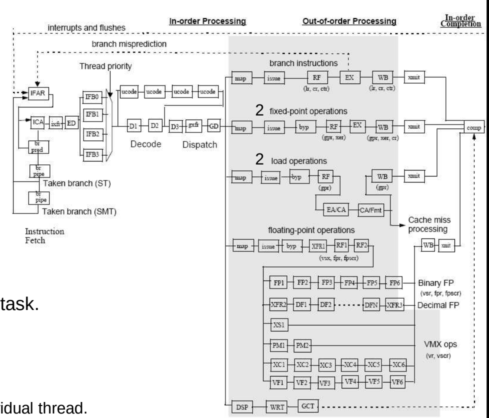

Up to 4 thread's instruction streams executing in parallel on one core.
- 4 threads share the same cache(s) and "pipes".
- SMT4 Performance enabled via Faster L2 cache.
- 8-core chip can support **32 threads**.
- 2-core POWER6 chip supports only **4 threads**.

SMT works because other thread's instructions execute during
- Another thread's cache miss delays.
- When pipe stage not already used by another.

Cores have a lot more capacity than is used by a single task.
- SMT4 common to have 1.5-2X more capacity over ST.

However, when core executing multiple threads, each thread's execution speed is slower.
- SMT4's 2X capacity increase also means ½ throughput for individual thread.

Cores automatically switch between ST, SMT2, SMT4.

More on Cache: http://en.wikipedia.org/wiki/CPU_cache

Cache Design Enhanced… Lateral Cast-out
(In preparation for explaining "TurboCore")

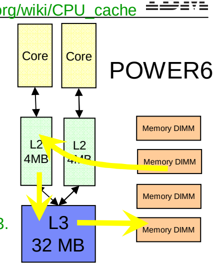

Fast L2 for rapid cache fills of L1s.
More Cache means
- **Better Single-Threaded Performance**
- **More System Capacity**
- **Acts to offset lower frequency.**

As data read into L2, other data stored out to L3. Called a "Cast-Out".

To make room for cache fills, block in L2 is "Cast-Out" to L3. (Data cast-out of L3 written to memory if changed.)

In order to make room for block cast **into** L3, a block may be "Cast-Out" out to another L3.

This is most useful when other cores on chip are temporarily or permanently idle.

Lateral Cast-Out exists on all chips **by default**, but only amongst licensed/active chiplets.

(Now onto "**TurboCore**".)

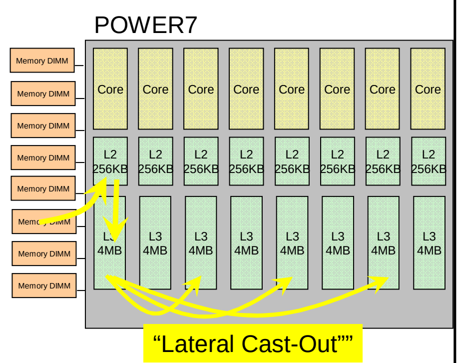

 Power your planet.

"TurboCore"

More Cache And Higher Frequency.

A chip with 4 Active Cores can more easily be cooled.
This, in turn, allows a higher frequency.
(3.86 GHz → 4.14 GHz)

8-Chiplet chip has 8 L3 caches, so remaining 4 remain enabled.

Best results expected in multi-partition environment with partition's packaged within individual chips.

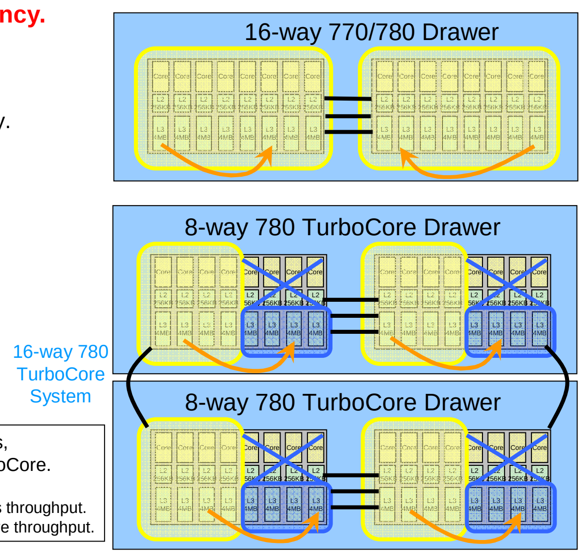

In some workloads and partition placements, cross-chip traffic can offset benefits of TurboCore.
Example:
16-core CPW … 4-chip TurboCore has slightly less throughput.
8-core CPW … 2-chip TurboCore has slightly more throughput.

 Power your planet.

The Conclusion Chart …

Think of POWER7 Performance as having …

- Better Core-to-Core Capacity Producing Desired Growth
- Massive improvement in “socket” capacity …. 8 cores per chip
  - A lot more capacity within a package (card, drawer).
- A lower frequency (than POWER6), but still better
  - Capacity improvements result from other improvements
  - Many single-threaded workloads execute better than POWER6.
- Lower energy consumption per core
  - Even more energy savings if you exploit IBM Systems Director Active Energy Manager.

Application Performance Characteristics

- Batch performance characteristics; single vs multithread
  - Multithreaded applications can automatically take advantage of new POWER features
    - Typically I/O bound
  - Single threaded applications will need some tuning
    - Typically uses only a single core at a time (e.g. max 6% of a sixteen way)
    - Typically I/O bound
- CPU or disk I/O bound?
  - Low overall CPU utilization, high disk utilization, lots of memory paging
  - Tuning techniques available to make an I/O bound job CPU bound
- Batch jobs waiting for something else?
  - Wait for I/Os, e.g. disk configuration # disk arms, disk speed, IOA, disk cache, etc.
  - Wait for processors, CPU queuing, higher priorities, virtualization, capping, etc.
  - Wait for record locks and seizes, DB design, application architecture
  - Wait for many other potential wait situations, hundreds of wait situations

Multithread vs single thread applications - Example

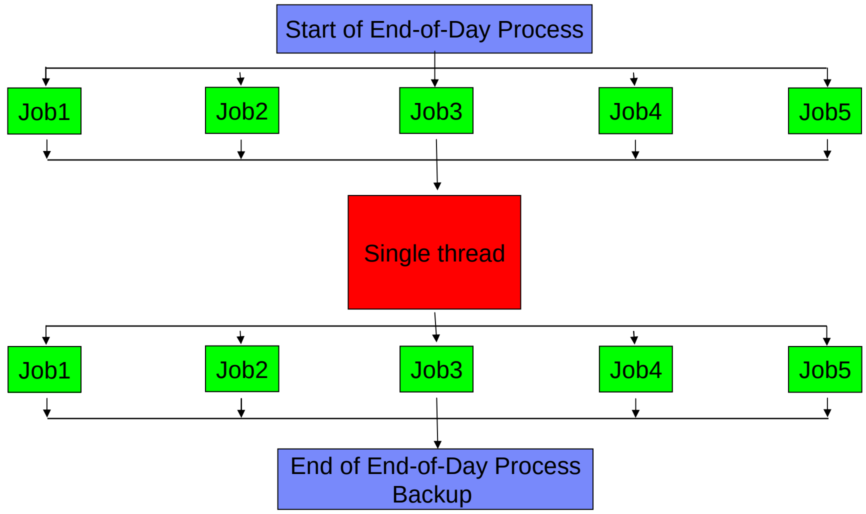

- Some applications cannot be run in multithread mode, e.g. data integrity, data dependency, etc
- Processor utilization in single threaded is typically below 20%
- Performance improvements in single threaded jobs are limited to one processor unless DB SMP is used

Single versus multi-threaded Application runtime

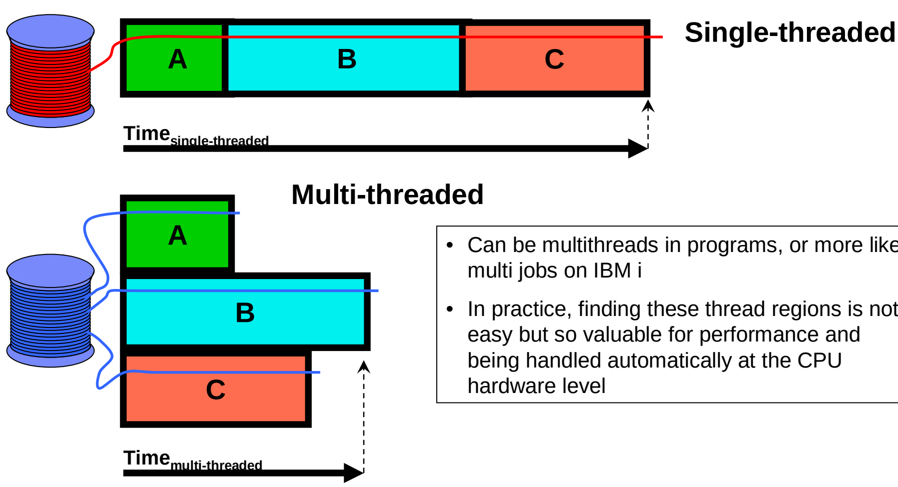

- Can be multithreads in programs, or more likely multi jobs on IBM i
- In practice, finding these thread regions is not easy but so valuable for performance and being handled automatically at the CPU hardware level

IBM Power Systems

How to multithread an application?

- Use DB2 SMP feature to parallelize I/O and index operations
- Use Holey inserts to improve concurrent ADDs
- Split input stream into multiple concurrent batch jobs so that each job works on separate parts of the input data
- From single thread:
- To multiple threads:

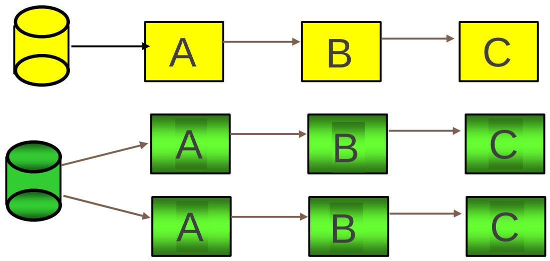

- use iStream and BSTAR from Cinimex at http://www.cinimex.co.uk/products/bstar.html
  – Automatically multithreads applications w/o application programming required

IBM Power Systems

Finding the optimal number of threads

- The first challenge is to evenly distribute work across threads.
- In this simple example, 8 threads are running concurrently, but the work is not evenly distributed
- CPU resources are under-utilized.

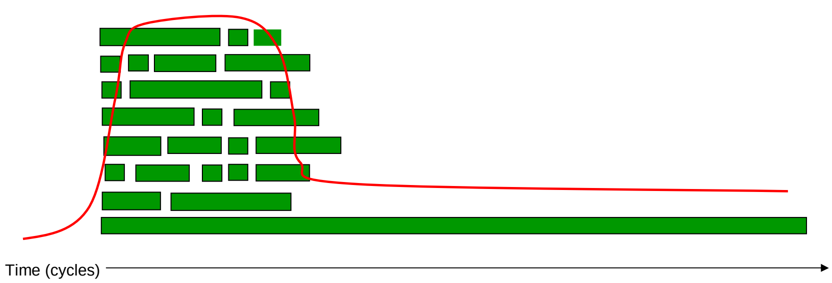

Find the optimal number of threads

- On Power 7, for the highest overall capacity,

number of concurrent threads = number of cores * 4

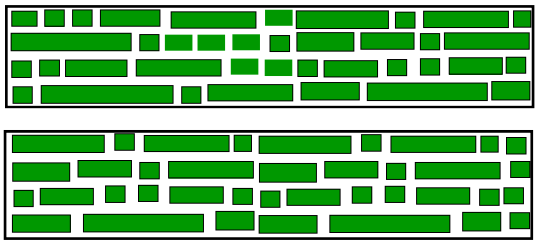

2 cores

8 concurrent threads

- For the best runtime for individual threads,

number of concurrent threads = number of cores

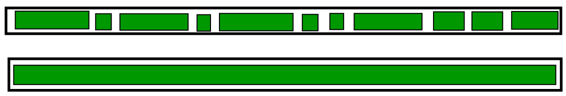

2 cores

2 concurrent threads

High priority work

Power 6 CPU Queuing Time -> Power 7 Dispatched CPU Time

Power 6, 5.0 Ghz,
3 PUs, 3 Virtual Processor
SMT 2 -> 6 threads dispatched

Power 7, 3.5 Ghz,
2.5 PUs, 3 Virtual Processor
SMT 4 -> 12 threads dispatched

Note: It appears there are additional CPU resources available on Power 7.
Keep in mind that thread execution time is longer when in SMT 4 mode.

How to improve I/O bound jobs?
- make them CPU bound..

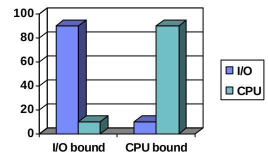

- Goal: Eliminate all or reduce most of the I/O waits
- Use SETOBJACC for high volume accessed files and indexes
- Avoid open/close of files, use RETRN vs SETON *LR in RPG
- Remove deleted records – RGZPFM
- Utilize OVRDBF REUSEDLT(*NO) to get both concurrent writes and blocked inserts enhance concurrent writes
- Optimize Load/ Copy/ Save processes by multithreading and/or other techniques
- Change random reads to sequential reads
- Higher blocking for input/output only operations – SEQONLY((*YES #rcds)) >= 128KB Increase logical page size in DB operations
- Avoid opening Output files as I-O files to take advantage of logical blocking.
- Eliminate FRCRATIO(1)
- CRTLF PAGESIZ(64 – 512KB)
- Use Expert cache in job pools
- For High volume ADDs, use DB2 SMP for parallel index maintenance and use ALLOCATE(*YES)
- Use SQL to create DB rather than DDS
- If journaling
  - Use larger commit cycles
  - Use HA Journal Performance Enhancement feature (SS1 opt 42)
  - Reduce JRN entry size with RCVSIZOPT(*RMVINTENT)

Tuning applications and system for best performance?

- Memory pools, activity level, expert cache, SETOBJACC, etc.
  - Proper memory pool sizing and tuning
    - Move all available memory to the batch pool, or
    - Use QPFRADJ(2 or 3)
    - Make sure memory pool is enabled for expert cache

- Application tuning, e.g. SEQONLY((*YES nbrof rcds)), NBRRCDS(nbrofrcds)
  - SEQONLY is for program buffer (for input only, for output only)
  - NBRRCDS is used for SLIC buffer
  - Try to set as large as possible, e.g. 128KB or even 256KB

- SQL vs native I/O
  - SQL can take advantage of multithreading in IBM i, e.g. parallel ADD, large page support, etc.

- Journal caching
  - Reduces write waits and number of I/Os significantly
  - IBM i option- fee

- DB2 SMP
  - Enables many multithreaded features for the DB, e.g. index maintenance, index builds
  - IBM i option - fee

Analyze Performance by collecting performance data

- What performance tools to use?
  - Collection services – 5 minutes interval for the complete run duration
  - Job Watcher – to identify wait components
  - PEX Analyzer with disk read/writes start/end ops
  - Disk Watcher to review disk and IO related performance

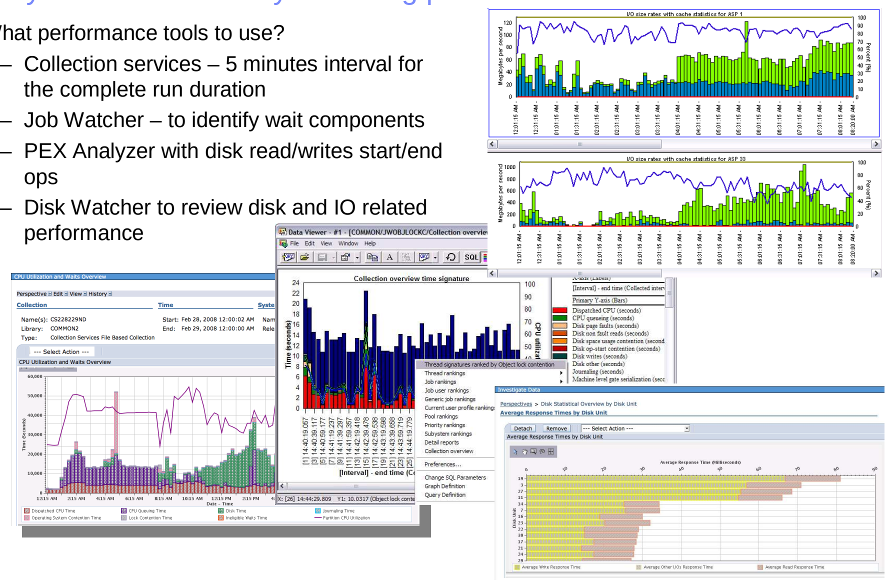

IBM Power Systems

iDoctor tools (Job Watcher and PEX Analyzer) Take You …

……………………..from this level……………..………

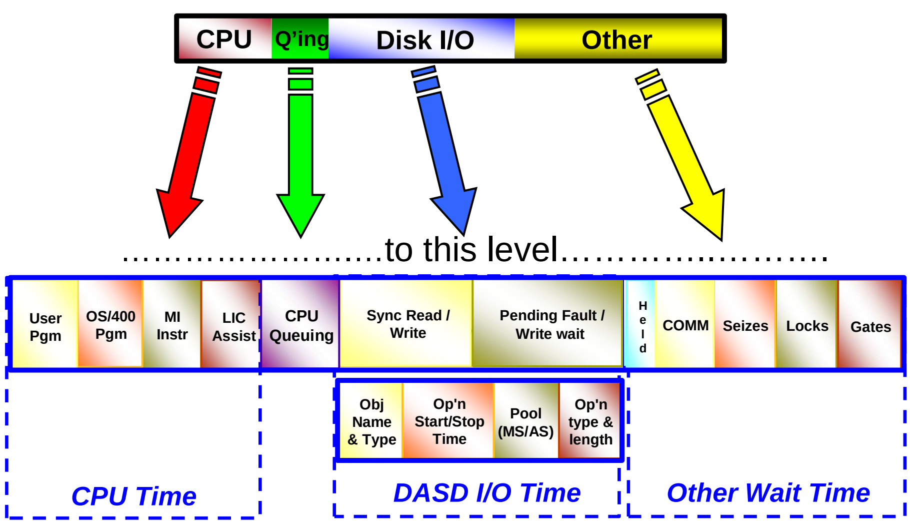

……………………..to this level…………..…………

Waits Overview – iDoctor Job Watcher

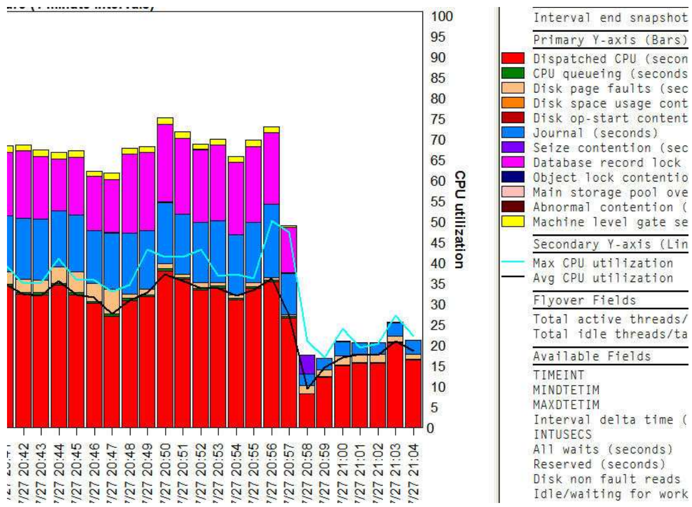

 Power your planet.

Waits Overview – iDoctor Job Watcher – Example Physical IO

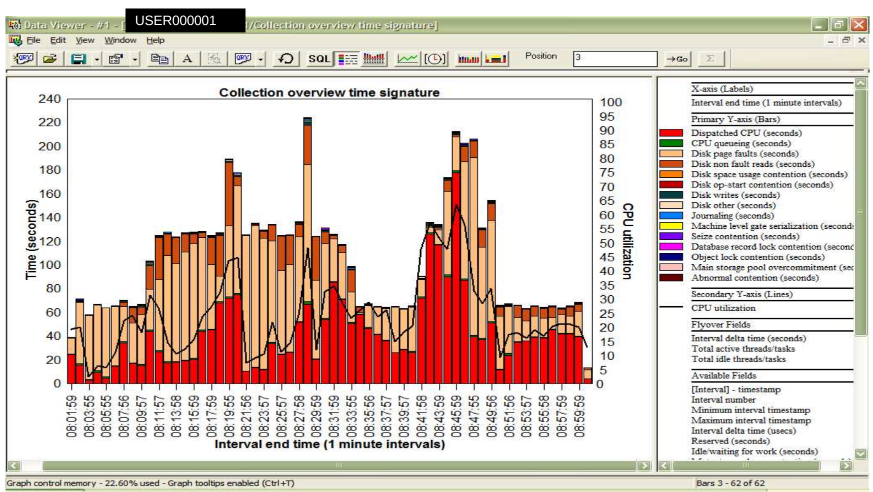

 Power your planet.

Waits Overview – iDoctor Job Watcher – Example Physical IO

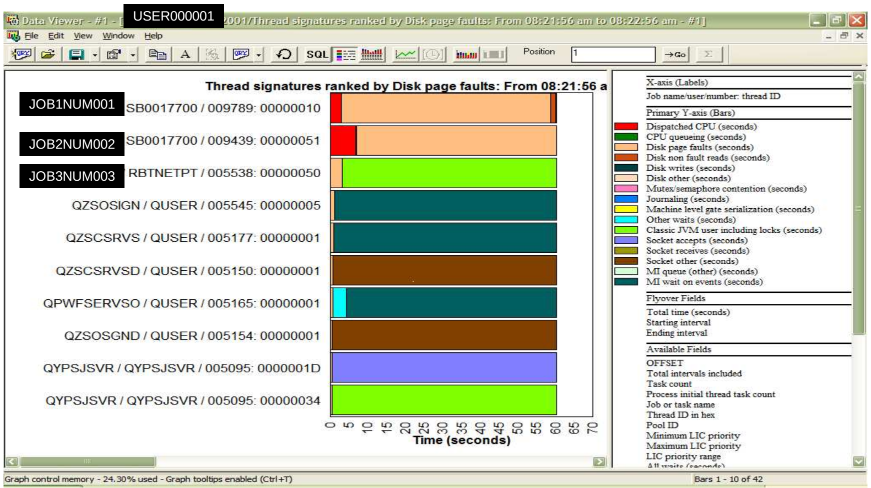

 Power your planet.

Waits Overview – iDoctor Job Watcher – Objects Waited On

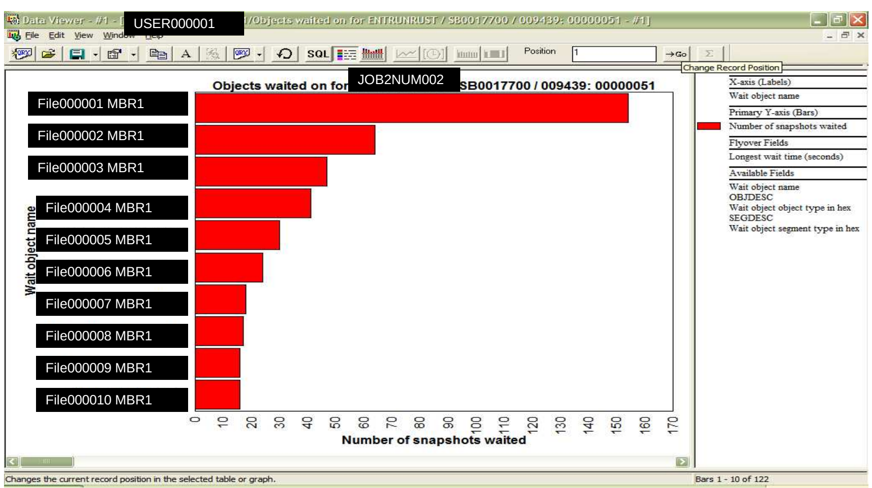

 Power your planet.

Waits Overview – iDoctor Job Watcher – Before and After

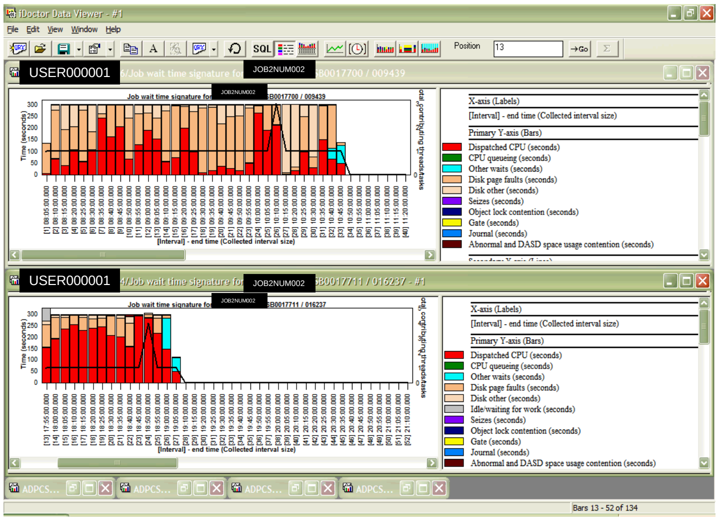

 Power your planet.

Other optimization examples from recent Lab Services engagements…..

 Power your planet.

Application scalability and tuning engagement for Business Partner

- Results
  - Maximum tps increased from 2000 -> 6139

- Changes made during Performance and Scalability Test
  - Created multiple copies of transaction file to reduce contention (gate wait) when inserting records to the transaction files.
  - Pre-allocate log file to reduce seize contention when files extend.
  - Switch from varchar to fixed length field in log file since most of the accesses are the max length, to reduce seize contention on auxiliary space.
  - Turn on journal caching.
  - Increase Journal Recovery Ratio - jorecra threshold to reduce seize contention on index
  - Contention on an application user queue - balance workload simulation. Will need to consider adding more queues in the future.
  - Change frcwrtratio from 1 to *none and specify Seqonly *no to reduce dasd wait time and to reduce contention when multiple jobs force a file to disk
  - Other application 'clean up' activities

    set index size to *max1tb for a few files

    set reuse deleted to *yes

Before tuning
2000 tps, cpu utilization 20%

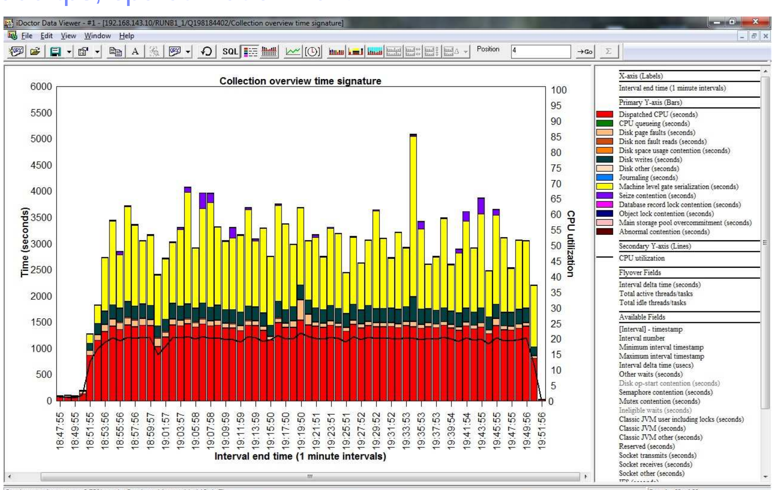

After tuning
6139 tps, cpu utilization 40%

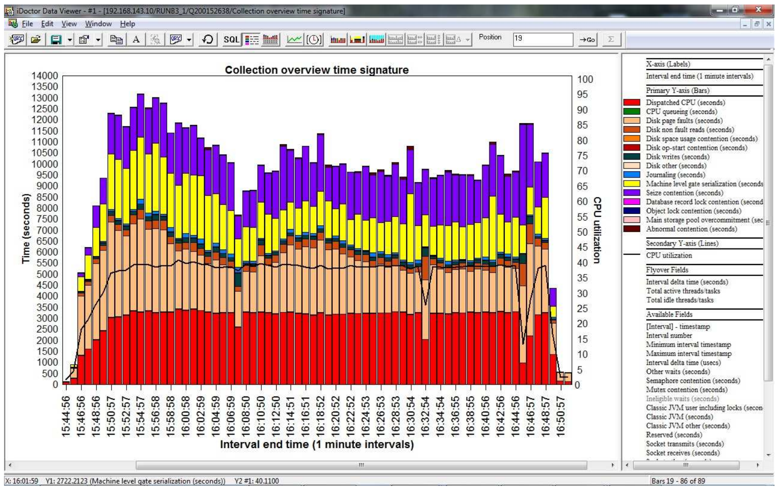

End of Day Batch – scalability testing and tuning for growth

- 2X, 3X, 4X growth scenarios, some results:

2X – End of Day	2X – End of Month	4X – End of Day	4X – End of Month
Power 6, 10 cores 220gb memory	Power 6, 10 cores 220gb memory	Power 7, 32 cores 586gb memory Appl and system tuning changes	Power 7, 32 cores 586gb memory Appl and system tuning changes
5 hrs 6 min	12 hrs 23 min	< 4 hrs	7 hrs 4 min

- Changes made during Performance and Scalability Test
  - Application changes:
    - Reduced full opens of db files
    - Increased parallelism in batch stream
    - Used setobjacc to load key files in memory
  - System and application settings
    - Turn on Journal cache
    - Re-use deleted records to take advantage of concurrent write
    - DB2 fix to reduce impact of seize when records added to file
    - Increased SMAPP recovery threshold

IBM Power Systems

2X – EOM
Before adding hardware and tuning changes

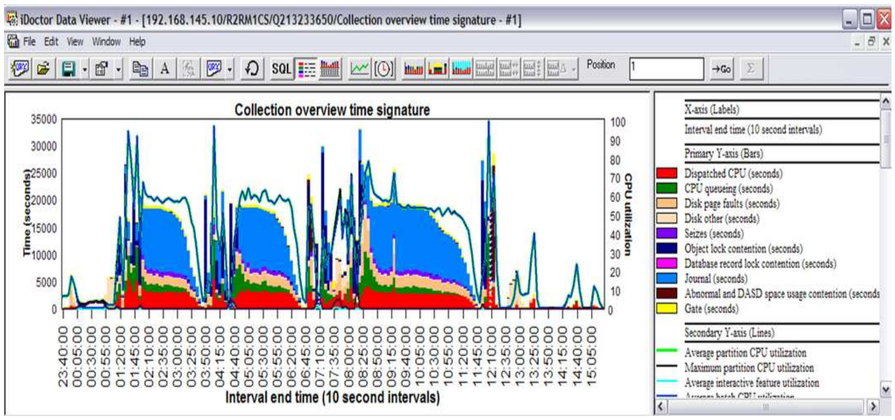

 Power your planet.

IBM Power Systems

4X - EOM
Additional cpu and memory, tuning changes

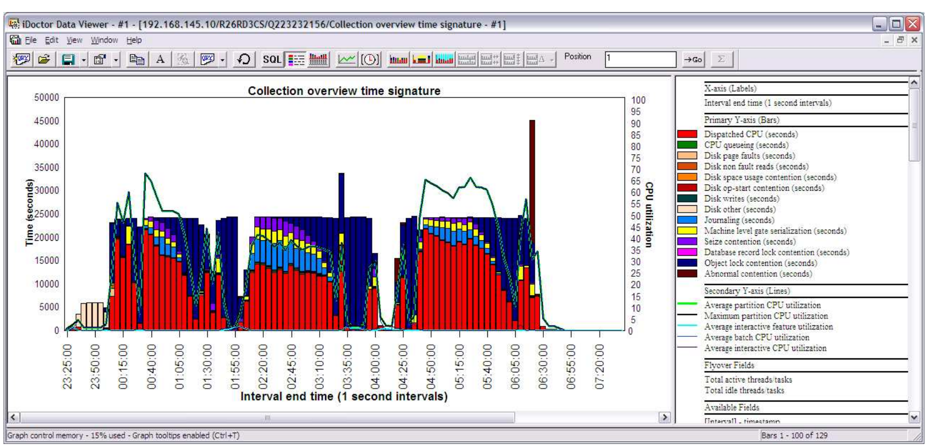

 Power your planet.

IBM

2X – EOD
Before adding hardware and tuning changes

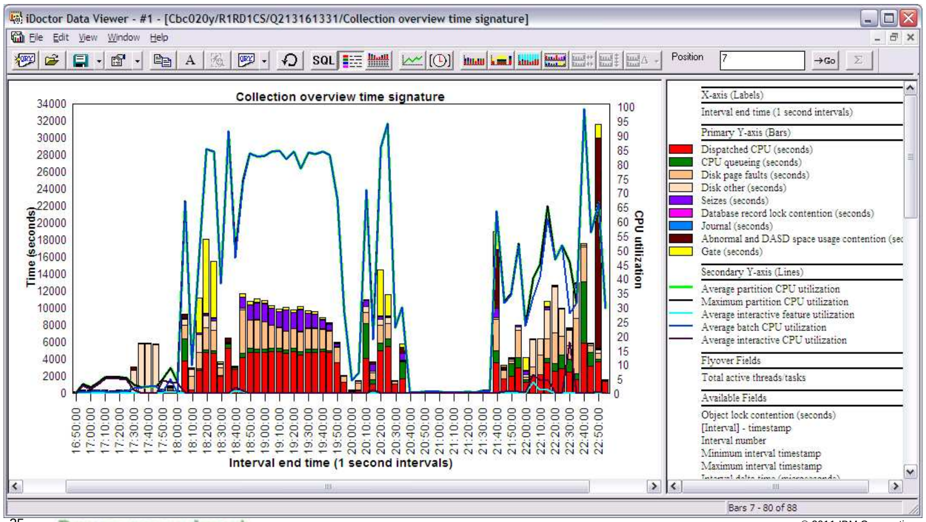

 Power your planet.

IBM

4X - EOD
Additional cpu and memory, tuning changes

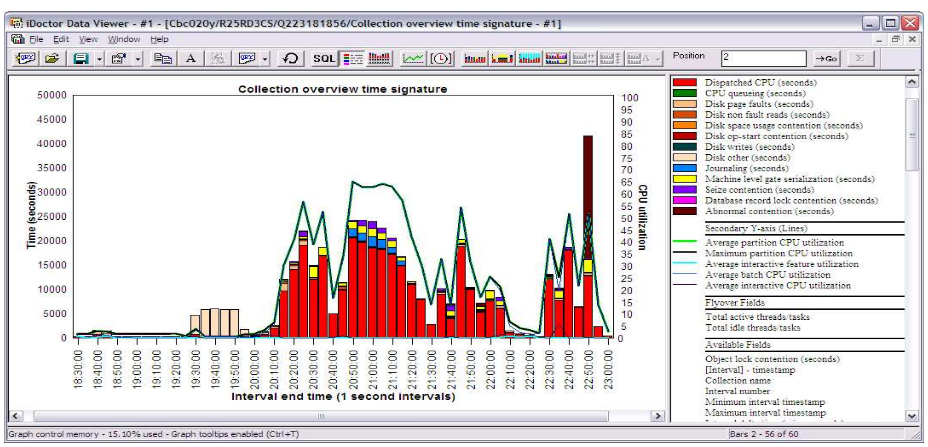

 Power your planet.

Application and System Tuning to achieve SLAs for quarterly pricing update

- Results:
  - Increased number of cpu cores by 50%
  - Tuned application and system
  - 86% improvement in throughput

- Tuning changes made during Performance and Scalability Test
  - Application changes:
    - Reduce date conversions
    - Change random to sequential IO
  - System and application settings
    - Tuned/reduced the number of concurrent jobs
    - Tuned memory settings
    - Apply PTF to take advantage of skip-locked-row sql processing

Before tuning application, system and adding hardware 8.7 million updates per hour

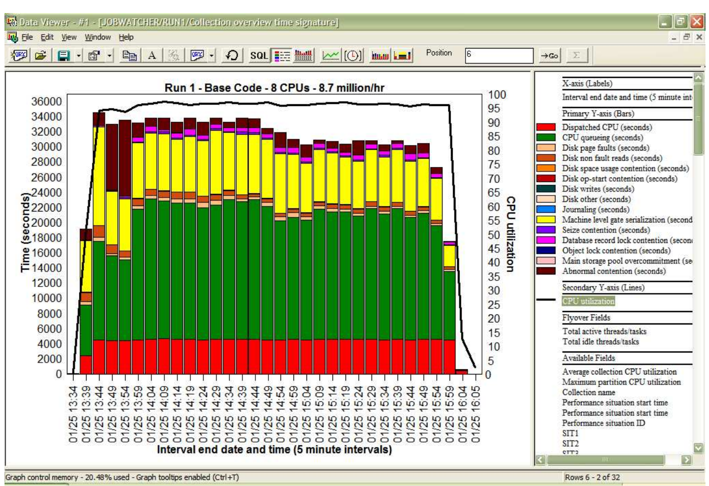

After application and system tuning changes, adding cpu 16.2 million updates per hour

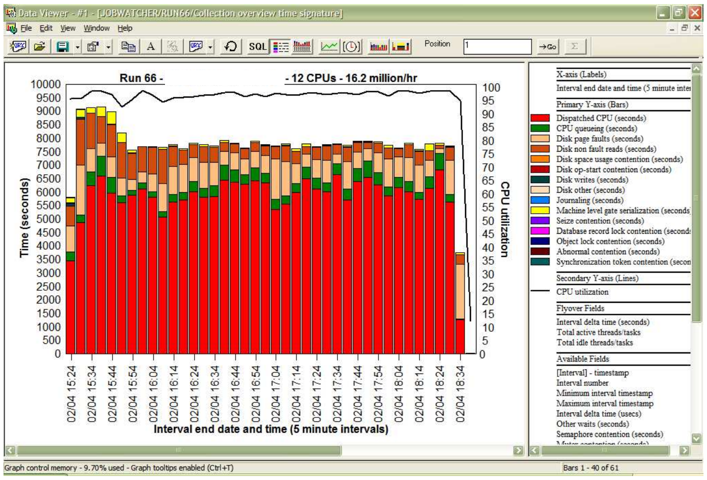

Application and System Tuning to achieve Scaling of Upto 2000 Users

- Results:
  - Increased scaling by 2X
  - Tuned application and system,
  - Overall response time profile (APDEX) improved from 6% acceptable to 91% acceptable

- Tuning changes made during Performance and Scalability Test
  - Tuned application specific ‘prestart’ job settings
  - Removed unnecessary activation group creation
  - Utilized OVRDBF REUSEDLT(*NO) to get both concurrent writes and blocked inserts
  - Optimized server job structure to reduce CL Program Initialization queuing
  - Tuned memory settings

Before tuning application, system and adding hardware – Scaling Upto 2000 Users

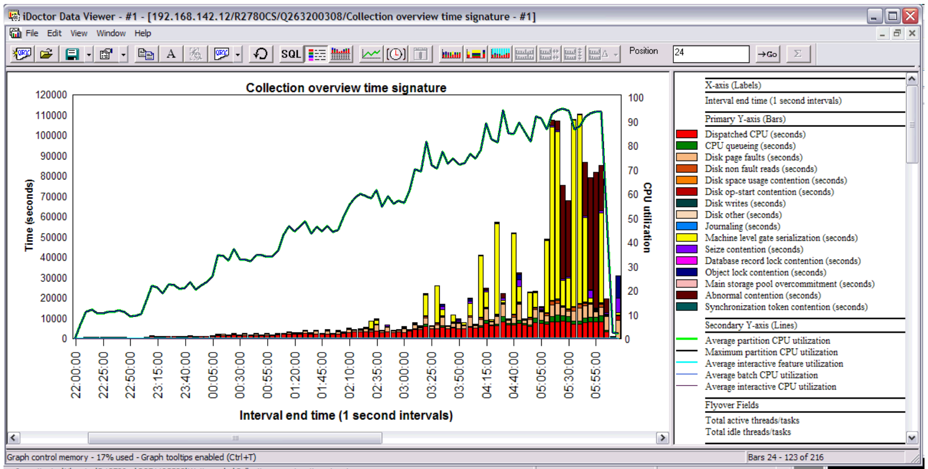

 Power your planet.

After tuning application, system and adding memory – Scaling Upto 2000 Users

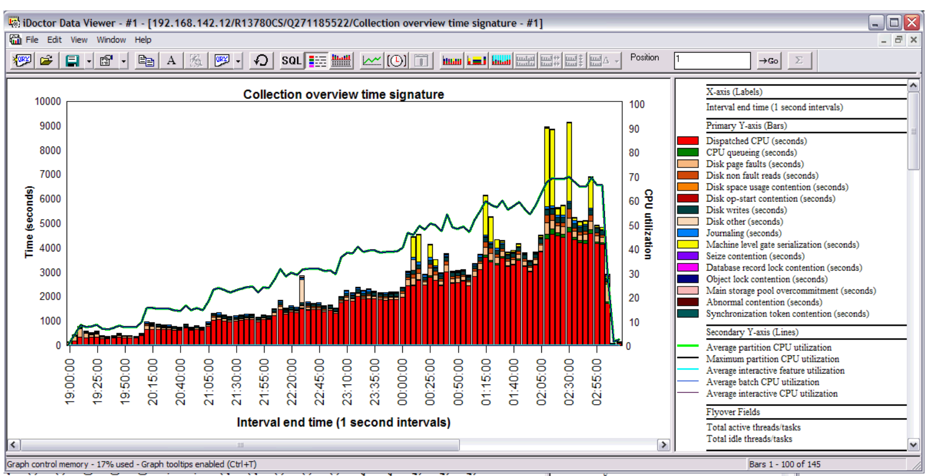

 Power your planet.

Before/After tuning application, system and adding hardware – Scaling Upto 2000 Users

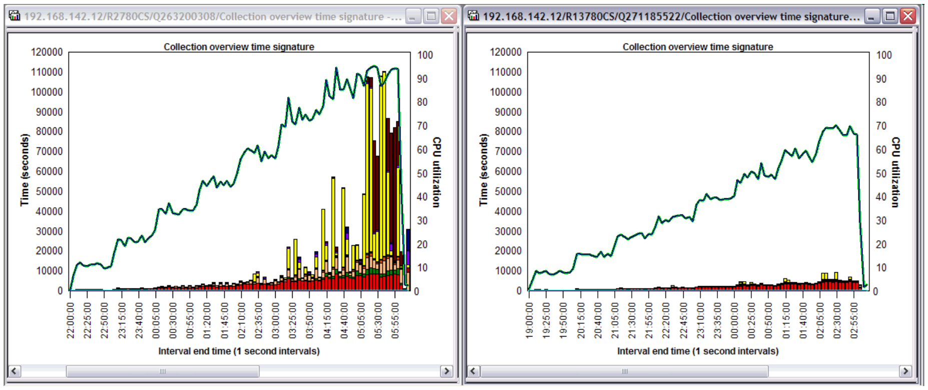

Response Time Acceptability Measure (APDEX)

Transaction	Ranking	240	480	720	960	1200	1440	1680	1920		240	480	720	960	1200	1440	1680	1920
Overall		0.98_0	0.99_0	0.98_0	0.96_0	0.86_0	0.63_0	0.41_0	0.06_0		0.99_0	1_0	1_0	0.99_0	1_0	0.97_0	0.97_0	0.91_0

Application and System Tuning to achieve Scaling of 7X Accounts

- Results:
  - Objective was to scale to 3X accounts
  - Tuned application and system, and achieved 7X scaling
- Tuning changes made during Performance and Scalability Test
  - Partitioned single transaction log file into multiple log files
  - Resolved Object Lock contention on SEQ Object with PTF to reduce CPU and contention due to object authorization checking at the object level.
  - Optimized infrastructure of front end server job queue depth and number of back end server jobs.

Before tuning application, system and adding hardware – Scaling EOD Upto 7X Accounts – Run Time Not Acceptable

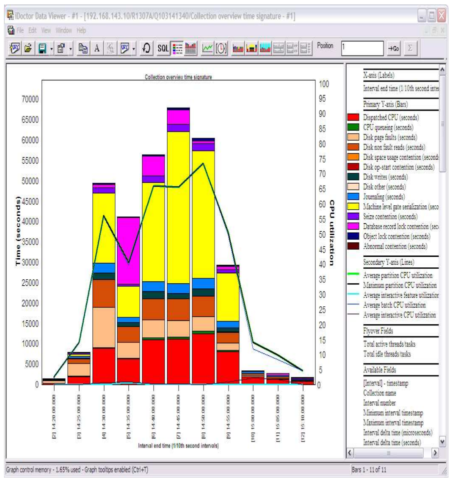

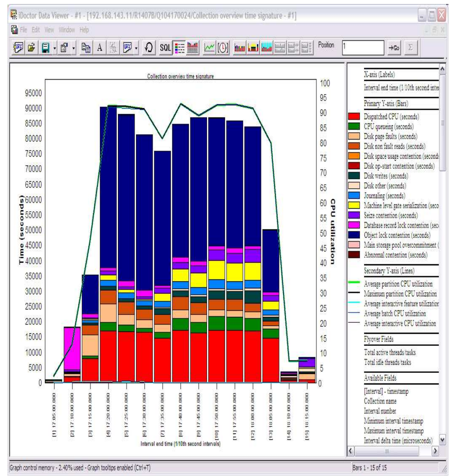

After tuning application, system and adding hardware – Scaled EOD Upto 7X Accounts With Acceptable Run Time.

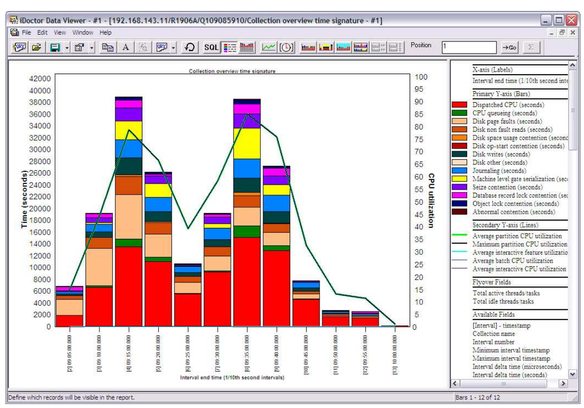

Hardware configuration optimization

- # of processors allocated, # of VPs
  - Ensure you always have enough CPU capacity for the complete run, e.g. <100%

- SSD – Solid Sate Drives vs HDDs performance characteristics and when to use
  - Use SSDs for select DB objects or hot pages traced from previous batch runs
  - Use the QIBMSSD tool to determine if SSDs will help,
  - use TRCASPBAL to trace and collect information of disk usage
  - And STRASPBAL to move hot objects onto SSDs, or back to HDDs
  - Investigate hot data and
  - Run TRCASPBAL and STRASPBAL *HSM (6.1.1) to move hot data to SSDs
    - Run STRASPBAL *SSSM (7.1) to move hot to SSDs + move cold data back to HDDs
    - Move complete files or libraries with CRTxF or CHGxF or CREATE/ALTER TABLE/INDEX
      - CRTLIB ASP(ssd-ASP), CREATE SCHEMA x IN ASP n

- Proper disk arm and disk IOA caching options
  - The more disk adaptors, the better
  - The more cache, the better
  - RAID5 typically provides better write performance
  - Use the WLE tool to actually use collected disk performance data and estimate most performance optimal disk configuration
  - Disk adapter cache for read and write ops
  - The larger the cache, the better the performance

- Power and energy saving options vs Performance options
  - Watch out for energy saving options set in HMC or Active Energy Manager – AEM
  - Can significantly influence performance

Mixed SSD + HDD – cold and hot data

- It is typical for data bases to have a large percentage of data which is infrequently used (“cold”) and a small percentage of data which is frequently used (“hot”)

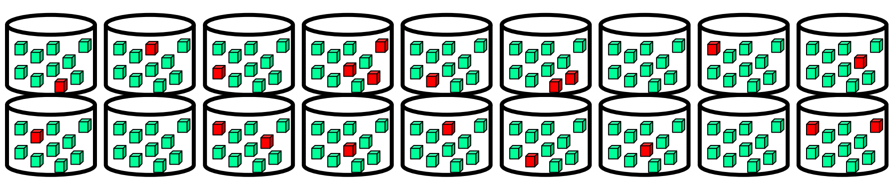

Hot data may be only 10-20% capacity, but represent 80-90% activity

- SSD offers best price performance when focused on “hot” data
- HDD offers best storage cost, so focus it on “cold” data …. sort of a hierarchal approach

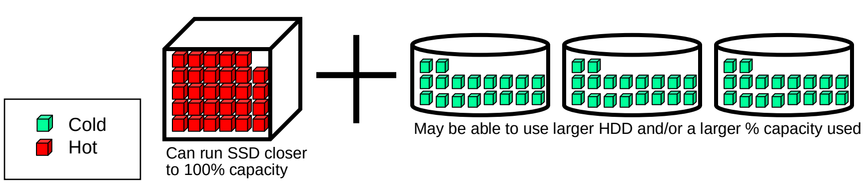

Energy Saving 101

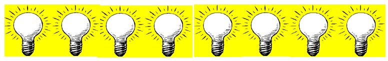

If I have 8 lights (i.e. Cores)
that produce 8000 lumens, and I only need 4000 …

My options include …

- **Turn some off … Cease clocking of some cores (OS-based)**
  - **Nap cores (End clocking core, cache remains active)**
    - **OS returns unused cores to hypervisor for napping.**
  - **Sleep chiplets (End clocking core and cache)**
    - **Sleep used if cores not licensed or if not used for an extended period.**

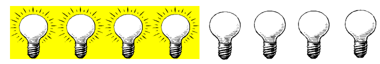

- **Dim them … Slow the Frequency (TPMD-based)**
  - **Dynamic Power Save option in Systems Director.**

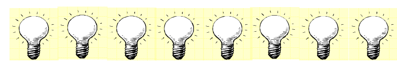

Frequency Adjusting Modes … Three options:

- **Static Power Save ….** Frequency Lowered by Fixed Percentage below nominal to save power.
- **Dynamic Power Save - Normal**
  Variably adjusts frequency to reduce power. Frequency returns to nominal when required by processor utilization.
- **Dynamic Power Save – Favor Performance**
  Same as “Normal”, but frequency can reach a super-nominal frequency if power/cooling limits allow.

Summary and references

- Performance of POWER6 and POWER7 systems are excellent – take advantage
- Some applications may not be able to take advantage of the POWER6 or POWER7 architecture features
  - Check out best practices and apply system and application tuning techniques
- IBM can help Clients, BP and ISVs to analyze and optimize their solutions to take advantage of POWER6 and POWER7
  - Let us know if you have a question or need some help with performance
    Send an email to: stgls@us.ibm.com

IBM Power Systems

- Questions

51 Power your planet. © 2011 IBM Corporation

IBM Power Systems

52 Power your planet. © 2011 IBM Corporation

Performance and Scalability Services in Rochester

- The IBM i Performance and Scalability Services Center can provide facilities and hardware IN ROCHESTER to assist you in testing hardware or software changes
  - “Traditional” benchmarks
  - Release-to-release upgrades
  - Assess application performance when migrating to a new release of IBM I
  - Stress test your system
  - Determine impact of application changes
  - Proofs of Concept (e.g. HA alternatives; SSD analysis, external storage, etc.)
  - Evaluate application scalability
  - Capacity planning
- … all with the availability of Lab Services IBM i experts and development personnel
- To request any of these services, submit at:

http://www-03.ibm.com/systems/services/labservices/psscontact.html

IBM Systems Lab Services and Training

Our Mission and Profile

- Support the IBM Systems Agenda and accelerate the adoption of new products and solutions
- Maximize performance of our clients’ existing IBM systems
- Deliver technical training, conferences, and other services tailored to meet client needs
- Team with IBM Service Providers to optimize the deployment of IBM solutions (GTS, GBS, SWG Lab Services and our IBM Business Partners)

Our Competitive Advantage

- Leverage relationships with the IBM development labs to build deep technical skills and exploit the expertise of our developers
- Combined expertise of Lab Services and the Training for Systems team
- Skills can be deployed worldwide to assure all client needs can be met

Successful worldwide history:
17 years in Americas, 9 years in Europe/Middle East/Africa,
5 years in Asia Pacific

www.ibm.com/systems/services/labservices stgls@us.ibm.com

Mainframe Systems

Power Systems

System x & Bladecenter

System Storage

IT Infrastructure Optimization

Data Center Services

Training Services

IBM Systems Lab Services and Training Power Services

Key Offerings

- High Availability Services on Power Systems (including Advanced Copy Services for PowerHA™ on IBM i)
- Systems Director Services
- PowerCare Services
- Performance and Scalability services (including system, application, and database tuning)
- Virtualization Services for AIX® on Power Systems™
- Application and database modernization consulting (SOA implementation)
- Linux® on Power consulting, custom application development, implementation, and optimization services
- Security on Power consulting and implementation services
- System consolidation and migration service
- High Performance Computing consulting and implementation services
- SAP® on IBM i consulting
- Power Blades on BladeCenter (including VIOS on i and blades running IBM i implementation)
- Smart Analytics services (including DB2® Web Query implementation and consulting)
- Public, private, customized and self-paced virtual training
- Power Systems Technical University

Americas, WW Contacts

Mark Even
even@us.ibm.com, 507-253-1313
IBM i

Frank Kriss
kriss@us.ibm.com, 507-253-1354
IBM i, High Availability

Karen Anderson
kanders@us.ibm.com, 972-561-6337
IBM i Vouchers

Stephen Brandenburg
sbranden@us.ibm.com, 301-803-6199
PowerVouchers, Virtualization Program, AIX SWOT/SWAT, AIX

Allen Johnston
allenrj@us.ibm.com, 704-340-9165
PowerCare

Dawn May
dmmay@us.ibm.com, 507-253-2121
Power Performance and Scalability Center

www.ibm.com/systems/services/labservices stgls@us.ibm.com

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. 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 document. The furnishing of this document 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 future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

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

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

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

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document 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 document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A full list of U.S. trademarks owned by IBM may be found at: http://www.**ibm.com**/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
AltiVec is a trademark of Freescale Semiconductor, Inc.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

Revised December 2, 2010

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. 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 Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 2, 2010

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. 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 Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	**http://hmmer.janelia.org/** http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised December 2, 2010

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.

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, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

==

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007