





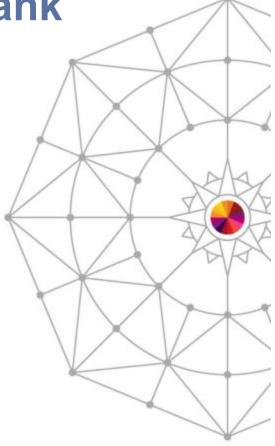
Cheryl Watson and Frank Kyne Watson & Walker, Inc.

technical@watsonwalker.com

www.watsonwalker.com

March 13, 2014 Session 15397







Abstract



Cheryl Watson, best known for her newsletter on MVS Performance (*Cheryl Watson's Tuning Letter*) will provide her favorite tips from her 23 years of writing. Our new associate, Frank Kyne, a long time author and editor of IBM's ITSO Redbooks, as well as a renowned teacher of the ITSO workshops, will provide his favorite tips from his 16 years of writing and teaching. And as an expert on outsourcing, and designer of a tool to resolve outsourcing disagreements (BoxScore), Cheryl will also present the best ways to prevent your outsourcing experience from being a disaster. Everyone is guaranteed to take away something of value from this session. You will receive a full copy of one of our recent *Tuning Letters*, and a drawing will provide one lucky attendee with the last 23 years of Tuning Letters on a DVD.





Agenda

- Who we are
- Frank's lukewarm tips
- Outsourcing gotchas
- Hot tips from Cheryl
- Drawing for DVD with with 23 years of Tuning Letters







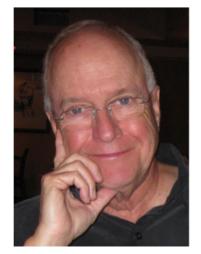
Who we are

- Watson & Walker, Inc. A SMALL Business!
 - Tom Walker, CEO and Publisher
 - Cheryl Watson, President, teacher, author, consultant
 - Frank Kyne, Editor, teacher, author, consultant
 - Linda May, Business Manager
 - Consultants: Scott Barry of SBBWorks and others
 - Recognized ISV and IBM Partner in IBM PartnerWorld
 - Loving what we do every day!



S H A R E Tethnology - Connections - Nasults

Watson & Walker, Inc. – A SMALL Business!



Tom Walker



Cheryl Watson



Frank Kyne



Watson & Walker Timeline



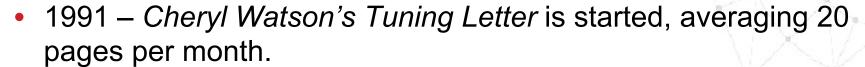
- 1965 Cheryl Watson graduated with a BS in Math & Physics from Portland State University; starts programming in autocoder at Consolidated Freightways on an IBM 1401.
- 1966-1986 Cheryl worked at several locations and companies (several startups, Amdahl, EDS, Morino Associates) primarily teaching, documenting, developing software, and consulting in MVS performance, chargeback, and capacity planning. Active in SHARE and CMG.



Watson & Walker Timeline



- 1986 Tom Walker retires from 21 years in software development, having sold his share of the timesharing company, Dialcom Inc., to ITT.
- 1987 Tom and Cheryl meet, and Tom finances jointly owned Watson & Walker to promote Cheryl's popular classes on performance and capacity planning. Also began consulting on outsourcing and performance.



 1999-2013 – Classes are suspended, but software products BoxScore and GoalTender are developed. Cheryl continues to speak at conferences and consult with customers.



Frank Kyne Timeline



- 1978-1985 Frank starts as an Operator in Irish Life Assurance Company in Dublin, Ireland. After 5 years in Ops, he moves to Tech Support as VM System programmer.
- 1985-1998 Frank joins IBM Ireland as trainee MVS system programmer in software lab in Dublin. Moves on to lead MVS services group, working with nearly every MVS site in Ireland and supporting DR tests, SMS migrations, sysplex implementations, release migrations, floor sweeping, coffee making. . . .



Frank Kyne Timeline



 1998-2014 – Frank moves to ITSO as project leader, responsible for Redbooks and workshops on Parallel Sysplex. During that time, he led project to create prototype for z/OS Health Checker, taught classes on 6 continents, created Customer Configuration Collector, worked on sysplex and availability engagements with 20+ customers, owned production of GDPS product manuals, and led or contributed to more Redbooks than he can remember.



Watson & Walker Timeline



- 2009 Cheryl participates in ITSO Redbook residency. with Frank as the project leader. (She's VERY impressed!)
- 2014 Frank joins Watson & Walker to become new Editor and author of Cheryl Watson's Tuning Letter, providing consulting to customers, and developing classes.





What we do

- Cheryl Watson's Tuning Letter
 - Began in 1991 as a 20-page monthly print newsletter.
 - Now sent via PDF with one subscription per physical site (special prices for multiple sites). Also includes yearly DVD with all issues, presentations, and CPU Charts since 1991.
 - Now published quarterly, but still providing over 240 pages a year with advice on z/OS. (Over 5,000 pages!)
 - Unique in that we have no advertising and we give our own advice, rather than simply stating facts.
 - Cheryl Watson's System z CPU Charts used by many companies for contracts and MIPS ratings – a complete summary of current processors.



What we do

- Cheryl Watson's Advisor Series
 - BoxScore
 - SAS-based product to compare the speed and capacity of a processor after a hardware change, software change, application tuning effort, or configuration change.
 - Shows how individual job steps and transactions behaved.



- Used by both outsourcers and those outsourced to determine normalization factor for billing.
- Used to verify IBM's zPCR estimates.
- GoalTender
 - Provides Workload Manager (WLM) analysis to help tune your WLM policies.
 - Provides Cheryl's warnings and reports.





Our Website

- www.watsonwalker.com
 - Contains Tuning Letter subscription information, prices, order forms, sample issues, and table of contents for 23 years.
 - Downloadable presentations, published articles, and some free software.
 - SMF 101 contains articles from TL that are directed to new sysprogs, and are free to the public.
 - Subscribe to free Cheryl's List and get emails (monthly or so). You can also see the archives with summary of each Tuning Letter and other information we think is important.
 - BoxScore and GoalTender information, brochures, pricing, and order forms.





Frank's lukewarm tips **CF Response times**



- z/OS can do one of two things while a CF request is running:
 - For a synchronous request, it spins, burning z/OS CPU until the CF responds
 - For an asynchronous request, it runs some other work. But there is a z/OS CPU cost associated with the task switching



 The z/OS Heuristic algorithm aims to minimize z/OS CPU cost of each request. If the request is expected to run z/OS CPU Time

30

CF Serv Time

for a "long" time, it will convert the request to an asynchronous one.



Synchronous

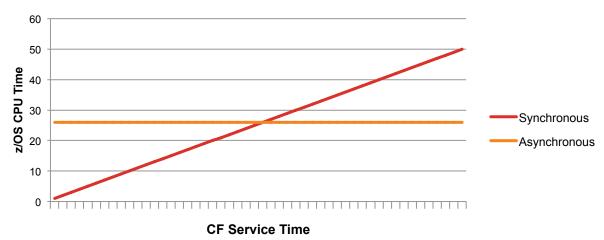
Asynchronous

Frank's lukewarm tips **CF** Response times



How does this relate to cost? As you can see in the chart, the lower the CF service time, the lower the z/OS CPU cost associated with running that request.....

z/OS CPU Cost of CF Requests







Frank's lukewarm tips **CF Response times**



- How do you minimize CF service times:
 - Use fastest link technology available for your configuration
 - When using 12X Infiniband, try to use HCA3 adapters and run them in IFB3 mode.
 - Use dedicated CF engines.
 - If dedicated CF engines are not possible, enable Coupling Thin Interrupts on CF if CF is on zEC12/zBC12 with Driver 15
 - See excellent White Paper (WP102400) by Barbara Weiler, available on IBM Techdocs web site.
 - And Gary King's "Coupling Technology and Overview Planning" session 15203.



Frank's lukewarm tips



Gary King's CF Overhead chart (based on 9 CF requests per second per z/OS MIPS)

The CF CPU
type and link type
indicate what sort
of response time
you can expect.
The better the
response time,
the lower the
overhead

Host	z10 BC	z10 EC	z114	z196	zBC12	zEC12
CF						
z10 BC ISC3	16%	18%	17%	21%	19%	24%
z10 BC 1x IFB	13%	14%	14%	17%	18%	19%
z10 BC 12x IFB	12%	13%	13%	15%	15%	17%
z10 BC ICB4	10%	11%	NA	NA.	NA	NA.
z10 EC ISC3	16%	17%	17%	21%	19%	24%
z10 EC 1x IFB	13%	14%	14%	17%	17%	19%
z10 EC 12x IFB	11%	12%	12%	14%	14%	16%
z10 EC ICB4	10%	10%	NA	NA.	NA	NA.
z114 ISC3	16%	18%	17%	21%	19%	24%
z114 1x IFB	13%	14%	14%	17%	17%	19%
z1 14 12x IFB	12%	13%	12%	15%	15%	17%
z1 14 12x IFB3	NA	NA	10%	12%	12%	13%
z196 ISC3	16%	17%	17%	21%	19%	24%
z196 1x IFB	13%	14%	13%	16%	16%	18%
z196 12x IFB	11%	12%	11%	14%	14%	15%
z196 12x IFB3	NA	NA.	9%	11%	10%	12%
zBC12 ISC3	16%	17%	17%	21%	19%	24%
zBC12 1x IFB	14%	15%	14%	18%	17%	20%
zBC12 12x IFB	13%	13%	12%	15%	14%	17%
zBC12 12x IFB3	NA	NA.	10%	11%	11%	12%
ÆC12 ISC3	16%	17%	17%	21%	19%	24%
zEC12 1x IFB	13%	13%	13%	16%	16%	18%
zEC12 12x IFB	11%	11%	11%	13%	13%	15%
zEC12 12x IFB3	NA	NA	9%	10%	10%	11%

Frank's lukewarm tips



	ISC3	PSIFB 1X	ICB4	PSIFB 12X	PSIFB 12X IFB3	ICP
z10						
Lock	20-30	14-18	8-12	11-15	N/A	3-8
Cache/List	25-40	18-25	10-16	15-20	N/A	6-10
z196						
Lock	20-30	14-17	N/A	10-14	5-8	2-8
Cache/List	25-40	16-25	N/A	14-18	7-9	4-9
zEC12						
Lock	20-30	12-16	N/A	10-14	5-8	2-6
Cache/List	25-40	14-24	N/A	13-17	7-9	4-8

For more information about CF configuration options, see Gary King's excellent presentation – session 15203



Frank's lukewarm tips System-Managed CF Duplexing



- Some CF exploiters rely on System Managed Duplexing to protect structure contents from CF outage (MQ Shared Queues, for example).
- Some CF exploiters rely on System Managed Duplexing to protect structure contents in case of a failure that impacts the structure AND one or more connectors (database lock structures, for example).
- A common scenario for using SM Duplexing with IRLM is when you place a lock structure in an internal, rather than a failure-isolated, CF.
 - This is typically done because an internal CF is perceived as being lower cost than an external one.

Frank's lukewarm tips **System-Managed CF Duplexing**



- However, there is both a z/OS CPU and a CF CPU cost associated with duplexing a request.
 - If you duplex very few requests, the total CPU cost will be low.
 - If you duplex many, the total CPU cost will be higher.
 - At some point, the lines cross, and the cost of using an internal CF combined with SM duplexing becomes higher than the cost of an external CF without duplexing.



Frank's lukewarm tips **System-Managed CF Duplexing**



- The difficulty is in knowing where the lines cross. You need to know:
 - The comparative cost of a simplex vs a duplex request
 - How many duplex requests you are processing a peak times
 - The software cost of a MIPS
- Because of the performance impact of SM Duplexing (plus) higher CF link and subchannel utilization), using SM Duplexing for a very high number of requests has a double impact – financial AND performance.



Frank's lukewarm tips CF gets no love



- How many people are in your storage management group? And how many tools/products do they have to monitor and manage disk performance?
- How large is your CF performance management team? And how many tools/products do they have?



Frank's lukewarm tips CF gets no love



- If you are doing data sharing, you are probably doing 2-3x as many CF requests as disk I/Os.
- If you are doing DB2 data sharing, most of the CF requests happen DURING the transaction - while most DB2 disk I/ Os happen BEFORE or AFTER the transaction, so DB2 can be more sensitive to CF performance than to disk performance.
- A 10 mics response time increase might increase CF service times by 200%, but might only have a 1% impact on disk I/O times.
- Make sure that you are applying your valuable performance skills where they will deliver the most value

Frank's lukewarm tips

What a teenager can teach you about software pricing

- Every time my daughter comes home from the shops, laden down with shopping bags, she patiently explains to me (again!) how much she has saved me by emptying the shelves during a sale.
- What does this have to do with software pricing? The ability to aggregate multiple CECs and use PSLC pricing can result in significant savings.
- In order to be able to qualify for aggregation, the same sysplex must be using >50% of the used MVS MIPS (CP + specialty engines) on each CEC that you want to aggregate.



Frank's lukewarm tips

What a teenager can teach you about software pricing

- What is your natural response to an edict to reduce SW costs? To REDUCE the amount of work being run.
- However, depending on the relative CPU consumption of your different sysplexes, INCREASING the amount of work running in the largest sysplex might SAVE you money if it makes the difference between being able to aggregate a CEC or not.
 - Move work from one sysplex to another
 - Move work between shifts
 - Move work from other platforms (web serving, for example)







- The C compiler has traditionally provided the ability to specify what CEC architecture the program will run on, letting you get the maximum benefit from architectural enhancements in CEC design.
 - But PL/I and COBOL traditionally did not provide the ability to specify what architecture level the resulting program would run on.
- IBM Enterprise COBOL for z/OS 5.1 adds this ability, with the base level being for a 2084.
- IBM Enterprise PL/I for z/OS 4.1 added this ability to PL/I





Frank's lukewarm tips Compiler enhancements in PL/I and COBOL

- Some customers have experienced significant (>10%) CPU savings by re-compiling their COBOL and PL/I programs with an ARCH level that matches their installed CEC. But:
 - You cannot just recompile all your programs and expect identical results.
 - Remember to allow for fallback situations such as if your DR site CEC is at an older architecture level than the production site.
- For more information about migration considerations:
 - Refer to Tom Ross' COBOL sessions at this SHARE sessions 15236, 15246, and 15247
 - Peter Elderon PL/I session 15284



Frank's lukewarm tips Catalog processing



- VLF provides the ability to cache catalog records more. intelligently than In Storage Cache (ISC).
- However, the default when you define a new catalog is to use ISC.
 - If you want a catalog to benefit from VLF, you need to explicitly identify the catalog in the COFVLFxx member of Parmlib



- You can EASILY get a summary of which catalogs are using ISC and which are using VLF by issuing the F CATALOG, ALLOCATED command.
 - Try to make sure that at least your busy catalogs are using VIF

Frank's lukewarm tips Catalog processing



- At the Boston SHARE, Terri Menendez (session 14142) provided some very interesting benchmark results comparing VLF-managed catalogs to RLS-managed catalogs
 - RLS-managed catalogs showed significant elapsed time and CPU time savings for direct processing
 - RLS showed decreased elapsed time but small increase in CPU time for sequential catalog processing



 However, Terri's measurements did not include Enhanced Catalog Sharing. ITSO ran similar experiment to Terri's and found that enabling ECS for a catalog delivered roughly half the elapsed time improvement that RLS provided – but CPU savings were much smaller than RLS – something to help you out until you are ready to move to RLS-managed catalogs.

Frank's lukewarm tips **DSS Compression**



- DFSMSdss currently provides two options to compress the output from a DUMP command:
 - The COMPRESS option has been there forever (if not longer).
 - COMPRESS uses software compression (i.e. CPU-intensive) to compress the output file.
 - The HWCOMPRESS option uses the CMPSC instruction (prezEDC hardware compression) to compress the output file. It also uses a different compression alrgorithm that should result in better compression ratios.



Frank's lukewarm tips **DSS Compression**



- Because HWCOMPRESS uses CMPSC, you might expect that it would use less CPU time than COMPRESS. However, the objective of HWCOMPRESS is to provide better compression, NOT necessarily to save CPU, so depending on the type of data being dumped, using HWCOMPRESS might actually cost you MORE CPU than COMPRESS.
 - Test both options with your data and select the most appropriate one for YOUR data.
- If you have zEDC, DFSMSdss will support zEDC later this year certainly worth investigating if you use either of the DFSMSdss compression options today.







Both CICS and DB2 provide the option to compress their SMF records. This is effective for reducing the volume of data sent to SMF, however it uses software compression, so is CPU intensive.

IF you have zEDC, you can use it to compress the SMF data before it is written to the log stream, and most of the CPU cost of compressing/ decompressing the SMF records is offloaded to the zEDC Express card.



Reminder – zEDC support for BSAM/QSAM will be available at the end of March.

The WSC tool, zBNA (System z Batch Network Analyzer), shows candidate data sets for zEDC. See John Burg's session 15207 on zBNA. And Glenn Anderson covered SMF Logstreams and zEDC in his excellent session 15219 "Top New z/OS Performance Functions **Every Sysprog Should Understand**"

Frank's lukewarm tips Large pages



Since forever, z/OS has used 4KB pages.

z/OS memory sizes are a little larger now (up to 1TB so far) than when I was a lad. But the size of the Translation Lookaside Buffer (TLB), and speed of access to memory (when you get a TLB miss) have not scaled with memory size or processor speed.

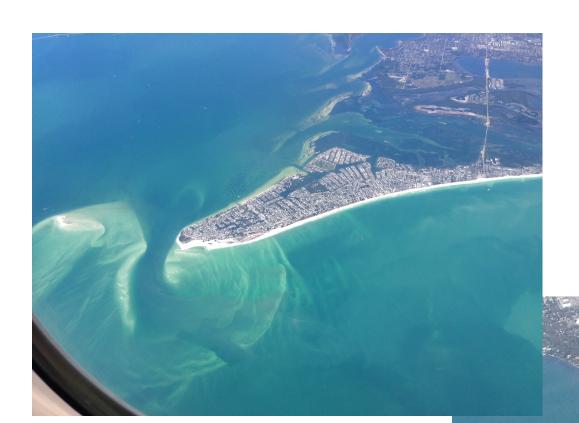
As a result, managing z/OS memory can be relatively more expensive now than in the past.



If you have lots of z/OS memory and are using DB2, IMS, and/or Java, investigate the use of large pages (fixed for DB2, pageable for Java) IBM's estimate of potential savings range from 0 to 6% of overall used 7/OS MIPS

For more info, see Elpida's excellent "z/OS Large Memory – Size Does Matter" session 15142







: SHAKE · . . • in Anaheim

A couple of photos taken by Cheryl while leaving Sarasota this Saturday.



Outsourcing Gotchas

- Benefits of outsourcing:
 - Knowledgeable outsourcer can use elements of scale to reduce the cost of the software, hardware, and staff.
 - Customer no longer needs to staff, house, or care about a part of the business they don't care about. It's a win-win!
- Downside of outsourcing:
 - Only the outsourcer is familiar with contract intricacies; this is all new for the customer. All contracts favor the outsourcer!
 - If customer doesn't retain some technical knowledge, then they can't disagree when contract isn't met.
- See Cheryl Watson's Tuning Letter 2013 No. 1 for full outsourcing article.





- What most agreements contain:
 - Price that customer pays for resources (CPU, DASD, etc.).
- What an agreement should also contain:
 - Baseline measurements (how much CPU, DASD, etc. is being used for current workload).
 - Service Level Agreements (SLAs).
 - Capacity plan.
 - Ongoing benchmark schedule and techniques.
 - Reporting schedule and techniques.
 - Right of access to data.
 - Penalties.





- Pricing Considerations
 - Some contracts charge by LPAR CPU usage
 - Problem Additional LPARs can increase CPU usage in the other I PARs
 - Problem Move to a different processor might not use the right normalization factor
 - Problem Change in CPU busy can change CPU time used by the LPAR
 - Some contracts charge by program CPU time
 - Problems Same as with LPARs.
 - Problem Not all programs change in the same manner
 - Problem Unexpected change in programs



Figure 10 - BoxScore Chart of CPU Changes



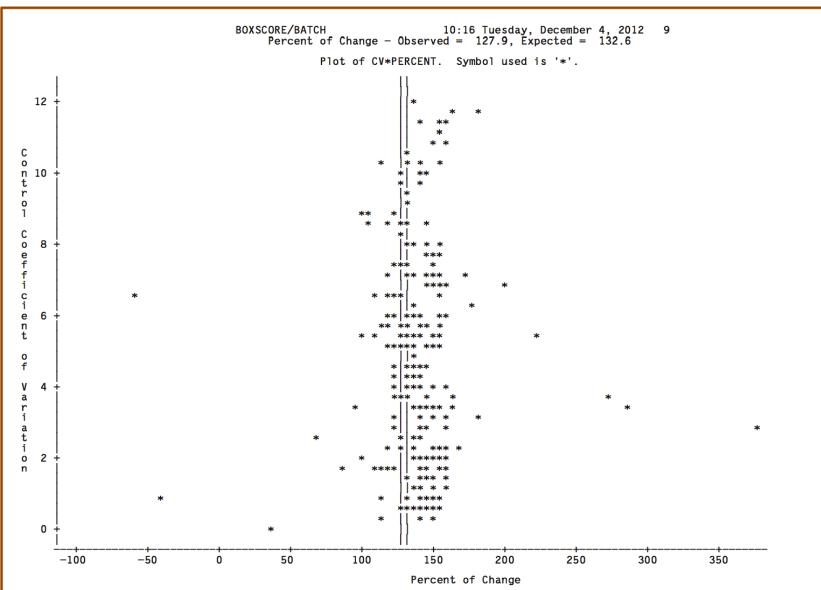
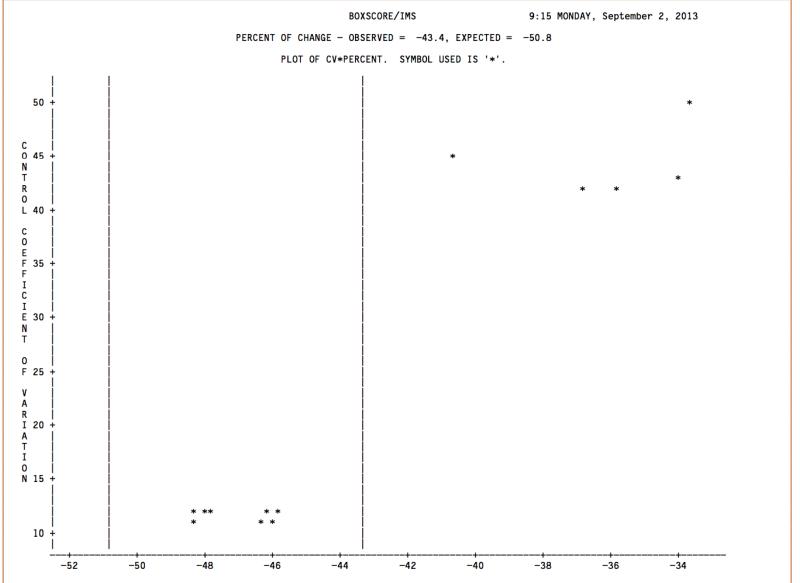


Figure 1 – BoxScore Plot of Percent of Change









- Pricing Considerations
 - Some contracts are set up to have serious penalties if usage exceeds a certain amount. This is to protect the outsourcer from unexpected capacity requirements and relieves them of commitment to meet SLAs. That's only fair.
 - But the contract must allow negotiation for new workloads, new technologies, new capacity growth with planning and without serious penalties. How can an outsourced customer initiate new technologies such as mobile, big data, and cloud without exorbitant prices?
 - Some contracts are essentially loans (usually benefits the customer VIP signing the contracts) – very low for a year, then excessive charges for succeeding year.





- Baseline Measurements
 - Must define when baseline is taken (length of interval, peak hours, peak days, peak weeks, peak month?).
 - Must define contents of baseline (which system address) spaces are included, which monitors are used, does baseline change during the year?)
 - Must define how baseline is measured (what tools, which fields are used)





- Service Level Agreements (SLAs) must be set for peak periods using percentiles (not averages).
 - Example: One contract said average TSO response was to be 0.4 seconds during the day. A year into the contract, users were complaining when response times at 10 am were over 6 seconds (but daily average was still .4). Old peak period response was 1 second, but the outsourcer was meeting the contract.





- Capacity Plan
 - Baseline should include not only recent year, but historical data (i.e. has business traditionally grown 5% a year?)
 - Before contract signed, a capacity plan should be made (and continued to be monitored)





- Ongoing Benchmarks and Reporting
 - The system WILL change! You need one or more benchmarks and/or reports to determine when changes occur.
 - Applications change, databases grow or shrink, files get tuned, PTFs are applied, work in other LPARs change, user hours change, etc. – the system seldom remains unchanged.
 - How will you know? Historical tracking by applications, by LPAR, by online system CPU usage, by I/O changes, by CPU per I/O, and more.
 - Some BoxScore customers run a benchmark every week.





- Right of Access
 - Customer must have access to raw data (e.g. SMF data) that is used for billing. Many outsourcers restrict this.
 - Customer must have right to audit and right to have their own consultants.
 - Customer must reject any gag restrictions. Almost all outsourcers restrict this. It's one of the only paths available when all communication breaks down.





Penalties

- There should be definite and useful penalties for not meeting SLAs.
- Poor example if SLAs are missed one day, then outsourcer has seven days to correct the problem (result – only one day of good service).
- Good example if SLAs are missed one day, then customer doesn't pay for that day.



- Poor example if customer workload increases by 20% one day and SLAs are missed, then outsourcer gets penalized.
- Good example if customer workload increases by 20% one day, then SLAs are ignored and customer pays more that day.











Document savings

- Measure before, measure after, determine difference. If you don't know where you are, you won't know if you changed.
- Determine the value of a CPU minute or hour to your management. Get them to agree.
- Learn the costs of your hardware, software, peripherals.
- Don't forget the cost of people. If you can save an hour of someone's time, then know what that's worth.
- Publish results of your savings don't hesitate to pat yourself on the back.





- Money Savings Tips
 - Reserve new capacity
 - After an upgrade, hold part of new capacity back goal should be to run between 95-100% busy during peak times.
 - If you don't, people may develop different techniques of working, and end up increasing the CPU usage.
 - If you don't, jobs will cost less until new capacity comes along and then they'll go up (about 3-5% change in CPU per job for every 10% change in CPU busy). Don't set normalization factor at this time!
 - Latent demand could increase CPU use and you'll need more capacity even sooner.





- Money Savings Tips
 - Re-evaluate pricing yearly
 - More savings are found by re-evaluating software pricing than almost anything else.
 - Change LPAR configurations to minimize software costs. See Kath Walsh's session 15200 this week.



- See Al Sherkow's sessions (Boston SHARE) on pricing. Investigate his software.
- Investigate capacity provisioning.
- Understand that chargeback to customer will never be accurate.
- Learn early and well that jobs run during peak time will cost more than those during off-peak.
- See my session 15274 on "Why is the CPU time so variable?"





- Money Saving Tips
 - Don't stay on old machines maintenance and new purchase/lease can often cost less than retaining non-current machines.
 - Consider adding zIIPs (not zAAPs) instead of adding more CPs. (More applications now move part of workload to specialty engines.) Try out ZAAPONZIIP.



 Don't be afraid to negotiate with all vendors. The more expensive the product, the better chance you have of reducing the price.





- My favorite (and most repeated) Tuning Letter recommendations:
 - Turn on HiperDispatch Can save CPU time, can cover your mistakes
 - Use zPCR Use it before any configuration change; can identify the best machine for an upgrade



- Set MSO=0.0 This can save CPU because it reduces swapping and dispatch switching
- Set DDCONS=NO This can reduce CPU time during shutdowns and interval breaks; improves online response time





- My favorite (and most repeated) Tuning Letter recommendations:
 - Keep the number of LPs to a minimum Even if HiperDispatch is turned on
 - Install and learn z/OSMF This is the future; don't cling to the old ways of doing things
 - Keep monitors at a large, rather than minimum frequency (e.g. use RMF intervals of 15 minutes instead of 10; use 2second RMF cycle time instead of 1-second; reduce refresh rates)
 - Use system-determined or large buffers in almost all cases.
 - Use RMF overview reports to identify just critical situations.





See you in Pittsburgh!







If you're in trouble, we can help!



- Watson & Walker, Inc.
 - Practical and educational newsletter Cheryl Watson's Tuning Letter
 - CPU Charts for all System z processors
 - BoxScore product to determine effect of changes
 - Consulting (Frank, Cheryl, other consultants):
 - Parallel Sysplex issues
 - Workload Manager (WLM) issues
 - Performance analysis
 - Capacity planning
 - Outsourcing contract review and negotiation
 - Courses (Frank, Cheryl, other teachers)
 - Same areas as above

