Value-4IT

GSE/UKCMG zCMPA Working Group

Performing Due Diligence For A zSeries Server Upgrade Proposal

Collaboration, Balancing & Synergy Thoughts

Vendor Interaction: Obvious issues arise if you totally outsource zSeries server sizing to the vendor; or if you deliver a final configuration to the vendor, without their input. Utilize various data profiles to find a fit-for-purpose configuration.
Observation: Safeguard the vendor "underwrites" with due-diligence & QA.

Hardware Resource vs. Software Cost: Does the Capacity Planner know all IBM software (E.g. MLC, IPLA, et al) & 3rd party cost ISV mechanisms? What is more important to your business; the Peak Hour, R4HA, SLA/KPI metrics, or all three?
Observation: Typically Sub-Capacity (SCRT) software pricing is always a factor.

Customer & Vendor Teamwork: Often the vendor has internal tools that can refine the "size & shape" of server configurations. Only the customer truly knows their business; data sources like CPU MF fully qualify real-life CPU usage.
Observation: Fully utilize all of the skills, data sources & tools at your disposal.

A Living Process: True synergies only evolve with long-term team-working, especially with vendors. Make sure your vendor is always there for "tweaks & repairs", as the perfect server configuration was never delivered on day #1!
Observation: Neither the customer nor the vendor benefits from poor results.

As a customer, never ever isolate yourself from the vendor; talk!

A Real-Life zSeries Server Upgrade: Overview

Current Configuration

*2 * z10 BC 2098-x0n Servers*
*6 * GP CPUs (3 per z10)*
~230 MSU Installed Capacity
*2 * zIIP Engines (1 per z10)*
No SYSPLEX Coupling
SCRT @ ~200 MSU Per Month
Fixed cost ELA @ ~<= 220 MSU
3rd Party (Warm-Site) DR Service
~10% Annual Capacity Increase

Customer-Vendor Handover

Customer Provided:

- *Type 7n records (1 year)*
- *SCRT reports (1 year)*
- *Capacity usage (1 year)*
- *Capacity forecast (3 years)*

Vendor Requested:

- *Budget for 1-3 year period*

Other Considerations:

- *Customer had a paid-for quarterly capacity & SCRT review with the vendor*

Proposed Configuration

*2 * zBC12 2828-x0n Servers*
*6 * GP CPUs (3 per zBC12)*
~300 MSU Installed Capacity
*2 * zIIP Engines (1 per zBC12)*
No SYSPLEX Coupling
SCRT @ ~200 MSU Per Month
Fixed cost ELA @ ~<= 220 MSU
3rd Party (Warm-Site) DR Service
~10% Annual Capacity Increase

The customer asked their trusted zSeries partner for an upgrade!

High Level zSeries Upgrade Proposal: QA Review

Ubiquitous
SMF Type 7n
Requested
SMF Type 113
(CPU MF)

SCRT Reports
CPU Forecasts

PERFTECHPRO
zDynaCap

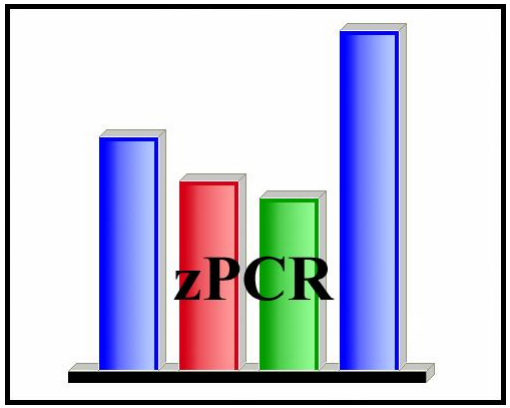

Market place software includes: ASG-PERFMAN, BMC Capacity Management for Mainframes, CA MICS Resource Management, EPV for z/OS, Tivoli Performance Modeler for z/OS, In-House via MXG, SAS, WPS, et al

The customer had no software tools installed as they had "smart-sourced" capacity planning to their partner. Within 1 day, we had analysed 3 Months of SMF, producing CPU & SCRT forecasts, highlighting WLM & zIIP tuning opportunities. Customer had not activated SMF 113 (CPU MF) & so this facility was switched on. After 1 Month we further refined the "Capacity & Performance" data with CPU MF real-life business workload data. ~2 days work over a ~1 Month period reduced server upgrade costs by ~50%. The customer now performs capacity planning themselves…

The overall process is mandatory; software tools are arbitrary!

zSeries Upgrade Review: SCRT Observation #1

z10 BC Server #1 - EWLC		z10 BC Server #2 - EWLC		Data Centre Summary - 2 * z10 BC Servers	
Product Name	**SCRT R4HA**	**Product Name**	**SCRT R4HA**		
CICS TS for z/OS	~100 MSU	CICS TS for z/OS	~90 MSU	*Observation: MSU costs reduce as usage increases; the most expensive MSU costs are in the lower SCRT/MLC bands. Aggregate MSU wherever possible into "product" containers. Why was this customer deploying 2 uncoupled z10 BC servers with no SYSPLEX coupling?*	
COBOL V4	~50 MSU	COBOL V4	~90 MSU		
DB2 UDB for z/OS	~100 MSU	DB2 UDB for z/OS	~90 MSU		
Enterprise PL/I for z/OS		Enterprise PL/I for z/OS	~25 MSU		
IMS Database Manager		IMS Database Manager	~25 MSU		
WS MQ Base for z/OS	~100 MSU	WS MQ Base for z/OS	~90 MSU		
z/OS V1 Base	~100 MSU	z/OS V1 Base	~90 MSU		
Monthly SCRT Total	~£92,000.00	***Monthly SCRT Total***	~£105,000.00	***Monthly SCRT Total***	~£197,000.00
Annual SCRT Total	~£1,104,000.00	***Annual SCRT Total***	~£1,260,000.00	***Annual SCRT Total***	~£2,364,000.00
zBC12 Server #1 - AEWLC		**zBC12 Server #2 - AEWLC**		**Data Centre Summary - 2 * zBC12 Servers**	
Product Name	**SCRT R4HA**	**Product Name**	**SCRT R4HA**		
CICS TS for z/OS	~100 MSU	CICS TS for z/OS	~90 MSU	*Observation: The vendor didn't provide the customer with any software ROI benefits for just upgrading the hardware (z10 to zBC12) technology. Clearly the benefits were significant, but was the hardrware configuration still the best customer option?*	
COBOL V4	~50 MSU	COBOL V4	~90 MSU		
DB2 UDB for z/OS	~100 MSU	DB2 UDB for z/OS	~90 MSU		
Enterprise PL/I for z/OS		Enterprise PL/I for z/OS	~25 MSU		
IMS Database Manager		IMS Database Manager	~25 MSU		
WS MQ Base for z/OS	~100 MSU	WS MQ Base for z/OS	~90 MSU		
z/OS V1 Base	~100 MSU	z/OS V1 Base	~90 MSU		
Monthly SCRT Total	~£80,000.00	***Monthly SCRT Total***	~£90,000.00	***Monthly SCRT Total***	***~£170,000.00***
Annual SCRT Total	~£960,000.00	***Annual SCRT Total***	~£1,080,000.00	***Annual SCRT Total***	***~£2,040,000.00***
z10 BC to zEC12 Server Upgrade Software Cost Benefit Summary				***Monthly SCRT Savings***	***~£27,000.00***
				Annual SCRT Savings	***~£324,000.00***

IBM MLC software costs are significant; understand them, fully!

zSeries Upgrade Review: zIIP & WLM Usage

WLM Class	CP-SU-0	CP-SU-1	CP-SU-2	CP-SU-3	CP-SU-4	CP-SU-5	CP-SU-SD	CP-SU	zIIP	zIIP->CP
BATCHHOT	0.00	0.00	0.00	824.56	0.00	0.00	0.00	824.56	0.00	0.00
BATCHMED	0.00	0.00	0.00	2753.91	0.00	0.00	0.00	2753.91	0.00	0.00
BATCHLOW	0.00	0.00	0.00	83.36	0.00	0.00	0.00	83.36	0.00	0.00
DB2PROD1	0.00	0.00	0.00	12304.04	0.00	0.00	0.00	12304.04	462.68	0.80
DB2PROD2	0.00	0.00	0.00	0.00	35.13	0.00	0.00	35.13	0.00	0.00
DB2L247	0.00	0.00	0.00	0.00	0.00	0.00	21027.41	21027.41	7.64	0.00
OLPROT	0.00	11688.43	0.00	0.00	0.00	0.00	0.00	11688.43	0.00	0.00
OLHIGH	0.00	0.00	25552.46	0.00	0.00	0.00	0.00	25552.46	***8976.19***	***2734.80***
OLLOW		0.00	0.00	127.25	0.00	0.00	0.00	127.25	0.00	0.00
STCHIGH	0.00	0.00	0.00	137.26	0.00	0.00	0.00	137.26	0.00	0.00
STCMED	0.00	3386.54	0.00	0.00	0.00	0.00	0.00	3386.54	484.08	4.02
STCLOW	0.00	0.00	0.00	94.58	0.00	0.00	0.00	94.58	0.00	0.00
TSO	0.00	0.00	3.61	0.00	0.00	0.00	0.00	3.61	0.00	0.00
SYSTEM	2075.85	0.00	0.00	0.00	0.00	0.00	0.00	2075.85	0.00	0.00
SYSSTC	6344.00	0.00	0.00	0.00	0.00	0.00	0.00	6344.00	2259.45	15.06
Review WLM (Type 72) SMF - Another Perspective; zIIP Eligible CPU SU Rediredted Back To General Purpose CP										***30.47%***

- *Tuning the ZIIPAWMT parameter can optimize zIIP usage, but sometimes more zIIP capacity is required. Some say 50% zIIP utilization is "goodness" and allocating an equal number (E.g. 1:1) of General Purpose and zIIP engines per LPAR are good "rules of thumb".*
- *zIIP specialty engines run at full speed & supported ratio for zxC12 is 2:1 (zIIP to GP).*
- *Arguably a workload runs for zero cost on a zIIP (no software charges for CPU usage).*
- *As time goes by, more workloads (SRB enclave) become zIIP eligible; check ISV software.*

Don't overlook zIIP capacity planning & performance monitoring

zSeries Evolution: Relative Nest Intensity (RNI)

With CPU MF (SMF 113) data (z10+), the ability to gain an insight into the interaction of workload & hardware exists. RNI (z196+) based methodology for workload matching is now the default in zPCR.

z10 BC

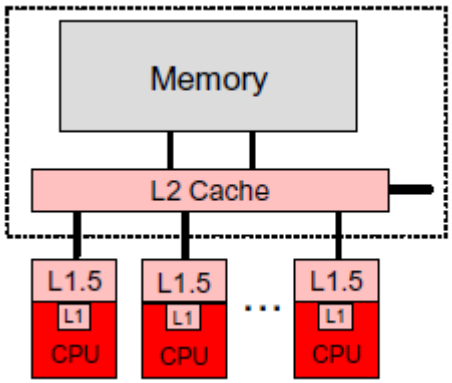

Nest

- *CPU @ 3.5 GHz*
- *Cache Hierarchy:*
  - ✓ L1 private 64 KB i + 128 KB d
  - ✓ L1.5 private 3 MB
  - ✓ L2 shared 24 MB/Book

RNI Observations

- *Activity beyond private cache(s) is the most sensitive area (due to longer latencies)*
- *RNI: reflects activity distribution & latency to chip & book level caches & memory*
- *Level 1 cache miss percentage important*
- *Fewer address spaces = reduced contention*
- ✓ *Low RNI (closer to chip) means faster execution*

zBC12

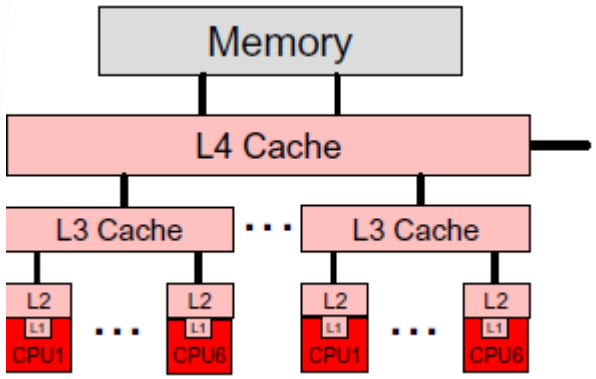

- *CPU @ 4.2 GHz:*
  - ✓ Enhanced out-of-order
- *Cache Hierarchy:*
  - ✓ L1 private 64 KB i + 96 KB d
  - ✓ L2 private 1 MB i + 1 MB d
  - ✓ L3 shared 24 MB/Chip
  - ✓ L4 shared 192 MB/Book

Workload behaviour will change as zSeries server chips evolve…

zSeries Upgrade Review: Multi-Faceted Process

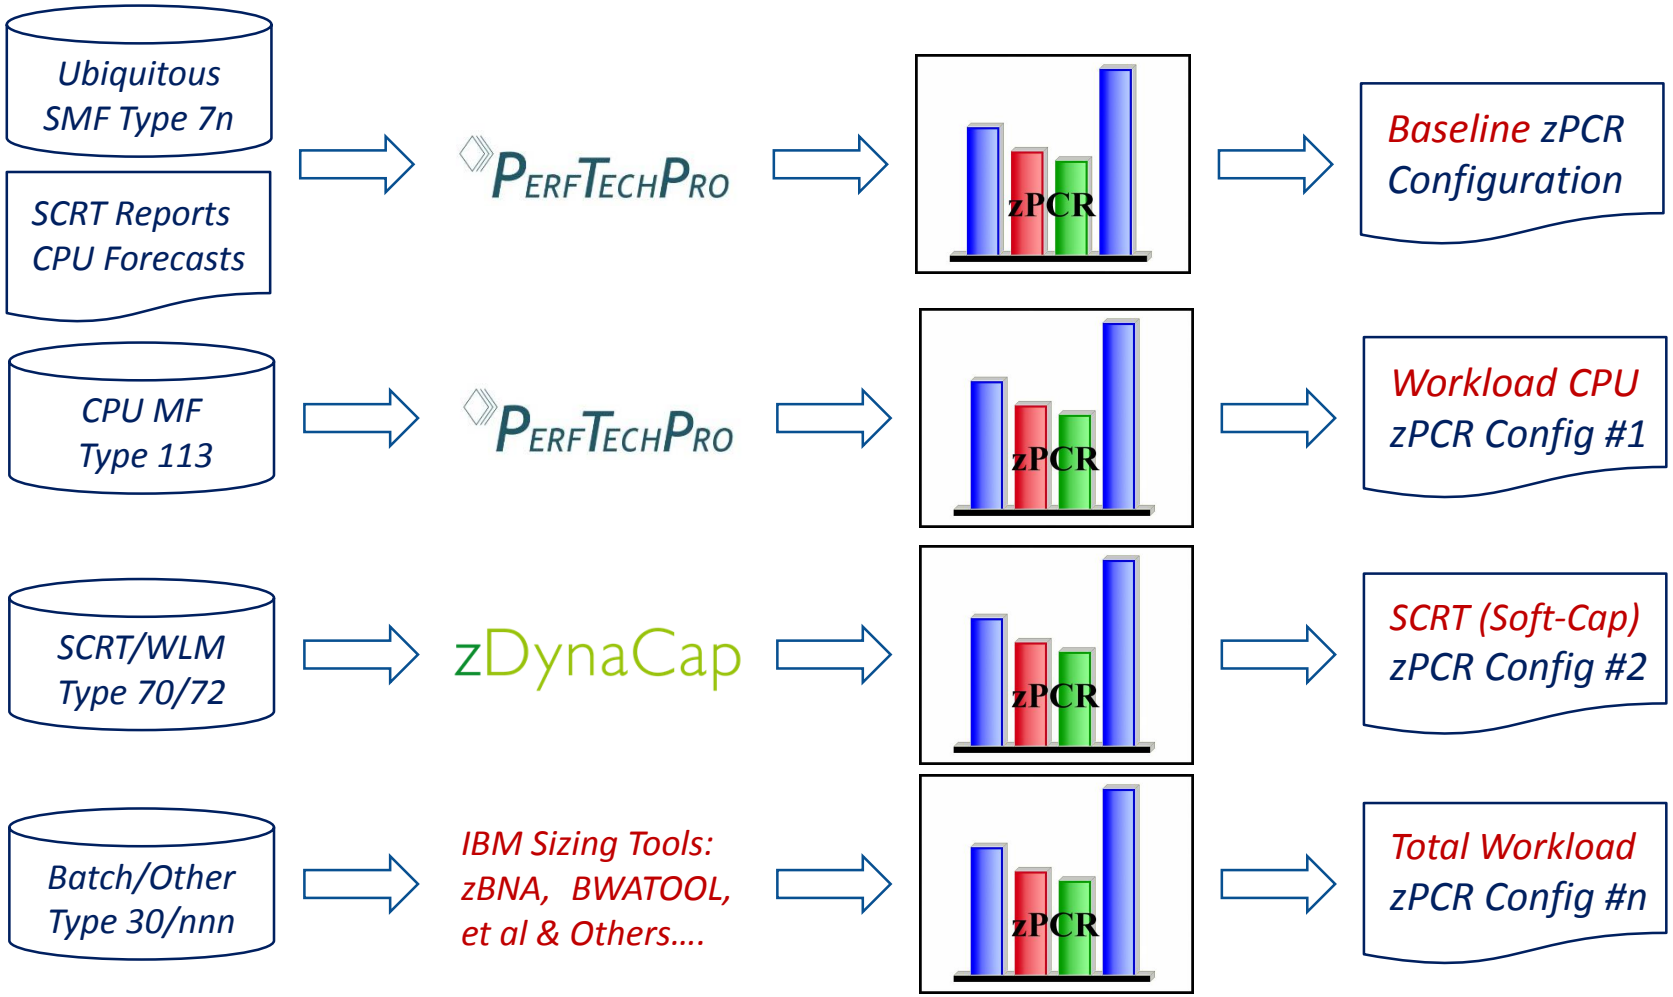

Practice doesn't make perfect; perfect process might get close!

A Real-Life zSeries Server Upgrade: Summary

Proposed Configuration

*2 * zBC12 2828-x0n Servers*
*6 * GP CPUs (3 per zBC12)*
~300 MSU Installed Capacity
*2 * zIIP Engines (1 per zBC12)*
No SYSPLEX Coupling
SCRT @ ~200 MSU Per Month
Fixed cost ELA @ ~<= 220 MSU
3rd Party (Warm-Site) DR Service
~10% Annual Capacity Increase

Customer Lessons Learnt

- *Capacity Planning must be performed internally.*
- *CPU MF data is required for accurate configurations.*
- *Comprehensive knowledge of WLC-SCRT is mandatory.*
- *Leverage from IBM vendor tools to further understand the business workload & further refine capacity plan.*
- *Ask the right questions of your vendor, in terms of qualification/quantification.*

Target Configuration

*1 * zBC12 2828-x0n Servers*
*6 * GP CPUs (25% spare capacity)*
~300 MSU Installed Capacity
*2 * zIIP Engines (30% spare capacity)*
No SYSPLEX Coupling
SCRT @ ~200 MSU Per Month
No ELA; monthly variable SCRT
3rd Party (Warm-Site) DR Service
~10% Annual Capacity Increase

CPU MF, WLM, R4HA & zIIP analysis delivered ~30% TCO savings

zSeries Upgrade Review: SCRT Observation #2

Workload Consolidation - 1 zBC12 Server (6 GP + 2 zIIP Engines) With Optimized Soft-Capping Capacity Balancing Option				
zBC12 Server #1 - AEWLC		***zDynaCap Soft-Capping (~10% Savings)***		*Observation: The customer had no requirement to separate their workloads, the only reason for separate servers was historical, based upon maximum number of CPU engines supported. Consolidating the workloads generated significant SCRT cost savings, even projecting worst case scenario for MSU resource usage. Because workloads were being combined, optimized soft-capping was beneficial, both for SCRT cost & overall performance; differentiating between time critical & non time critical workloads.*
Product Name	***SCRT R4HA***	***Product Name***	***SCRT R4HA***	
CICS TS for z/OS	~190 MSU	CICS TS for z/OS	~175 MSU	
COBOL V4	~150 MSU	COBOL V4	~135 MSU	
DB2 UDB for z/OS	~190 MSU	DB2 UDB for z/OS	~175 MSU	
Enterprise PL/I for z/OS	~25 MSU	Enterprise PL/I for z/OS	~25 MSU	
IMS Database Manager	~25 MSU	IMS Database Manager	~25 MSU	
WS MQ Base for z/OS	~190 MSU	WS MQ Base for z/OS	~175 MSU	
z/OS V1 Base	~190 MSU	z/OS V1 Base	~175 MSU	
Monthly SCRT Total	~£135,000.00	***Monthly SCRT Total***	~£127,000.00	
Annual SCRT Total	~£1,620,000.00	***Annual SCRT Total***	~£1,524,000.00	
Consolidated Workload Savings		***Optimized Soft-Capping Savings***		
Monthly SCRT Savings	***~£35,000.00***	***Monthly SCRT Savings***	***~£8,000.00***	
Annual SCRT Savings	***~£420,000.00***	***Annual SCRT Savings***	***~£96,000.00***	
For MLC software alone, the projected cost reduction was ~£500,000 per annum, by consolidating workloads with soft-capping!				

Only pay the lowest MSU cost, consider soft-capping benefits…

zSeries Capacity Planning: Going Forward #1

- *zPCR (Processor Capacity Reference) is a Windows PC based tool, designed to provide capacity planning insight for IBM System z processors running various z/OS, z/VM, z/VSE, Linux, zAware, & CFCC workload environments on partitioned hardware. Capacity results are based on IBM's most recently published LSPR data for z/OS. Capacity is presented relative to a user-selected Reference-CPU, which may be assigned any capacity scaling-factor & metric.*
- *zSoftCap (Software Migration Capacity Planning Aid) is a Windows PC based tool, designed to assess the effect on IBM System z processor capacity, when planning to upgrade to a more current operating system version and/or major subsystems versions (E.g. Batch, CICS, DB2, IMS, Web & System). zSoftCap assumes that the hardware configuration remains constant while the software version or release changes. The capacity implication of an upgrade for the software components can be assessed independently or in any combination.*
- *zBNA (System z Batch Network Analysis) is a Windows PC based tool, designed to understand the batch window, for example:*
  - ✓ *Perform "what if" analysis & estimate the CPU upgrade effect on batch window*
  - ✓ *Identify job time sequences based on a graphical view (filter by CPU time, job class, et al)*
  - ✓ *Review the resource consumption of all the batch jobs (down to individual step level)*
  - ✓ *Identify candidate jobs for running on different processors*
  - ✓ *Identify jobs with speed of engine concerns (top tasks %)*

Generally available (IBM web site) tools for further refinement

zSeries Capacity Planning: Going Forward #2

- *zCP3000 (Performance Analysis & Capacity Planning) is a Windows PC based tool, designed for performance analysis & capacity planning simulations for IBM System z processors, running various SCP & workload environments. It can also be used to graphically analyse logically partitioned processors & DASD configurations. Input from customer data via a separate tool (I.E. z/OS SMF via CP2KEXTR, VM Monitor via CP3KVMXT, VSE CPUMON via VSE2EDF).*
- *zPSG (Processor Selection Guide) is a Windows PC based tool, providing sizing approximations for IBM System z processors intended to host a new application, implemented using popular, commercially available software products (E.g. WebSphere, DB2, ODM, Linux Apache Server).*
- *BWATOOL (Batch Workload Analysis Tool) is a Windows PC based tool analysing SMF type 30/70 data, producing a report showing how long batch jobs run on the currently installed processor. Both CPU time & elapsed time are reported. Similar results can then be projected for any IBM System z processor model.*
- *zMCAT (Migration Capacity Analysis Tool) is a Windows PC based tool, comparing performance of production workloads before & after migration of the system image to a new processor, even if the number of engines on the processor has changed. Workloads for which performance is to be analysed must be carefully chosen because the power comparison may vary considerably due to differing use of system services, I/O rate, instruction mix, storage reference patterns, et al. This is why customer experiences are unique from an internal throughput ratio (ITRR) based on LSPR benchmark data.*

Internal (IBM sales team) tools for further refinement

Results via Collaboration, Balancing & Synergy

Vendor Interaction: The vendor has specialist knowledge of the zSeries Server & associated software costs, with access to internal sales tools. Collaborate with the vendor to QA & refine the capacity plan, until both parties reach agreement. Observation: Ideally customers challenge the vendor with meaningful questions.

Hardware Resource vs. Software Cost: The customer must fully understand the various IBM software pricing mechanisms, especially WLC (SCRT) related. Ideally this skill is located in the technical department & not the commercial/legal team. Observation: IBM MLC software costs are a significant factor of zServer TCO.

Customer & Vendor Teamwork: Nobody knows the business or zServer workload better than the customer. The vendor should safeguard the customer has access to all available IBM zServer knowledge & tools for the capacity planning process. Observation: The vendor must educate the customer during server upgrade bids.

A Living Process: zSeries server technologies evolve, generating a requirement to refine capacity planning processes. Instead of waiting for the compelling event, both the customer & vendor should periodically review real-life zSeries statistics. Observation: Be careful how the small things grow; make sure everybody wins.

The best relationships are based on trust, respect & compromise

zSeries Server Capacity Planning: Useful Web Links

- ✓ *zPCR (Processor Capacity Reference):*
www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381
- ✓ *zSoftCap (Software Migration Capacity Planning Aid):*
www-01.ibm.com/support/docview.wss?uid=tss1prs268
- ✓ *zBNA (System z Batch Network Analysis):*
www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132
- ✓ *Setting Up and Using the IBM System z CPU Measurement Facility with z/OS (REDP-4727-00):*
www.redbooks.ibm.com/abstracts/redp4727.html?Open
- ✓ *z/OS CPU MF Enablement Education:*
www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4922
- ✓ *CPU MF Overview and WSC (z196) Experiences: SHARE 2012:*
www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TC000066
- ✓ *zTPM (Tivoli Performance Modeler - Can be used as an IBM Sales Tool):*
www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=OC&subtype=NA&htmlfid=897/ENUS5698-A18
- ✓ *Value-4IT Capacity Planning Blog Entry:*
www.value-4it.com/wp/?p=106
- ✓ *PerfTechPro Capacity Planning & Performance Measurement Solution:*
www.perftechpro.com/products.html
- ✓ *zDynaCap Automated Soft-Capping & Capacity Balancing Solution:*
www.zit-consulting.com/zdynacap-2

Even if you're not a customer, ask us a question, we'll try to help!

Backup Slide: MLC Bands & WLM MSU Utilization Metric

EWLC (z800*, z890*, z9 BC, z 10 BC)		AEWLC (z114, zBC12)		VWLC (z900, z990, z9 EC, z 10 EC)		AWLC (z196, zEC12)	
MLC Level	MSU Range	MLC Level	MSU Range	MLC Level	MSU Range	MLC Level	MSU Range
Base	3	Base	3	Level 0	4-45	Base	3
Level 1	4-17	Level 1	4-17	Level 1	46-175	Level 0	4-45
Level 2	18-30	Level 2	18-30	Level 2	176-315	Level 1	46-175
Level 3	31-45	Level 3	31-45	Level 3	316-575	Level 2	176-315
Level 4	46-87	Level 4	46-87	Level 4	576-875	Level 3	316-575
Level 5	88-175	Level 5	88-175	Level 5	876-1315	Level 4	576-875
Level 6	176-260	Level 6	176-260	Level 6	1316-1975	Level 5	876-1315
Level 7	261+	Level 7	261-315	Level 7	1976+	Level 6	1316-1975
* Qualified Parallel Sysplex		Level 8	316+			Level 7	1976+

*WLM (Workload Manager) is responsible for taking MSU utilization samples for each LPAR in 10-second intervals. Every 5 minutes, WLM documents the highest observed MSU sample value from the 10-second interval samples. This process always keeps track of the past 48 updates taken for each LPAR. When the 49th reading is taken, the 1st reading is deleted, and so on. These 48 values continually represent a total of 5 minutes * 48 readings = 240 minutes or the past 4 hours (I.E. R4HA). WLM stores the average of these 48 values in the WLM control block RCT.RCTLACS. Each time RMF (or BMC CMF equivalent) creates a Type 70 record, the SMF70LAC field represents the average of all 48 MSU values for the respective LPAR a particular Type 70 record represents. Hence, we have the "4 Hour Rolling Average". RMF gets the value populated in SMF70LAC from RCT.RCTLACS at the time the record is created.*

SCRT also uses the Type 70 field SMF70WLA to ensure that the values recorded in SMF70LAC do not exceed the maximum available MSU capacity assigned to an LPAR. If this ever happens (due to a soft capping or otherwise) SCRT uses the value in SMF70WLA instead of SMF70LAC. Values in SMF70WLA represent the total capacity available to the LPAR.