Value-4IT

GSE/UKCMG zCMPA Working Group

IBM z13 Technology Transition Offerings

AWLC Price Performance Benefits

zSeries Server CPU Chip Evolution: What Next?

IBM Mainframe Charter 2003: IBM lowered MSU values incorporated in z990 microcode by ~10%, resulting in savings for IBM zSeries MSU based software products pricing. The reduced MSU metrics do not change machine performance. Observation: z10 was the last server to deliver this "technology dividend".

Server	GHz	GA Date	# Cores	SOI nm	Max CPUs	SW MSU	MIPS/MSU	# Threads
z900	0.77	2000	1	189	16	535	5.9	1
z990	1.20	2003	2	130	32	1512	6.6	1
z9 EC	1.70	2005	2	90	54	2409	7.3	1
z10 EC	*4.40*	2008	4	65	64	3739	7.7	1
z196	5.20	2010	4	45	80	6140	***AWLC***	1
zEC12	5.50	2012	6	32	101	9194	***AWLC***	1
z13	*5.00*	2015	8	22	141	13078	***AWLC***	***2***

A Living Process: For several generations of zSeries Servers (I.E. z990, z9,z10), the ~10% technology dividend delivered a simple mechanism for increased CPU power & lower MLC software costs. From z196 onwards, AWLC is the only option. Observation: The CPU second is the only reliable power metric; SMT considerations

CPU chip technology will introduce new WLC pricing mechanisms!

zEC12 vs. z13 CPU Chip: Cache Improvements

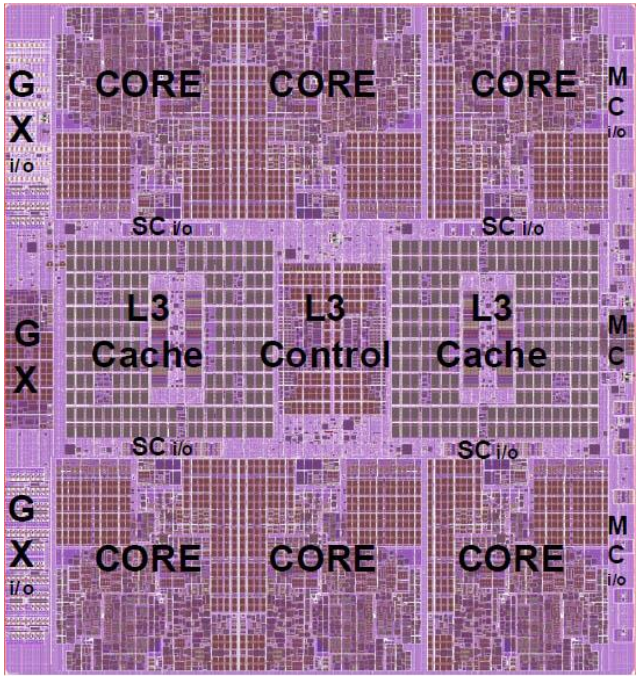

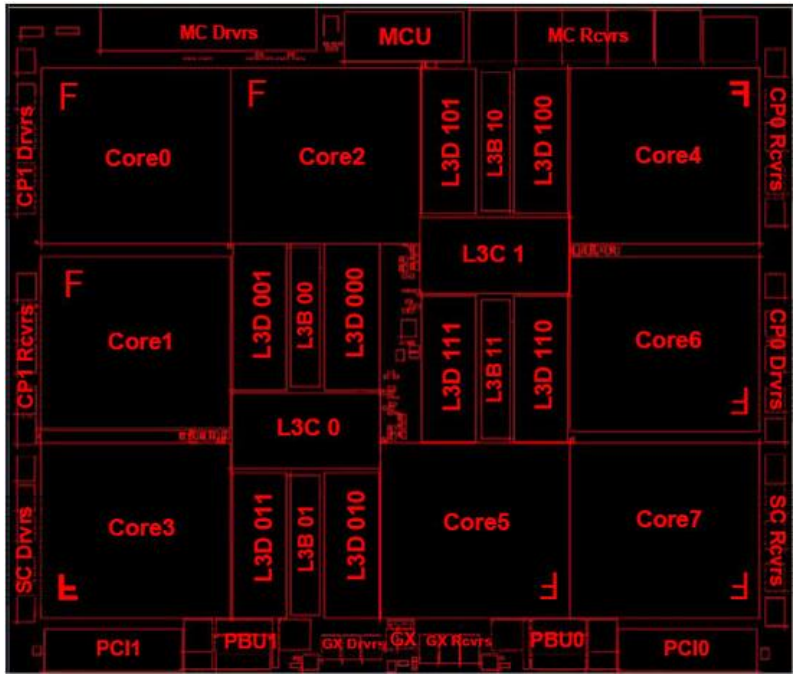

❖ *zEC12 Cache Hierarchy:*

- ✓ L1 private 64 KB i + 96 KB d
- ✓ L2 private 1 MB i + 1 MB d
- ✓ L3 shared 48 MB
- ✓ L4 shared 384 MB

❖ *z13 Cache Hierarchy:*

- ✓ L1 private 96 KB i + 128 KB d
- ✓ L2 private 2 MB i + 2 MB d
- ✓ L3 shared 64 MB
- ✓ L4 shared 480 MB plus
  224 MB NIC Directory

CPU chips now work smarter not harder; delivering better RNI!

IBM z13: Technology Update Pricing (TU3)

When stand-alone z13 servers are priced with AWLC, or when all the servers in an aggregated Sysplex or Complex are z13 servers priced with AWLC, these servers receive a reduction to AWLC pricing which is called Technology Update Pricing for the IBM z13 (TU3). Quantity of z13 Full Capacity MSUs for a stand-alone server, or the sum of Full Capacity MSUs in an actively coupled Parallel Sysplex or Loosely Coupled Complex made up entirely of z13 servers:

MSU Range	AWLC % Discount
4-45	4%
46-315	8%
316-1315	9%
1316-2676	10%
2677-5476	12%
5477+	14%

Example 1: Two Standalone z13 Servers

Machine #1 is a z13 model 411 server with 281 MSU, so it receives an 8.0% reduction in AWLC billing. Machine #2 is a z13 model 718 server with 2584 MSU, so it receives a 10.0% reduction in AWLC billing.

Example 2: Three Parallel Sysplex Coupled z13 Servers

Machine #1 is a z13 model 710 server with 1632 MSU Machine #2 is a System z13 model 711 server with 1764 MSU & Machine #3 is a z13 model 712 with 1891 MSU, a cumulative total of 5287 MSU. The Parallel Sysplex will receive a 12.0% reduction in aggregated AWLC billing.

Put simply, TU3 eligibility is configured General Purpose Processor (GPP) MSU; excluding MSU capacity related to specialty engines, unassigned capacity, & temporary capacity such as, but not limited to, "Capacity Back Up", "Capacity for Planned Events" & "On/Off CoD".

A simple AWLC update; as always, aggregated MSU is best.

IBM z13: AWLC Sysplex Transition Charges (TC2)

When two or more machines exist in an aggregated Sysplex or Complex & at z13, zEC12, or zBC12 server & at least one is a z196 or z114 server, with no older technology machines included, they will receive a reduction to AWLC pricing across the aggregated Sysplex or Complex. This reduction provides a portion of the benefit related to the Technology Update Pricing for AWLC (TU1) based upon the proportion of zEC12 or zBC12 server capacity in the Sysplex or Complex:

zxC12 & z13 MSU %	AWLC % Discount
0-20%	0.5%
21-40%	1.5%
41-60%	3.0%
61-80%	4.0%
81-<100%	4.5%

Notes:

1. Add the total MSUs of all the servers in the Sysplex or Complex.

2. Divide sum of z13 & zXC12 MSUs by the total Sysplex/Complex MSUs.

3. Round to nearest full % (I.E. 20.50% ↑ 21.0% or 20.49% ↓ 20%).

*Example: Three Sysplex Coupled Servers (1*z13 & 1*zEC12 & 1*z196) z13 (2964-514) @ 1000 MSU, zEC12 (2827-519) @ 1096 MSU & z196 (2817-615) @ 1084 MSU, totalling 3180 MSU. z13+zXC12=2096 MSU. 2096/3180 MSU=65.91%, rounded ↑ 66.0%. The Parallel Sysplex will receive a 4.0% reduction in aggregated AWLC billing.*

Put simply, TC2 eligibility is configured General Purpose Processor (GPP) MSU; excluding MSU capacity related to specialty engines, unassigned capacity, & temporary capacity such as, but not limited to, "Capacity Back Up", "Capacity for Planned Events" & "On/Off CoD".

AWLC discounts, even when upgrading from z196/zxC12 to z13.

IBM z13: AWLC Sysplex Transition Charges (TC3)

When two or more machines exist in an aggregated Sysplex or Complex & at least one is a z13 server & at least one is a zEC12 or zBC12 server, with no older technology machines included, they will receive a reduction to AWLC pricing across the aggregated Sysplex or Complex. This reduction provides a portion of the benefit related to the IBM z13 TU3 offering, based on the total Full Capacity MSU of all z13, zEC12, & zBC12 Machines in the Sysplex or Complex:

MSU Range	AWLC % Discount
4-45	2.8%
46-315	5.6%
316-1315	6.3%
1316-2676	7.0%
2677-5476	8.4%
5477+	9.8%

*Example: Three Parallel Sysplex Coupled Servers (2 * z13 + 1 * zEC12)*
Machine #1 is a z13 model 515 server with 1056 MSU Machine #2 is a System z13 model 520 server with 1333 MSU & Machine #3 is a zEC12 model 709 with 1350 MSU, a cumulative total of 3739 MSU. The Parallel Sysplex will receive a 8.4% reduction in aggregated AWLC billing.

Put simply, TC3 eligibility is configured General Purpose Processor (GPP) MSU; excluding MSU capacity related to specialty engines, unassigned capacity, & temporary capacity such as, but not limited to, "Capacity Back Up", "Capacity for Planned Events" & "On/Off CoD".

Good AWLC discounts, even when upgrading from zxC12 to z13.

IBM MLC (WLC) Pricing: Structural Enhancements

IBM Collocated Application Pricing (ICAP): Previously, new applications (zNALC) required a separate LPAR to avoid increase in other MLC software charges. ICAP facilitates new eligible applications, be charged as if they are running in a dedicated environment. Technically they are integrated with other (non-eligible) workloads. Software supporting the new application will not impact the charges for other MLC software collocated in the same LPAR. ICAP appears as an evolution of the Mobile Workload Pricing (MWP) for z/OS pricing mechanism. ICAP will use an enhanced MWRT, implemented as a z/OS application. ICAP applies to z13, zXC12, z196/z114 servers. IBM anticipates that ICAP will deliver zNALC type price benefit, discounting ~50% of ICAP eligible software MSU.
Bottom Line: Simplified low cost new workload charging, without the need for a dedicated LPAR.

Country Multiplex Pricing (CMP): A major evolution, essentially eliminating prior Sysplex pricing rules, requiring that systems be interconnected and/or sharing the same data in order to be eligible for aggregation of MLC software pricing charges. A Multiplex is defined as the collection of all z Systems within a country. Therefore, sub-capacity usage will be measured & reported as a single machine, regardless of the connectivity or data sharing configurations. A new sub-capacity reporting tool is being implemented & clients should expect a transition period as the new pricing model is implemented. This should allow flexibility to move & run work anywhere, eradicating multiple workload peaks when workloads move between machines. Ultimately the cost of growth is reduced with one price per product based on MSC capacity growth anywhere in the country.
Bottom Line: Country aggregation of MSU capacity, could deliver good MLC savings.

Simpler & cost efficient MLC pricing for new & existing workloads.

zSeries Workload License Charges Cost Saving Options #1?

2001	*Variable Workload License Charges (VWLC):* Sub-Capacity Workload License Charges (WLC) to grow hardware capacity & manage workload utilization without increasing software charges. *~2-5% cost savings. No cost to implement; commercial & process changes.*	**Application Tuning** **ASM COBOL SQL VSAM**
	Soft-Capping: Manage Defined Capacity (DC) MSU metrics via WLM & PR/SM functions, paying SCRT invoice for WLC software via R4HA or DC MSU metric, whichever is the lowest. *~2-5% cost savings. No cost to implement; commercial & process changes.*	
2004	*System z Application Assist Processor (zAAP):* Offload MSU workload (E.g. Java and related workloads) to specialty engine, avoiding sub-capacity MSU charges (software runs for free). *~5-10% MSU reduction per engine. Cost option, but zAAP specialty engines run at full speed.*	
2006	*System z Integrated Information Processor (zIIP):* Offload MSU workload (E.g. DB2, SRB Enclave, Etc.) to specialty engine, avoiding sub-capacity MSU charges (software runs for free). *~5-10% MSU reduction per engine. Cost option, but zIIP specialty engines run at full speed.*	
	Group Capacity Limit (GCL): Software cost control, while taking advantage of available server MSU capacity (white-space), for a pre-defined LPAR groups aggregate Defined Capacity. *~0-5% cost savings. No cost to implement; looping transaction can impact CPC performance.*	
	Advanced Soft-Capping: Utilize a 3rd party software tool to dynamically optimize available MSU capacity between CPC, LPAR & CPU resources, prioritizing by workload classification. *~5-15% cost savings. Cost option, with short-term ROI. Transparent implementation.*	

A myriad of mostly low cost options to optimize WLC TCO…

zSeries Workload License Charges Cost Saving Options #2?

2007	*System z New Application License Charges (zNALC):* Separate LPAR running a qualified new workload (E.g. WebSphere Application Server for z/OS, Domino, SAP, PeopleSoft, Siebel, Etc.). *~90% z/OS product MLC discount. Software workload cost; commercial & process changes.*	**Application Tuning** ASM COBOL SQL VSAM
2010	*Advanced Workload License Charges (AWLC):* A Sub-Capacity eligible pricing metric, delivering additional price performance for z196 & upwards servers, superseding VWLC. *~2-10% cost savings. No cost to implement; commercial & process changes.*	
2013	*Capacity Provisioning Manager (CPM):* Dynamically manage Defined Capacity (DC) & Group Capacity Limit (GCL) MSU settings, as per soft-capping, workload (WLM) & provisioning rules. *~0-5% cost savings. No cost to implement; process changes.*	
2014	*Mobile Workload Pricing for z/OS (MWP):* For mobile transaction workload eligible products (I.E. CICS, DB2, IMS, WebSphere/MQ); subtract 60% of MSU from GPP MSU (R4HA) total. *~0-15% cost savings. No cost to implement; commercial & process changes.*	
2015	*IBM Collocated Application Pricing (ICAP):* New eligible (TBD – CICS, DB2, IMS, WebSphere/MQ) workloads in existing LPAR, without impacting existing LPAR MLC charges. *~0-15% cost savings. No cost to implement; commercial & process changes.*	
	Country Multiplex Pricing (CMP): Aggregate country wide MSU metrics for all workloads, not just coupled Sysplex/Multiplex workloads, potentially avoiding concurrent R4HA peaks. *~0-5% cost savings. No cost to implement; commercial & process changes.*	

A myriad of mostly low cost options to optimize WLC TCO…

z13 Chip Design & Software Pricing: Useful Web Links

- ✓ *IBM z13 Hardware Design – Feeds & Speeds:*
http://mainframeinsights.com/ibm-z13-hardware-value-market-differentiation/
- ✓ *IBM z13 Software Pricing Technology Transition Offerings:*
http://www-01.ibm.com/common/ssi/rep_ca/1/877/ENUSZP15-0001/ENUSZP15-0001.PDF
- ✓ *IBM z13 Announcement Review – Clipper Group:*
http://www.clipper.com/research/TCG2015001.pdf
- ✓ *IBM System z Simultaneous Multithreading Observations:*
http://www.ibmsystemsmag.com/mainframe/trends/IBM-Research/smt_mainframe/

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

Vendor Update: PerfTechPro (PTP) zAnalytics™ Version 2

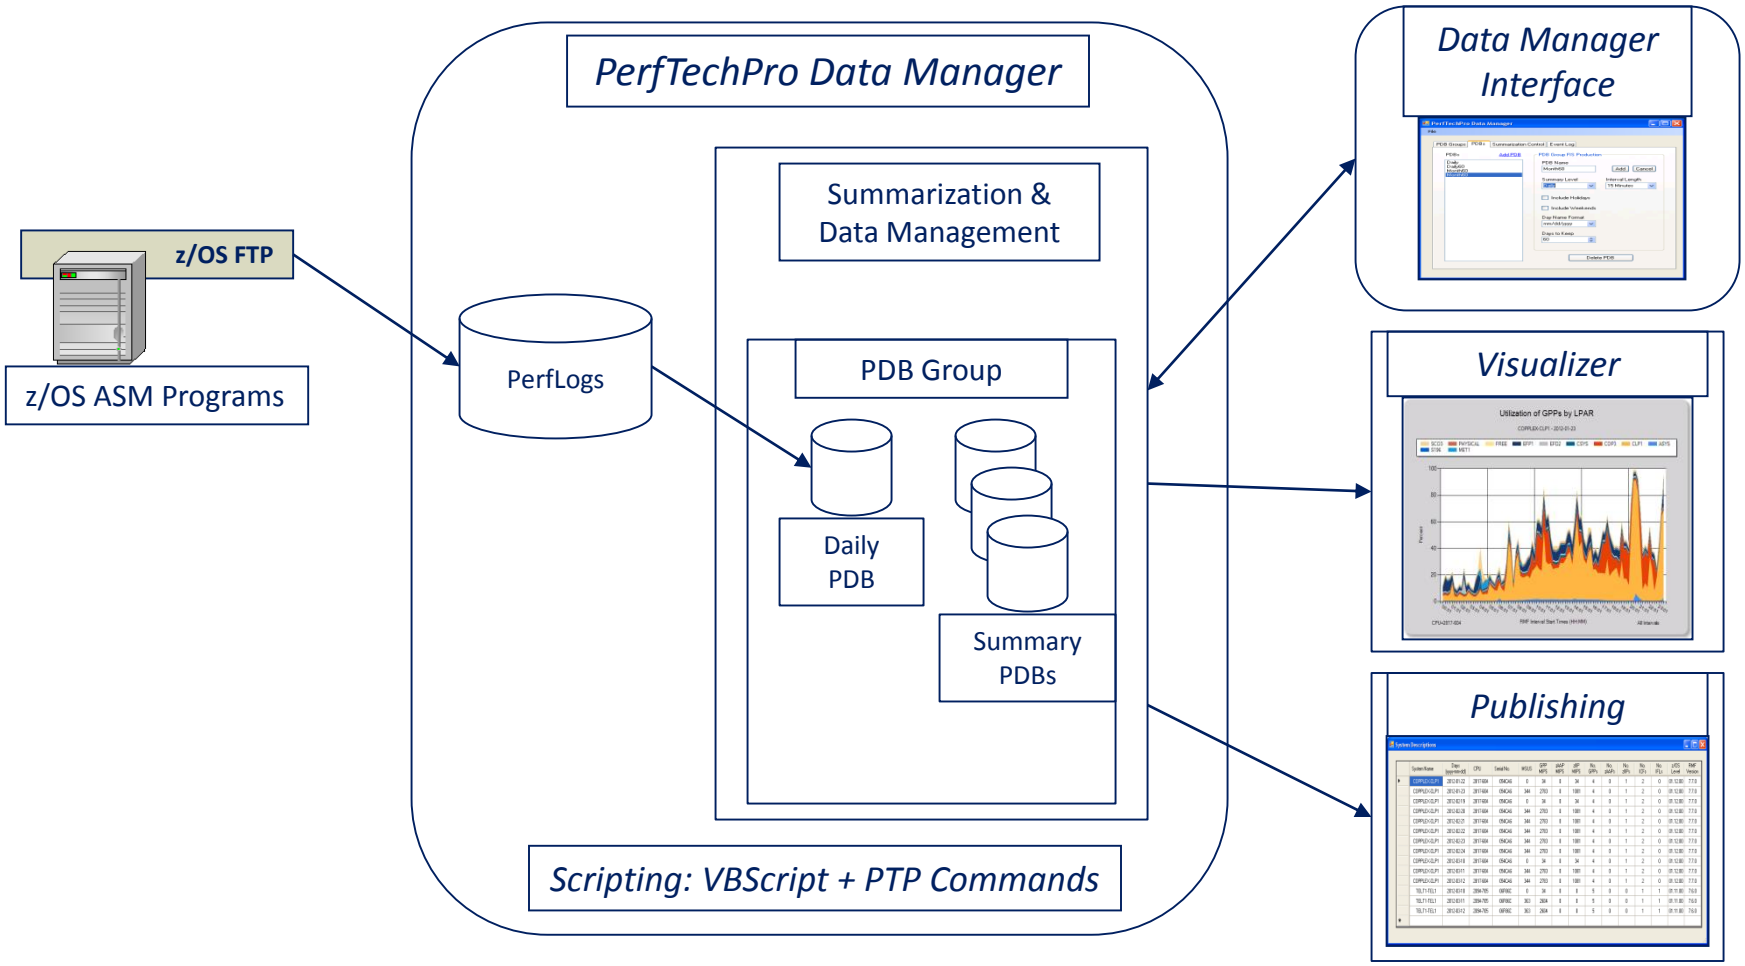

A fully-rounded capacity & performance management solution

Vendor Update: PerfTechPro (PTP) zAnalytics™ Version 2

Current Stand-Alone Architecture

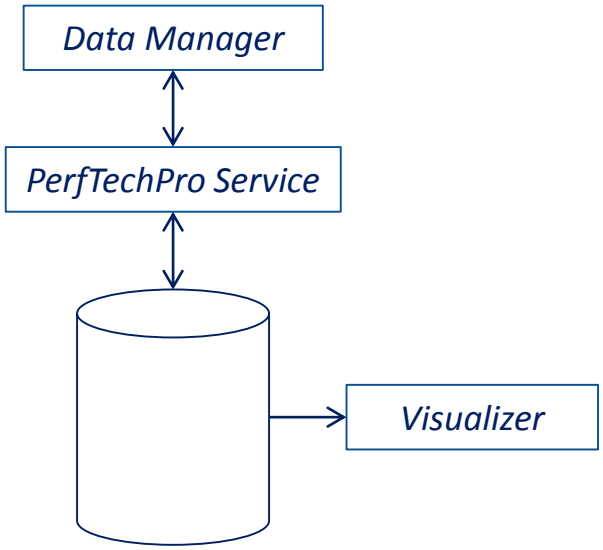

New Client Server Architecture

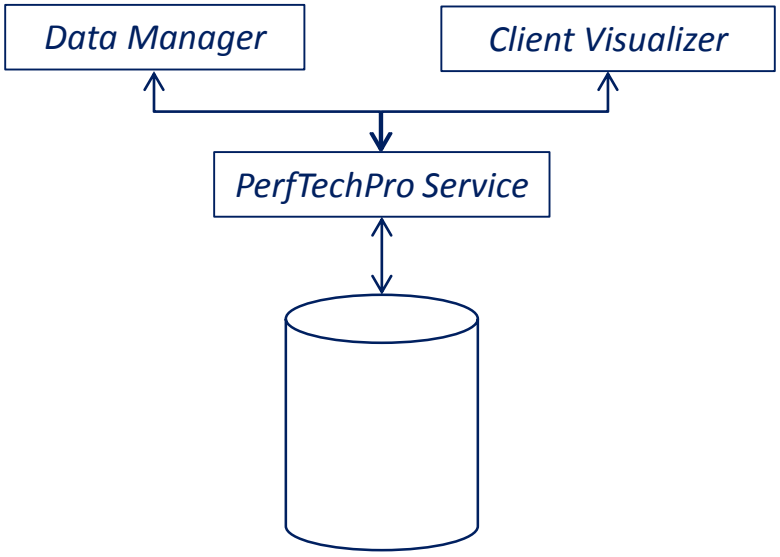

Client Server Architecture Benefits:

- *Network Efficient (E.g. Data Transfer)*
- *Web Service Based (E.g. Browser Access)*
- *Remote Support (E.g. Out Of Corporate Office Access)*
- *Multiple Database Support (E.g. SQL Server, MySQL, Oracle)*

PTP: simple & efficient usage, corporate network attached or not

Vendor Update: PerfTechPro (PTP) zAnalytics™ Version 2

Model Parameterization

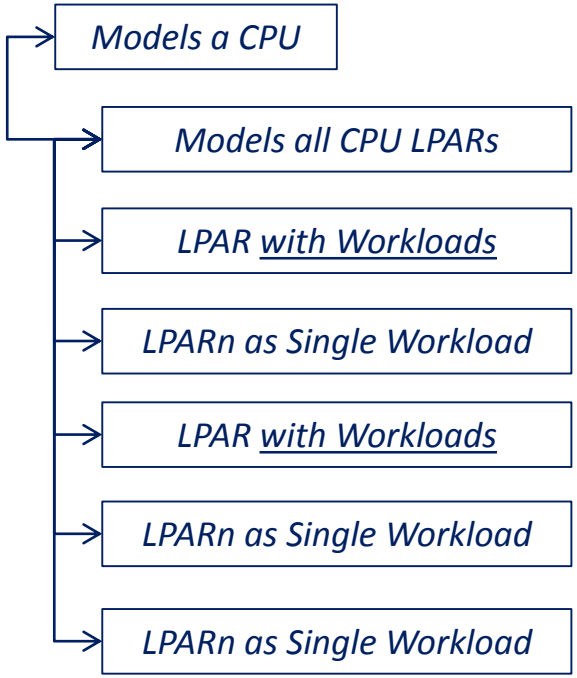

Model Execution

- ❖ ***Simulation Mode***
- ✓ Models contention for LPAR assigned GPP, zAAP & zIIP Logical Processors (LPs)
- ✓ Models contention for PR/SM managed physical GPPs, zAAPs & zIIPs
- ✓ Models DASD I/O (E.g. Rates, Contention, et al)
- ❖ ***Model Changes***
- ✓ LPAR controls (E.g. Weight, Capping, et al)
- ✓ Workload characteristics (E.g. CPU Utilization, DASD I/O Rate, Memory Usage, WLM, et al)
- ✓ CPU Server Changes (E.g. #GPPs, #zAAPs, #zIIPs, Model/Series CPU Engine Changes, et al)
- ✓ zPCR Interface (E.g. MIPS Ratings, Official IBM Process, et al)
- ❖ ***Model Outputs***
- ✓ Compare measured system, base & new models
- ✓ Save models (I.E. compare predicted vs. actual performance)

PTP: complete end-to-end zSeries Server coverage & zPCR input

Vendor Update: zIT Consulting, zDynaCap/zPrice Manager

- ❖ *zDynaCap: Plug & Play Dynamic WLM Based MSU Optimization (CPC, LPAR, LPAR Group)*
- ✓ *Enhanced & improved MSU optimization algorithms (E.g. Optimized MSU for R4HA & performance benefits)*
- ✓ *Enhanced real-time 5 Minute interval GUI reporting (I.E. As per WLM sampling; better decision making)*

- ❖ *zPrice Manager: Granular Full MSU & WLC Optimization (CPC, LPAR, LPAR Group, WLM, WLC Product)*
- ✓ *Enhanced & improved MSU optimization algorithms (E.g. Optimized MSU for R4HA & performance benefits)*
- ✓ *Enhanced real-time 5 Minute interval GUI reporting (I.E. As per WLM sampling; better decision making)*
- ✓ *Automatic LPAR detection of defined, but inactive LPARs*
- ✓ *Storage of LPAR information resource data in REXX variables (I.E. Callable in REXX procedures for granular management of WLC workloads)*
- ✓ *Optimized IP network management (I.E. zSeries Server & HMC SE communications management)*
- ✓ *Additional workload management rule types (E.g. Greater flexibility for customer workload management)*
- ✓ *Country Multiplex Pricing (CMP) pricing support (E.g. Full WLC/R4HA management)*
- ✓ *Full Mobile Workload Pricing (MWP) & IBM Collocated Application Pricing (ICAP) support, for workloads running in separate LPAR structures*
- ✓ *Full Mobile Workload Pricing (MWP) & IBM Collocated Application Pricing (ICAP) support, for workloads running in the same collocated LPAR structure (Q2 2015)*

Intelligent, granular & fully Optimized WLC MSU management

Backup Slide: MLC Bands & WLM MSU Utilization Metric

EWLC (z800*, z890*, z9 BC, z 10 BC)		AEWLC (z114, zBC12, z13)		VWLC (z900, z990, z9 EC, z 10 EC)		AWLC (z196, zEC12, z13)	
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, & 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.