



How Do You Measure that New z13?

**National CMG Meeting
November 8, 2016**

Kathy Walsh
IBM Distinguished Engineer
IBM

Agenda



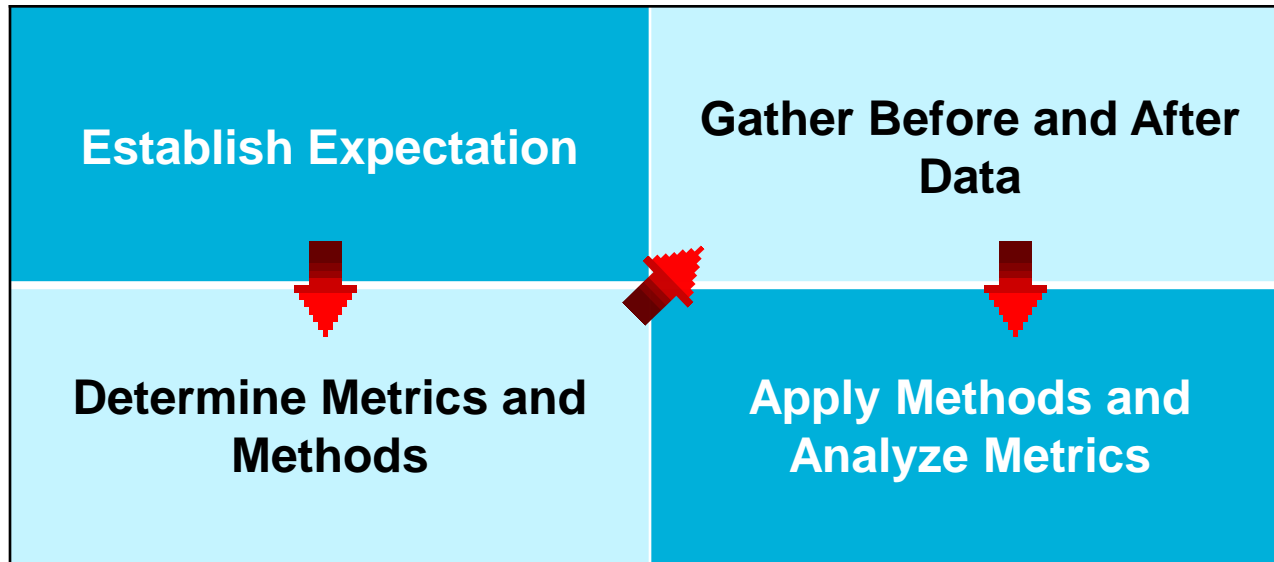
- Measurement challenges
- Establishing expectations
- Understanding metrics
- Describing Methods
- Examples

Overview



- Evaluate performance
 - Did I get what I paid for?
 - How is my workload taking advantage of the z13
- Evaluation Process
 - A series of processes
 - Successively more detailed
 - Using common metrics from production workloads
 - SMF Based Analysis





— Ground Rules

- No synthetic benchmarks
- Expected capacity is applied to workloads not specific jobs or transactions
- We expect winners and losers but on average expect workloads to match expectations

Synthetic Benchmarks



“I have run this job for years....”

Family	Processor	Frequency		LSPR Average Workload		Cache
		GHz	Delta	MIPS	Delta	
z9	710	1.7	N/A	4757	N/A	L1 256k i, 256k d L2 40 MB/Book
z10	710	4.4	158%	7105	49%	L1 64k i, 64k d L1.5 3 MB L2 48 MB/Book
z196	710	5.2	18%	9788	38%	L1 64k i, 64k d L2 1.5 MB L3 24 MB/chip L4 192 MB/Book
zEC12	710	5.5	6%	12179	24%	L1 64k I, 64k d L2 1 MB I, 1 MB d L3 48 MB/chip L4 384 MB/Book
z13	710	5.0	-9%	13515	11%	L1 96k I, 128k d L2 2 MB I, 2 MB d L3 64 MB/chip L4 480 MB/Node (2) 224 MB NIC Directory

Synthetic Benchmarks



- Most Common Types
 - Set of jobs run standalone
 - Loop of common instructions
- Issues
 - Environment doesn't match production
 - Sets of jobs outside of production don't exercise the hardware cache
 - Loop benchmarks essentially mimic the frequency of the processor
- None of these approaches exercise the more powerful aspects of server design
 - Pipelining, overlap execution, enhanced out of order execution, branch history tables, storage hierarch improvements, new instructions, etc.

That leaves us having to measure work running in a production environment

Establish Expectations



- LPARs need to align to 1 of 3 LSPR workload categories
 - LOW, AVERAGE, HIGH
- Determination of workload category is done with CPU MF HIS Data (SMF 113)
 - Focus on L1MP and RNI metrics
- LPAR configuration and workload category is paired with SMF data via the zPCR tool
- zPCR is used to set expectation:
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS1381>
 - Expectation is set with a +/- 5% Confidence Factor
- Data used for this effort is from close to time of migration
 - Should try and align business cycles (day of week, month end, quarter end)
 - Need to determine the workloads are “essentially” the same

Establish Expectation



- Run zPCR for the before and after environments to determine current capacity relationships
- Hold the workload choice constant (Low, Average, High)
 - Held constant to the Before value
 - In the After zPCR study may need to change to the Before value
 - This manipulation is valid for the expectation setting phase only
 - Capacity planning going forward should use workload as defined by SMF 113 for current environment

Introduction to Metrics - ETR



- External Throughput Rate (not typically used for this)
 - The rate at which a processing system executes a particular workload
 - Expressed as: Units of Work / Elapsed Time
 - Characterizes "system" capacity
 - All resources are potential inhibitors
 - CPs
 - Processor Memory
 - Channels, CUs, and I/O Devices
 - Enqueues
 - CFs

Introduction to Metrics - ITR



- Internal Throughput Rate (typically used for this)
 - The rate at which a processor executes a particular workload
 - Expressed as:
 - Unit of Work / CPU Busy Time
 - ETR / Utilization
 - Bears a direct relationship to work done
 - Can be measured with software
 - Characterizes "processor" capacity
 - ITRs are used by LSPR to measure IBM processors

Introduction to Metrics - COV



- In probability theory and statistics, the **coefficient of variation** (COV), also known as relative standard deviation (RSD), is a standardized measure of dispersion of a probability distribution or frequency distribution
 - Defined as the ratio of the [standard deviation](#) to the [mean](#)
 - It shows the extent of variability in relation to the mean of the population
- Use it in detailed analysis phases to understand if work is repeatable and stable
 - Look for a coefficient of variance of .104 - .300 though higher variances are sometimes needed to include more of a workloads transactions

ITRs and MIPS



- MIPS ratings are meaningless
- Comparison of MIPS ratings is an ITRR

– ITR Ratio

- Example:

Family	Processor	LSPR Average Workload		
		MIPS	Delta	ITRR
z9	710	4757	N/A	N/A
z10	710	7105	49%	1.49
z196	710	9788	38%	1.38
zEC12	710	12179	24%	1.24
z13	710	13515	11%	1.11



Checking

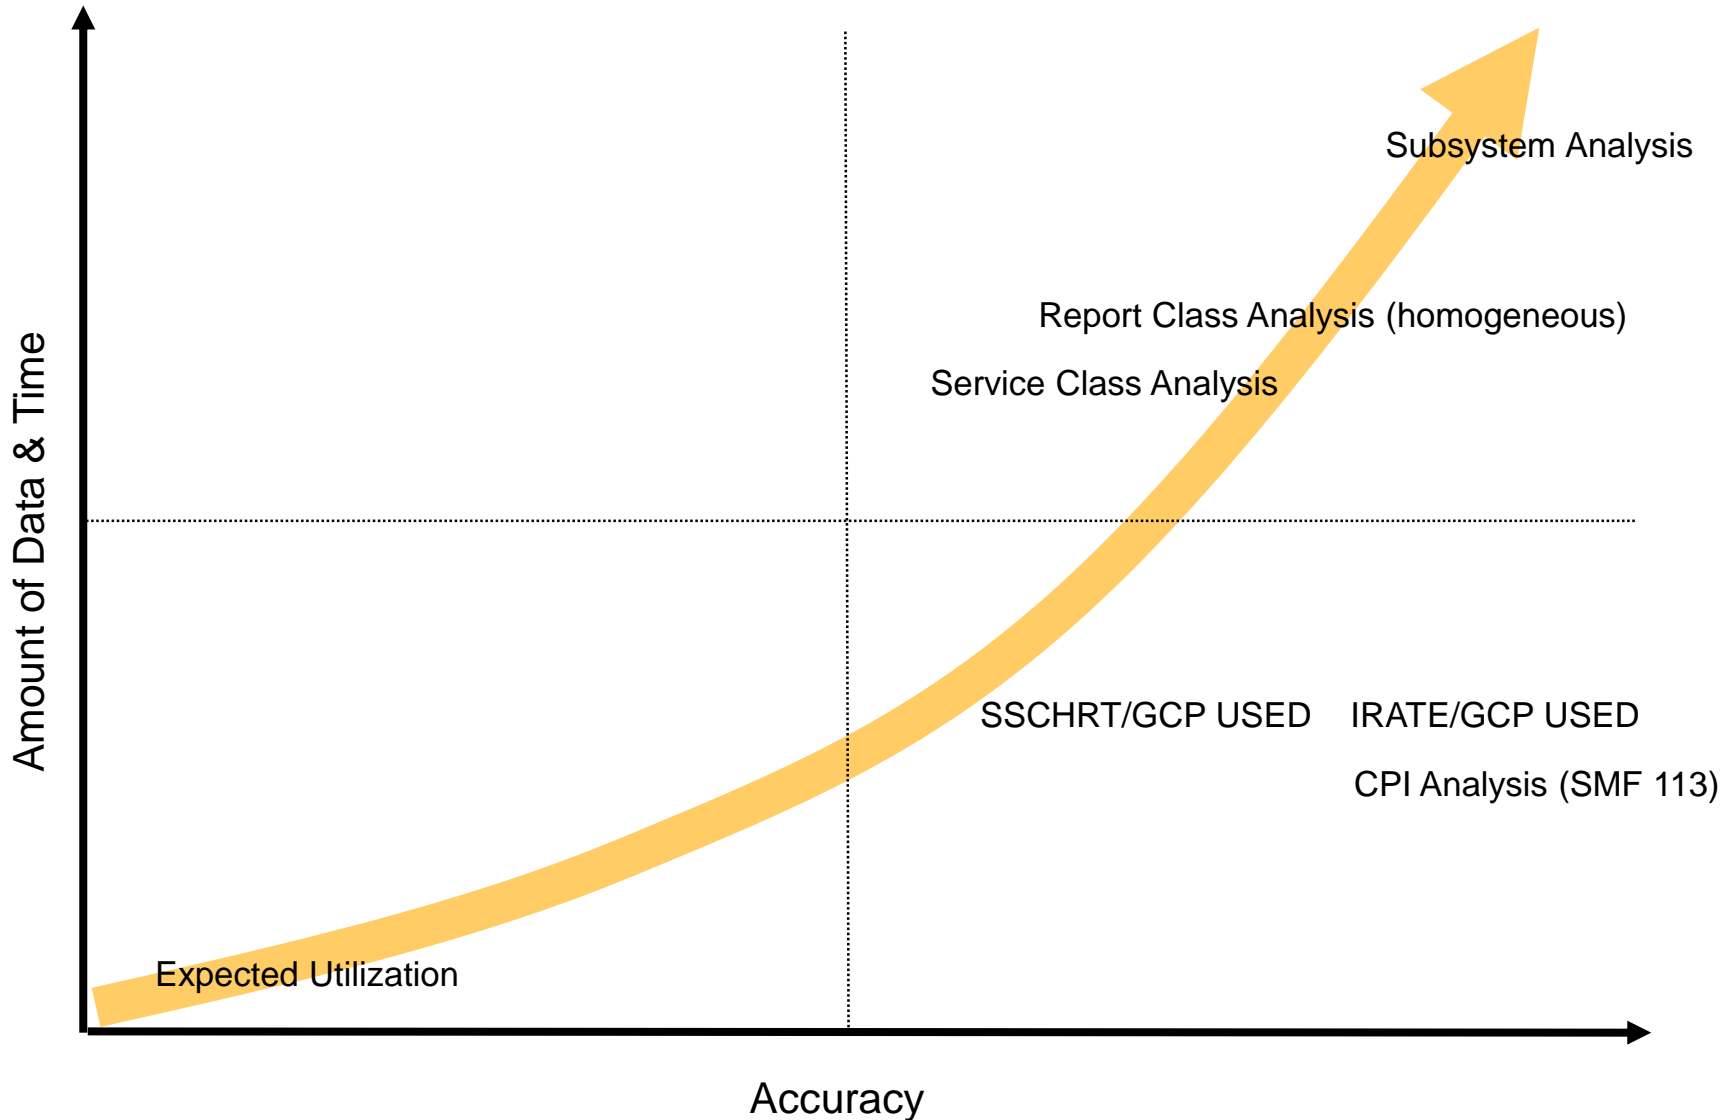
Overall CEC ITRR and Per CP ITRR



- Individual units of work are judged against the per CP ITRR
- Example:

Family	Processor	LSPR Average Workload				
		MIPS	Delta	ITRR	MIPS / CP	ITRR
z9	710	4757	N/A	N/A	476	N/A
z10	710	7105	49%	1.49	711	1.49
z196	710	9788	38%	1.38	979	1.38
zEC12	710	12179	24%	1.24	1218	1.24
z13	710	13515	11%	1.11	1352	1.11
z13	716	19666	62%	1.62	1229	1.01

Determine Metrics and Methods





■ High Level Metrics

Metric	Definition	Source	Level	Average ITRR	Median ITRR	Weighted
CPI	Cycles Per Instruction (GCP + zIIP)	SMF 113	LPAR	✓	✓	After LPAR CPU
IRATE / GCP	LPAR IO Interrupt RATE / LPAR GCP *	SMF 70 SMF 113	LPAR	✓	✓	After LPAR CPU
SSCH / GCP	DASD Start Subchannels / GCP	SMF72 SMF 113	LPAR	✓	✓	After LPAR CPU
LPAR Utilization	Utilization of the LPAR	SMF 70	LPAR			

* IRATE includes CTC and PCIE interrupts



■ Medium Level Metrics (Service Class or Report Class)

Metric	Definition	Source	Level	Average ITRR	Median ITRR	Weighted
CPU	Amount of zIIP and GCP CPU capacity used	SMF 72	Service or Report Class			
ENDED	Count of Ended Transactions	SMF 72	Service or Report Class			
SSCH	Count of DASD Start Subchannels	SMF 72	Service or Report Class			
CPU PER TRAN	CPU / ENDED (1/ITR)	Calculation	Service or Report Class	✓	✓	After CPU
SSCH PER GCP second	SSCH / CPU (ITR)	Calculation	Service or Report Class	✓	✓	After CPU

Low Level Metrics



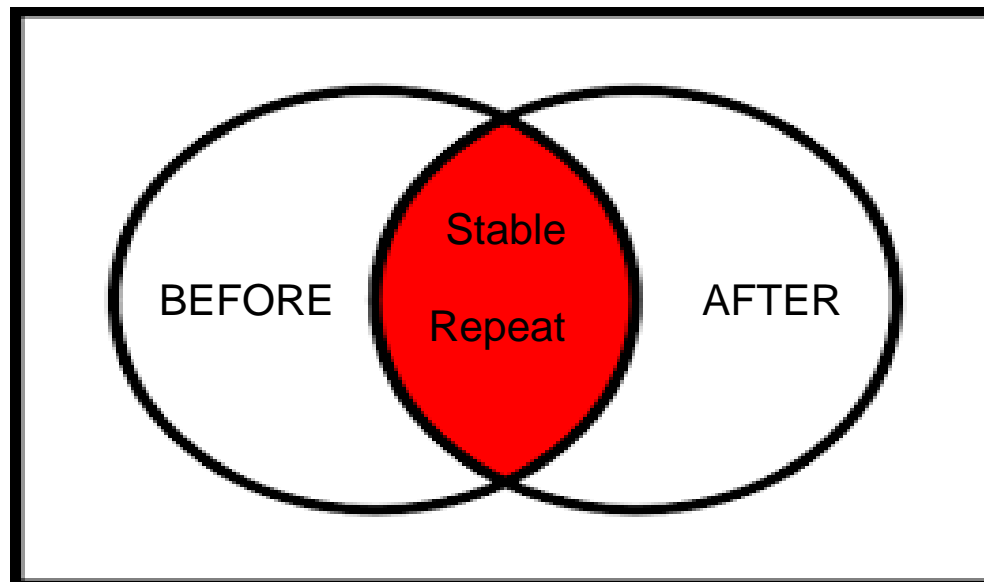
- Low Level Metrics (SMF 30, 110, 101:102, 120, IMS LOGs)
- Usually transaction level metrics

Metric	Definition	Source	Level	Average ITRR	COV	Weighted
CPU PER TRAN	CPU Per transaction (1/ITR)	SMF 110 IMS LOG	Address Space	✓	✓	After CPU
CPU	CPU by Address Space	SMF 30	Address Space			
EXCP	Count of EXCPs	SMF 30	Address Space			
CPU PER EXCP	CPU / EXCPS (1/ITR)	Calculation	Address Space	✓	✓	After CPU
CPU PER SQL CALL	CPU Per SQL Call (1/ITR)	SMF 101	Connection ID	✓		After CPU

Evaluation



- Dealing with Production Systems
 - Influenced by load factors and business cycle
 - Impacted by software and infrastructure changes
- Goal is to find repeatable, stable workloads on which to base the analysis
 - In both the before and after environments



Using Metrics and Methods



- Analysis is an iterative process
 - Skill level requirements get higher the deeper you go
 - More knowledge of design and influence on metrics is needed
 - Leverage knowledge gained in capacity planning for these workloads
- Need a transparent method to deal with outliers
 - Identify factors and describe how to handle
 - Needed for either the Before or After Data
- Objective is to describe workloads in terms of ITRRs
 - ITRRs should have statistical tests applied against them depending upon the metric
 - Average, Median, Coefficient of Variation
 - ITRR are weighted by CPU

Before and After Data



- Before and After data needs to be as close as possible to ensure consistency and same workload mix
- Quantity of data required varies
 - High Level Metrics - 3-5 full days before and after
 - Medium Level Metrics - 3-10 days before and after
 - Low Level Metrics - depends on the transaction rate
 - Subsystems like CICS, IMS, WAS, and DB2 usually 1 hour : 3 days is usually sufficient
 - Job related data (SMF 30) where 1 job represents 1 transaction often need 7-10 days
- May need to do multiple time frames
 - Example: Online vs Batch

Other Factors to Consider




- Relative Utilization of Environments
 - This can be accounted for when setting expectations
- Technology changes in new processor, including MCL changes, which are effective “out of the box”
 - LSPR workloads may not exploit all of these features or benchmarks not run at current MCL levels hence MIPS ratings don’t include
 - Impacts of CF technology in heavy data sharing environments if CFs are not upgraded with servers
 - Operating system and middleware maintenance changes made to support new processor hardware
 - Application changes made concurrently with processor change
- Need to apply “smell test” to results
 - A workload which sees an improvement of 300x with an expectation of 10% that work may need to be eliminated

Drive to a CEC Level Assessment



Expectation is 1.11



Common Elements	ITRR	Capacity in CPs	Relative Weight	Relative ITRR
A	1.02	0.2	.07	.0714
B	1.06	0.4	.13	.1378
C	1.14	0.6	.2	.228
D	1.12	0.8	.27	.3024
E	1.09	1.0	.33	.3597
Total		3.0	1	1.10

Start with the Building Blocks of an Element



- Let's take IRATE per GCP
- Collect data from Before and After and Summarize

IRATE HOUR/DAY	Before				After		
	8	9	10		29	30	31
8	28884	23971	19399		19873	19643	19535
9	32974	24382	19700		19666	21011	18374
10	29223	24995	16014		18902	18621	21779
11	27755	24224	15478		20825	20419	23909
12	29684	14960	15323		25048	23686	20163
13	28808	14027	14186		17865	16827	17395
14	28035	16030	13685		18235	17954	18645
15	23769	15795	14313		16587	17020	16931
16	18395	18083	11938		14739	15249	14752
mean	27503	19607	15560		19082	18937	19054

- Divide IRATE by the GCP in CPs and develop the ITRR called IRATE LPAR GCP

GCP (CPs) HOUR/DAY	Before				After		
	8	9	10		29	30	31
8	4.99	4.62	4.06		4.31	4.52	4.65
9	6.39	5.08	4.14		4.39	4.84	4.78
10	5.54	5.25	4.69		5.01	4.70	5.33
11	6.14	5.15	4.39		5.17	4.93	5.19
12	5.36	4.26	4.05		4.49	4.33	4.68
13	5.23	4.32	3.90		4.71	4.26	4.85
14	4.82	4.58	3.80		5.19	4.57	4.84
15	4.55	4.23	3.83		4.09	4.47	4.77
16	4.43	4.80	4.10		5.09	4.92	4.86
mean	5.27	4.70	4.11		4.72	4.61	4.88

Start with the Building Blocks of an Element



- Let's evaluate the IRATE per GCP metric

IRATE PER GCP	Before			After		
	DAY 8	DAY 9	DAY 10	DAY 29	DAY 30	DAY 31
8	5783.27	5188.24	4778.57	5650.77	5340.44	5168.20
9	5163.35	4795.40	4756.52	5503.55	5275.78	4789.51
10	5275.46	4762.93	3414.95	4668.44	4924.38	4928.89
11	4521.32	4700.37	3522.12	4902.60	5057.54	5470.90
12	5541.06	3508.74	3786.52	6574.33	6510.04	5264.76
13	5513.48	3248.86	3641.82	4750.13	5005.90	4511.70
14	5821.74	3499.88	3603.71	4378.64	4911.10	4781.72
15	5225.80	3734.51	3738.30	5158.46	4814.87	4490.50
16	4155.28	3766.25	2909.70	3781.26	4013.26	3960.97
mean	5222.31	4133.91	3794.69	5040.91	5094.81	4818.57
median	5275.46	3766.25	3641.82	4902.60	5005.90	4789.51
std	558.64	719.08	608.06	807.69	653.47	460.77
cov	0.1070	0.1739	0.1602	0.1602	0.1283	0.0956

- Let's start summarizing

Start with the Building Blocks of an Element



- For this LPAR copy before and after IRATE per GCP information

	DAY	CPI	Median CPI	DASD GCP ITR	Median DASD GCP ITR	IRATE GCP ITR	Median IRATE GCP ITR
	8					5222.31	5275.46
	9					4133.91	3766.25
	10					3794.69	3641.82
zEC12 Average						4383.64	4227.84
	29					5040.91	4902.60
	30					5094.81	5005.90
	31					4818.57	4789.51
z13 Average						4984.76	4899.34
Relative Capacity Ratio						1.14	1.16

- Go do the rest of the high level metrics

Start with the Building Blocks of an Element



- CPI metric is calculate as Before CPI / After CPI * (After GHz / Before GHz)
- Example for zEC12 to z13: $(3.18/2.64) * (5.0/5.5) = 1.10$

	DAY	CPI	Median CPI	DASD GCP ITR	Median DASD GCP ITR	IRATE GCP ITR	Median IRATE GCP ITR
	8	3.33	3.41	4478	4380	5222	5275
	9	3.13	3.23	3620	3419	4134	3766
	10	3.08	3.14	3291	3158	3795	3642
zEC12 Average		3.18	3.26	3796	3652	4384	4228
	29	2.70	2.67	4790	4425	5041	4903
	30	2.63	2.71	4323	4312	5095	5006
	31	2.58	2.45	3980	4150	4819	4790
z13 Average		2.64	2.61	4364	4296	4985	4899
Relative Capacity Ratio		1.10	1.14	1.15	1.18	1.14	1.16

- Now need to determine a single value for this LPAR

Start with the Building Blocks of an Element



- Take a simple average of the high level metrics
- Record the Amount of CPU used by the LPAR

	Average			Median				
	IRATE GCP ITR	DASD GCP ITR	CPI ITR	IRATE GCP ITR	DASD GCP ITR	CPI ITR	Overall Average	LPAR CPU
WSC1	1.14	1.15	1.10	1.16	1.18	1.14	1.15	474.4
WSC2								
WSC3								
WSC4								
Total / Weighted Average								1341.2

- Complete same analysis for the rest of the LPARs....

Finish the High Level Metric Examination



Expectation is 1.11

	Average			Median			Overall Average	LPAR CPU	Relative Weight	Total ITRR
	IRATE GCP ITR	DASD GCP ITR	CPI ITR	IRATE GCP ITR	DASD GCP ITR	CPI ITR				
WSC1	1.14	1.15	1.10	1.16	1.18	1.14	1.15	474.4	0.354	0.405
WSC2	1.07	1.02	0.97	1.02	1.03	1.06	1.03	41.9	0.031	0.032
WSC3	1.09	1.11	1.08	1.07	1.08	1.07	1.08	301.0	0.224	0.243
WSC4	1.14	1.18	1.17	1.08	1.13	1.13	1.14	523.9	0.391	0.445
Total / Weighted Average								1341.2	1.00	1.12

- Now decide if you're comfortable with the processor performance....
 - If not move to medium level metrics and measure again



ACTUAL

Start a Medium Level Metric Examination



- Create an inventory of the work in the LPAR

TYPE	SCLASS	CPs	CPSEC	ENDED	COMMENTS
ONL_WKL	DDFTRAN	170.7	55312.9	10306703	Transaction based
CIC_WKL	CICPROD	120.8	39129.7	1000	Started Task
BAT_WKL	BATCH	56.2	18196.9	7652	Transaction based
ONL_WKL	ONLPRODL	42.9	13907.1	22	
SYSTEM	SYSSTC	16.9	5464.8	331	
SYSTEM	SYSTEM	15.4	4975.3	0	
ONL_WKL	DDFLOW	14.4	4674.7	749342	Transaction based
STC_WKL	STCMED	13.3	4314.4	180	
BAT_WKL	BATCHHI	13	4219.3	881	
STC_WKL	STCHI	11.8	3831.5	0	
OMVS_WKL	OMVSPROD	4.4	1415.3	39999	Transaction based
STC_WKL	STCLOW	2.8	915.9	0	
TSO_WKL	TSOPROD	2	657.5	73701	Transaction based

- Transaction based workloads are best since the ENDED transaction rate can be used to create the ITRR CPU per Tran
- Understand which workloads are driving the capacity, they will be the focus of the analysis
- Workloads which use a small amount of CPU won't sway the overall answer and can be dropped

Start a Medium Level Metric Examination



- Create an inventory of the work in the LPAR

SCLASS	DAY	CPUPERTXN	CPs
DDFTRAN	7	0.006	200.7
DDFTRAN	8	0.0052	196.8
DDFTRAN	9	0.005	196
mean		0.005	197.833
median		0.005	
std		0.001	
cov		0.098	
DDFTRAN	29	0.0052	170.7
DDFTRAN	30	0.0047	181
DDFTRAN	31	0.0046	227.2
mean		0.005	192.967
median		0.005	
std		0.000	
cov		0.067	
itr		1.12	

- Process major workloads, focusing on those that can generate CPU per Tran numbers
- For other workloads without transactions look at metrics such as DASD SSCT / CPU
 - Only valid for workloads with significant IO loads
- Process the rest of the workloads and summarize....

Start a Medium Level Metric Examination




- Create the data at the LPAR Level

LPAR	WORK	ITR	CPs	Relative Weight	ITRR
WSC1	DDFTRAN	1.12	1.92	0.51	0.577
	CICPROD	1.14	1.21	0.32	0.370
	Batch	1.09	0.6	0.16	0.175
			3.73	1	1.12

- Then build out the other LPARs

LPAR	Overall ITRR	CPs	Relative Weight	Total ITRR
WSC1	1.12	3.73	0.341	0.382
WSC2				
WSC3	1.13	2.89	0.264	0.299
WSC4	1.15	4.31	0.394	0.453
Total / Weighted Average		10.93	1	1.13


Expectation is 1.11


ACTUAL

Summary



- Trying to determine the overall capacity delivered by the CEC
 - Within +/- 5% of the expectation as set by zPCR
- Focus on the workloads / LPARs which drive the capacity
- Need to look across time, but close to the change
 - For Medium and Low Level Metrics let the quantity of data drive you
 - CICS / IMS/ WAS may only need 1 hour before and after
- Look for homogeneous metrics
 - Report classes are better than service classes
 - AORs are better than TORs or FORs
- Important to check the metrics stability before and after
 - If the CPU per tran is widely different on the prior processor it becomes harder to trust the metric
 - If there is a winner or a loser clearly out of bounds then best to exclude it. For instance, expectation is 1.10 and certain work is showing 7.3. This indicates a workload change has occurred



Notices and Disclaimers

Copyright © 2016 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY. IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply."

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law



Notices and Disclaimers Con't.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. IBM EXPRESSLY DISCLAIMS ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, Aspera®, Bluemix, Blueworks Live, CICS, Clearcase, Cognos®, DOORS®, Emptoris®, Enterprise Document Management System™, FASP®, FileNet®, Global Business Services®, Global Technology Services®, IBM ExperienceOne™, IBM SmartCloud®, IBM Social Business®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®, Netcool®, OMEGAMON, OpenPower, PureAnalytics™, PureApplication®, pureCluster™, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, Smarter Commerce®, SoDA, SPSS, Sterling Commerce®, StoredIQ, Tealeaf®, Tivoli®, Trusteer®, Unica®, urban{code}®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.



Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):



Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SEs only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at: www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

IBM z13 versus zEC12 Hardware Comparison




zEC12

– CPU

- 5.5 GHz
- Enhanced Out-Of-Order

– Caches

- L1 private 64k i, 96k d
- L2 private 1 MB i + 1 MB d
- L3 shared 48 MB / chip
- L4 shared 384 MB / book



z13

– CPU

- 5.0 GHz
- Major pipeline enhancements

– Caches

- L1 private 96k i, 128k d
 - L2 private 2 MB i + 2 MB d
 - L3 shared 64 MB / chip
 - L4 shared 480 MB / node
- plus 224 MB L3 NIC Directory

