

## *Understanding Your Rolling 4 Hour Average to Tune Db2 and Lower Mainframe Costs*



**Craig S. Mullins**  
Mullins Consulting, Inc.  
<http://www.MullinsConsulting.com>

# This presentation was prepared by:

**Craig S. Mullins**

**Mullins Consulting, Inc.**

**President & Principal Consultant**

**IBM Champion for Data and AI**

**IBM Gold Consultant**

[mullc@craigsmullins.com](mailto:mullc@craigsmullins.com)

[www.mullinsconsulting.com](http://www.mullinsconsulting.com)



# Agenda

- **Pricing, Reducing Cost, and Tuning DB2 for z/OS**
  - General Terminology
  - Pricing models
  - VWLC
    - SubCapacity Report Tool (SCRT)
    - Understanding **IMSU** and **R4H**
  - But I thought it was usage pricing?
  - What does this mean for DB2?
  - DB2 tuning approaches to reduce the **R4H**...
    - ...and therefore reduce your monthly software bill.



# Terminology

Technical terms & abbreviations used in this presentation.

**MSU: (Million Service Units) : reference unit for Software billing.**

**IMSU (or ACTMSU): instantaneous MSU consumption**

**R4H: Rolling 4 Hours**

**DC: Defined Capacity**

**CPC and CEC**

**CPC: Central Processor Complex**

**CEC: Central Electronic Complex**

**CPC Capacity: z/Server full capacity**

**HMC: Hardware Management Console**



# MIPS vs. MSUs

## MIPS: an acronym for Millions of Instructions Per Second.

- Historically the number of instructions that could be processed in a second of compute time
- Used to measure general computing capacity
  - For mainframes and other computing systems
- Over time, MIPS was no longer an actual number of instructions but a way to measure computing capacity.
  - Still used by outsourcers (and others)
  - No longer really in vogue

WHAT  
IS  
MILLION OF  
INSTRUCTION  
PER  
SECOND

# MIPS vs. MSUs

**MSU:** million service units is a measurement of the amount of processing work that can be performed in an hour.

- One “service unit” originally related to an actual hardware performance measurement, but that is no longer the case.
- A service unit is an imprecise measurement
- 1 MSU is approximately 8.5 MIPS for those who care

**IBM publishes MSU ratings for every mainframe model.**

- <https://tinyurl.com/MSUbyModel>

**More details for those interested:**

- <https://cloudframe.com/mips-versus-msus-whats-the-difference/>

# Sample MSU Chart

Information presented in this table is current as of April 2022								
						z/OS	z/OS	z/OS
IBM System z16						V2R4	V2R4	V2R4
Processor	#CP	PCI**	MSU***			Low*	Average*	High*
3931-401	1	280	35			0.50	0.50	0.48
3931-402	2	538	68			0.98	0.96	0.90
3931-403	3	789	100			1.45	1.41	1.31
3931-404	4	1034	130			1.91	1.85	1.70
3931-405	5	1274	158			2.36	2.28	2.09
3931-406	6	1507	187			2.81	2.69	2.46
3931-407	7	1736	215			3.25	3.10	2.83
3931-408	8	1958	242			3.69	3.50	3.20
3931-409	9	2176	268			4.12	3.89	3.55
3931-410	10	2388	295			4.54	4.27	3.90
3931-411	11	2596	321			4.96	4.64	4.23
3931-412	12	2798	346			5.37	5.00	4.57
3931-413	13	2996	370			5.78	5.35	4.89
3931-414	14	3189	393			6.18	5.70	5.21
3931-415	15	3377	416			6.58	6.03	5.52
3931-416	16	3561	439			6.97	6.36	5.83
3931-417	17	3741	461			7.36	6.68	6.13
3931-418	18	3920	483			7.75	7.00	6.42
3931-419	19	4099	505			8.13	7.32	6.72
3931-420	20	4278	527			8.51	7.64	7.02
3931-421	21	4457	548			8.89	7.96	7.31
3931-422	22	4635	570			9.27	8.28	7.61
3931-423	23	4813	592			9.65	8.60	7.90
3931-424	24	4991	614			10.03	8.92	8.19
3931-425	25	5168	635			10.40	9.23	8.48
3931-426	26	5346	655			10.77	9.55	8.77
3931-427	27	5523	676			11.14	9.87	9.06
3931-428	28	5699	698			11.51	10.18	9.35
3931-429	29	5876	720			11.88	10.50	9.64
3931-430	30	6052	742			12.24	10.81	9.93
* LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category								
see workload description for more information								
** PCI stands for Processor capacity index. The PCI values were calculated by multiplying the								
LSPR average column by a common scaling factor associated with a particular LSPR table.								
Note the values appearing here were generated using zPCR so the full precision of each								
ITRR ratios represented.								
***MSUs are used for software pricing only; they are not a capacity metric								

<https://tinyurl.com/MSUbyModel>

# Licensing vs. Pricing

**It is important to understand the difference between *licensing* and *pricing*.**

- If you execute IBM zSeries software on a CPC, you must have a *license* to do so.
  - A license for an IBM monthly license charge product is specific to the product and a CPC with a particular serial number.
  - The license is sold in terms of **MSUs**.
  - If you are executing a product on a 1500 MSU CPC, you must have a 1500 MSU license, specifying the serial number of that CPC.
- The *price* for a product, that is, how much you pay IBM each month, depends on the pricing *metric* that is used for that product.
  - Examples of pricing metrics include PSLC and VWLC.



# MLC Products

## Monthly License Charge (MLC) Products

- Monthly license charges apply to many IBM software products
- Some of the most common include z/OS, Db2, CICS, IMS, MQSeries, and COBOL
  - There are others.
- Pricing and terms and conditions for MLC products are based on the pricing metric you select.
- Pricing metrics can roughly be grouped into two categories:
  - Full capacity
  - Sub-capacity



# Full vs. Sub Capacity

## Full capacity based pricing metrics

- Under a full capacity-based metric, all software charges are determined by the capacity of the CPC in which the product runs.
- **Parallel Sysplex License Charges (PSLC)** and **zSeries Entry License Charges (zELC)** are examples of full capacity based metrics.

## Sub-capacity capable pricing metrics

- Under a sub-capacity metric, software charges for certain products are based on the utilization capacity of the LPARs in which the product runs.
- **Workload License Charges (WLC)** and **Entry Workload License Charges (EWLC)** are examples of sub-capacity capable pricing metrics.

# IBM Software Pricing Metrics

	z13 z14 z15	zEC12 z196	z10EC z9EC	z990 z900	9672 MP2000	z13s z14 ZR1 z15 T02	zBC12 z114	z10 BC z9BC	z890	z800	MP3000
<b>Advanced Workload License Charges (AWLC) *+ Z125-8538</b>	x	x				x5	x5				
<b>Country Multiplex License Charges (CMLC) * Z126-6965</b>	x	x6				x	x6				
<b>Variable Workload License Charges (VWLC) *+ Z125-6516</b>		x1	x	x			x1	x2	x2	x2	
<b>Flat Workload License Charges (FWLC) ^</b>	x	x	x	x		x5	x5	x2	x2	x2	
<b>Advanced Entry Workload License Charges (AEWLC) * Z125-8755</b>						x	x				
<b>Entry Workload License Charges (EWLC) * Z125-6587</b>								x	x	x	
<b>Tiered Workload License Charges (TWLC) Z125-3901</b>						x	x	x	x		
<b>System z New Application License Charges (zNALC) *+ Z125-7454</b>	x	x	x	x		x	x	x	x	x	
<b>Parallel Sysplex License Charges (PSLC) + Z125-5205</b>	x	x	x	x	x	x3	x3	x3	x3	x3	
<b>Midrange Workload License Charges (MWLC) * Z125-7452</b>	x	x	x					x			
<b>zSeries Entry License Charges (zELC) Z125-3901</b>						x4	x4	x4	x4	x	
<b>Growth Opportunity License Charges (GOLC) Z125-3901</b>											x

- \* Metric eligible for Sub-Capacity charges
- + Metric eligible for Aggregation in a qualified Parallel Sysplex®
- ^ Metric available only in conjunction with AWLC, CMLC and VWLC
- 1 Customer may select VWLC pricing for a zEC12, zBC12, z196 or z114 machine only when it participates in a VWLC Sysplex under the AWLC Transition Charges for Sysplexes terms and conditions with at least one of the following machines: z10EC, z10BC, z9EC, z9BC.
- 2 WLC and FWLC are only available on z10BC, z9BC, z890 or z800 when that machine is participating in a qualified Parallel Sysplex environment.
- 3 PSLC are only available on z15 T02, z14 ZR1, z13s, zBC12, z114, z10BC, z9BC, z890 or z800 when that machine is participating in a qualified Parallel Sysplex environment.
- 4 The z15 T02 model A01, z14 ZR1 model A01, z13s model A01, zBC12 model A01, z114 model A01, z10BC model A01, z9BC model A01, and the z890 model 110 are priced using zSeries Entry License Charges (zELC)
- 5 AWLC and FWLC are only available on a z15 T02, z14 ZR1, z13s, zBC12 or z114 when that machine is participating in a qualified Parallel Sysplex environment.
- 6 Only applicable for existing CMP environments with Machines that were Currently Eligible when added to multiplex.

# Confused? Don't worry...



# Tailored Fit Pricing (TFP)


## One of IBM's latest mainframe software pricing methods

### > **TFP** – Tailored Fit Pricing

- Designed to be more like cloud pricing
- Based on overall hardware consumption, not LPAR R4HA

### > High-level

- A baseline of overall MSU consumption for the year is taken
- Agreement to consume more MSUs in the upcoming years
- Discounted rate for MSUs applied
- Monthly billing is predictable and the same regardless of monthly variations in usage



This is a high-level overview, only...

<https://www.ibm.com/it-infrastructure/z/pricing-tailored-fit>

# Sub-Capacity Pricing Advantages

**Your charges for products that use sub-capacity pricing are based on how much the LPARs in which the products run utilize system resources, rather than on the full capacity of the CPC.**

- You can purchase hardware capacity for future needs without incurring an immediate increase in your software bill.

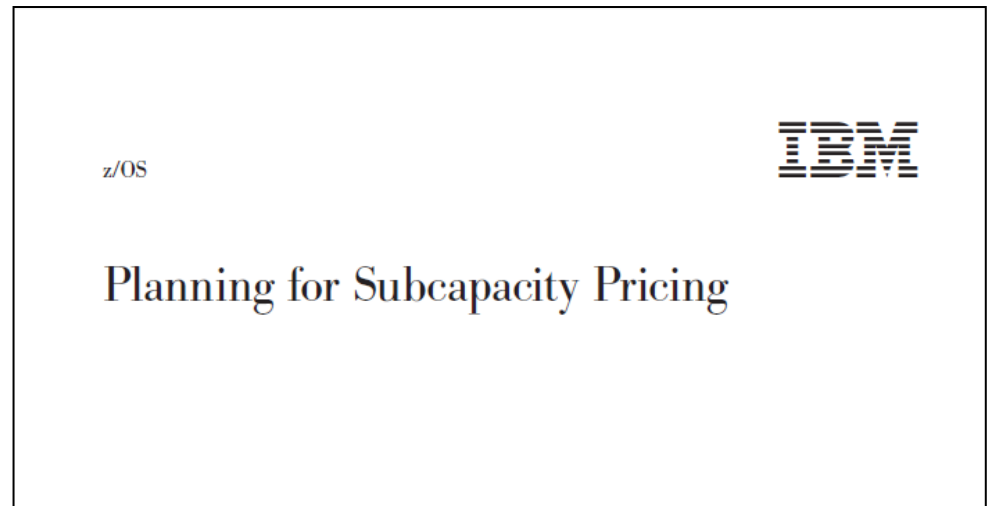
**If your usage decreases when business is slow, your software bill decreases with it.**

- If your usage is seasonal, your monthly software bills are lower during periods of lower usage.

**You pay for capacity on a rolling four-hour (R4HA) average, not on maximum capacity reached.**

# Moving to Sub-Capacity

- If you are not already using subcapacity pricing...
- Read *Planning for Subcapacity Pricing* (SA23-2301-30)
  - **Step by Step Guide**
    - Form a team
    - Prerequisites
    - Timeline
    - Software inventory
    - Run a planning tool
    - Get a cost analysis from IBM
    - Create cost chart
    - Review T&C's



# The 4 parts of Sub-Capacity

R4H, DC, IMSU, SCRT





# Part 1

# R4H

# The Rolling Four Hour Average (R4H or R4HA)

The **R4H** represents **the average consumption** (in **MSU**) of the LPAR during the last 4 hours  
 = Average of the last 48 **IMSU** metrics.




# Clarifying R4H

Every 5 minutes **IMSU** consumption is measured

The **R4H** is an average of the past 48 **IMSU** metrics

> 48 (metrics) / 4 (hours) = 12 metrics/hour



# How Does R4H Translate Into Billing?

This is not *product* usage based pricing; instead, it is a bit of a hybrid.

> Remember we use **IMSU** to calculate the **R4H**.

**So, assume you have three LPARs.**

- > **R4H** averages are calculated for each qualified product, **each hour**, for each LPAR, for the month.
- > A 30 day month has 720 hours → (24 x 30 = 720)

Hour	1	2	3	4	5	6	...	718	719	720
LPAR1	60	55	50	48	44	47	...	49	50	45
LPAR2	70	80	75	60	64	68	...	71	75	70
LPAR3	40	40	42	40	34	38	...	41	47	44

# How Does R4H Translate Into Billing? (continued)

## Now we also have various products in each LPAR

- > LPAR1: z/OS, DB2, MQ
- > LPAR2: z/OS, CICS, DB2, COBOL
- > LPAR3: z/OS, CICS, IMS, COBOL

## The products will be charged based on the peak period of the sum of the R4H for the LPARs in which they run

- > z/OS: LPAR1, LPAR2, LPAR3
- > CICS: LPAR2, LPAR3
- > DB2: LPAR1, LPAR2
- > MQ: LPAR1
- > COBOL: LPAR2, LPAR3
- > IMS: LPAR3

# How Does R4H Translate Into Billing? (continued)

We need R4H sums for all combination with products

Hour	1	2	3	4	5	6	...	718	719	720
LPAR1 MQ	<u>60</u>	55	50	48	44	47	...	49	50	45
→ LPAR2	70	<u>80</u>	75	60	64	68	...	71	75	70
→ LPAR3 IMS	40	40	42	40	34	38	...	41	<u>47</u>	44
...										
1+2+3 z/OS	170	<u>175</u>	167	148	142	153	...	161	172	159
1+2 DB2	130	<u>135</u>	125	108	108	115	...	120	125	115
2+3 CICS, COBOL	110	120	117	100	98	106	...	112	<u>122</u>	114

# How Does R4H Translate Into Billing? (continued)

**Q: So what is the billable MSU for each product?**

**A: It is the peak R4H for the LPAR or grouping of LPARs in which the product ran.**

Hour	1	2	3	4	5	6	...	718	719	720
LPAR1 MQ = 60	<u>60</u>	55	50	48	44	47	...	49	50	45
LPAR3 IMS = 47	40	40	42	40	34	38	...	41	<u>47</u>	44
1+2+3 z/OS = 175	170	<u>175</u>	167	148	142	153	...	161	172	159
1+2 DB2 = 135	130	<u>135</u>	125	108	108	115	...	120	125	115
2+3 CICS, COBOL = 122	110	120	117	100	98	106	...	112	<u>122</u>	114

# Part 2

## DC



# DC vs. WLM vs. PR/SM

## Defined Capacity (DC)

- › Set in the HMC and used to control billing
- › Does not enforce capping and is **not** mandatory

## z/OS Workload Manager (WLM)

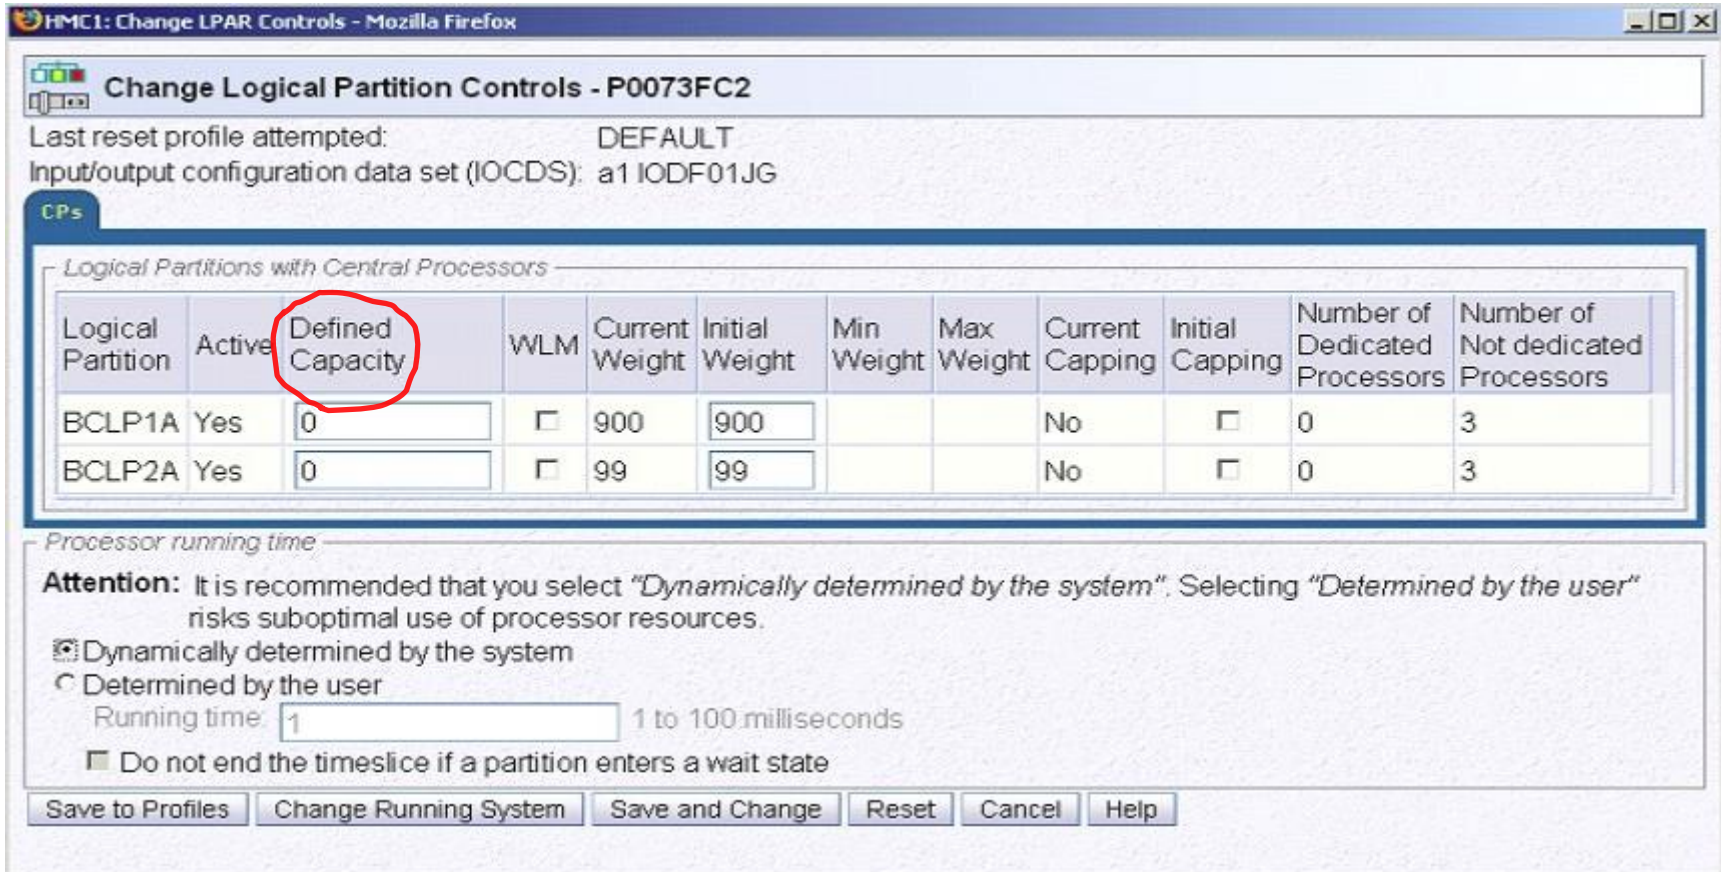
- › Calculates and monitors the R4H
- › If DC is set to a non-zero number, WLM monitors the R4H and ensures that the R4H is less than or equal to the DC
- › Of course, WLM also manages your workloads based on settings/goals (WLM operations are not within scope of this presentation)

## Processor Resource/Systems Manager (PR/SM)

- › Enforces the soft cap when WLM determines it is needed

# Defined Capacity

- Hardware Management Console (HMC) – you set the DC here



HMC1: Change LPAR Controls - Mozilla Firefox

Change Logical Partition Controls - P0073FC2

Last reset profile attempted: DEFAULT  
Input/output configuration data set (IOCDS): a1 IODF01JG

CPs

Logical Partitions with Central Processors

Logical Partition	Active	Defined Capacity	WLM	Current Weight	Initial Weight	Min Weight	Max Weight	Current Capping	Initial Capping	Number of Dedicated Processors	Number of Not dedicated Processors
BCLP1A	Yes	0	<input type="checkbox"/>	900	900			No	<input type="checkbox"/>	0	3
BCLP2A	Yes	0	<input type="checkbox"/>	99	99			No	<input type="checkbox"/>	0	3

Processor running time

**Attention:** It is recommended that you select "Dynamically determined by the system". Selecting "Determined by the user" risks suboptimal use of processor resources.


Dynamically determined by the system  
 Determined by the user  
Running time:  1 to 100 milliseconds  
 Do not end the timeslice if a partition enters a wait state

Save to Profiles Change Running System Save and Change Reset Cancel Help

# Part 3

# IMSU

► **IMSU** : Instantaneous consumption of MSU for the LPAR



# Part 4

# SCRT

# What is the SCRT?

- The Sub-Capacity Reporting Tool (SCRT) is a no-charge IBM tool that reports required license capacity for sub-capacity eligible products.
- The SCRT indicates the required license capacity (in MSUs) of each sub-capacity eligible product.
  - The SCRT cross-references LPAR utilization and product execution by LPAR to determine the maximum concurrent LPAR four-hour rolling average utilization — the highest combined utilization of LPARs where each product executes during the reporting period.
  - Sub-capacity products are charged based on the rolling four-hour (R4H) average utilization of the LPARs in which the sub-capacity products execute.
  - The sub-capacity report determines the required license capacity by examining, for each hour in the reporting period:
    - The four-hour rolling average utilization, by LPAR
    - Which eligible products were active in each LPAR

# Sections of the SCRT Report

The SCRT report is divided into sections. Each section contains specific type of information for a specific purpose:

B5 – Customer Information

C5 – Tool Information

D5 – Special Conditions

E5 – Product Summary Information

H5 – Detail Data Collection Information

L5 / M5 – Customer Verification Information

N5 – Detail LPAR Data

P5 – Product Detail

Q5 – Product Grid Snapshot

S5 – Defined Capacity Value Used

U5 – Group Capacity LPARs



# Section B5

## Customer Information

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 05 Mar 2010 - 10:55  
Name of Person Submitting Report: James User  
E-Mail Address of Report Submitter: James.User@softwareonz.com  
Phone Number of Report Submitter: 281 808-0270  
Customer Name SOFTWAREONZ  
Customer Number 216  
Machine Serial Number 99-026EE  
Machine Type and Model 2094-706  
Machine Rated Capacity (MSUs) 422  
Purchase Order Number (optional)  
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

Note: This report is expected to provide a "% data collected" > 95% and data reporting period beginning on the 2nd of the previous month and ending on the 1st of the current month.



# Customer Comments

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 05 Mar 2010 - 10:55  
Name of Person Submitting Report: James User  
E-Mail Address of Report Submitter: James.User@softwareonz.com  
Phone Number of Report Submitter: 281 808-0270  
Customer Name SOFTWAREONZ

Customer Number 216  
Machine Serial Number 99-026EE  
Machine Type and Model 2094-706  
Machine Rated Capacity (MSUs) 422  
Purchase Order Number (optional)  
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

Note: This report is expected to provide a "% data collected" > 95% and data reporting period beginning on the 2nd of the previous month and ending on the 1st of the current month.

The Customer Comments field can be used to describe anomalies that occur. For example, if you had a bug that caused a looping problem and consumed an extraordinary amount of workload, you can report it to IBM here to lobby for a decrease in your monthly bill.

# Section C5

## Tool Information

==C5=====	
TOOL INFORMATION	
Tool Release	18.2.0
Reporting Period	2 Oct, 2008 - 1 Nov, 2008 inclusive (31 days)
	Justification for low data collection (255 chars max)
% Data Collected for z/OS	93% (required)

### Section C5 is used for input & compliance purposes

SCRT requires System Management Facilities (SMF) data, specifically:

- SMF type 70, subtype 1 (CPU Activity) records and
- SMF type 89, subtypes 1 and 2 (Product Use) records

# Percent Data Collected

The Percent Data Collected, a field in the sub-capacity report, is calculated as the actual hours of data divided by expected hours of data. The number of expected hours varies depending on whether a month has 28, 30 or 31 days.

If your percentage collected is less than 95%, you must indicate why this is the case in the field labeled **Justification for Low Data Collection (up to 255 characters)**:

- Indicates where you must explain why the data collection for the matching operating system was less than 95%.
- You must either provide your explanation by editing the sub-capacity report using a spreadsheet application or you must provide your explanation through the license management support (LMS) application.
- The placeholder will appear in this column whenever you must explain why the data collection for the matching operating system was less than 95%.
- Your explanation must replace the placeholder.
- No explanation is required when the placeholder does not appear.

# Section P5

## Product Detail

```
==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name                Product ID    Highest    Date/Time                LPARD    LPARP    LPARWP
z/OS V1                     5694-A01     322       27 Feb 2010 - 06:00     24       294     4
DB2 UDB for z/OS V8         5625-DB2     322       27 Feb 2010 - 06:00     24       294     4
CICS TS for z/OS V3         5655-M15     318       27 Feb 2010 - 06:00     24       294     0
MQSeries for z/OS V6        5655-L82     322       27 Feb 2010 - 06:00     24       294     4
IMS V10                     5635-A01     318       27 Feb 2010 - 06:00     24       294     0
IBM Enterprise Cobol for z/OS and OS/390 V3 5655-G53     101       15 Feb 2010 - 10:00     101      0       0
Fault Analyzer for z/OS V8   5655-S15     296       16 Feb 2010 - 12:00     96       200     0
Fault Analyzer for z/OS V7   5655-R46     101       15 Feb 2010 - 10:00     101      0       0
Debug Tool Utilities Functions for z/OS V8 5655-S16     101       15 Feb 2010 - 10:00     101      0       0
Debug Tool Utilities Functions for z/OS V7 5655-R45     101       15 Feb 2010 - 10:00     101      0       0
IBM WebSphere Application Server for z/OS V6 5655-N01     24        05 Feb 2010 - 21:00     24       0       0
IBM WebSphere Application Server for z/OS V5 5655-I35     106       15 Feb 2010 - 10:00     101      0       5
Application Perf Analyzer for z/OS V8        5697-N63     101       15 Feb 2010 - 10:00     101      0       0
Application Perf Analyzer for z/OS V7        5697-N53     322       27 Feb 2010 - 06:00     24       294     4
```

# R4H or DC: What is Reported?

Each hour, the (R4H) rolling four-hour average is compared to the defined capacity (DC), if set.

If the DC is set, the SCRT uses the lower of the two values for the utilization value for the z/OS system for that hour.

If  $R4H < DC$  then R4H is used

If  $R4H > DC$  then DC is used

If the DC is set to zero (that is, no soft capping), the SCRT uses the R4H as the utilization value for the z/OS system for that hour.



Remember, the R4H is the average value of SMF70LAC in that hour.

# IBM Sub Capacity Reporting Tool (SCRT)

- SCRT (Sub Capacity Reporting Tool) is a tool used by IBM for z/OS billing.
- It uses the following SMF records:
  - SMF 70-1 and SMF 89-1 / 89-2
- The SCRT is computed on a monthly basis :
  - During the month (from the second day of the month at 0h00 to the first day at 24h00 of the following month)
- R4H averages are calculated for each MLC product, **each hour**, for each LPAR, for the month.

Hour	1	2	3	...	719	720
LPAR1 four-hour rolling avg util	60	55	50		50	45
LPAR2 four-hour rolling avg util	70	80	75		75	70
z/OS (in 1 and 2)	130	135	125		125	115
MQSeries and IMS (in LPAR 1 only)	60	55	50		50	45
CICS and DB2 (in LPAR 2 only)	70	80	75		75	70

# What About Those SMF Records?


## SMF 70-1

- Reports CPU activity for CPU / LPAR

## SMF 89-1 / 89-2

- Report software product usage and are used to calculate reduced sub-capacity software pricing
- Not all products cut 89 records
  - NO89 record products must be specified to the SCRT
- If 89 records are not cut by the product (e.g. COBOL) then it will be charged for the absolute peak **R4H** (or **DC**) for the month
- If 89 records a cut by the product, then it will be charged only for the time it was operational

# The Soft Capping Rule




When **Rolling 4 Hours (R4H)** becomes greater than or equal to **DC** then the LPAR is capped. That means that the **MSU** consumption will not be able to exceed **DC** anymore until the **R4H** becomes lower than the **DC**.




# Subcapacity with Soft-Capping: LPAR view

► **MSU** : Instantaneous consumption of MSU for the LPAR




# Supcapacity with Soft-Capping: LPAR view

- ▶ **IMSU** : Instantaneous consumption of MSU for the LPAR
- ▶ **R4H** : Average of IMSU in 4 consecutive hours.



# Subcapacity with Soft-Capping: LPAR view

- ▶ **IMSU** : Instantaneous consumption of MSU for the LPAR
- ▶ **R4H** : Average of IMSU in 4 consecutive hours.
- ▶ **DC** : Defined Capacity, billing limit that you don't want to exceed
  
- ▶ **But performance can be affected (capping → IMSU is brought back to the DC) !**



# Conclusions

## Pricing models and calculations can be confusing but...

- Subcapacity pricing can save you money!

## There are more nuances to VWLC than can be adequately covered in an introductory presentation of this nature

- ooCOD, GCL, etc.

## It is possible to impact the bill for multiple products by reducing the R4H for a single LPAR.

- Subcapacity pricing charges by **LPAR usage**, so if the R4H for an LPAR decreases, it is possible that the bill for *every* product in that LPAR will go down.

## Intelligent Db2 tuning can be helpful!

# A Tactic to Consider

## Remember the SMF 89 records that are cut by certain products?

- These records indicate when the product is up and running – and – when it is down and inoperative?
- When you are not using the software, if at all possible, shut it down so it won't be charged if the peak occurs
- Example:
  - Peak **R4H** regularly occurs during the batch cycle between Midnight and 4AM
  - CICS is not active between the hours of 10PM and 6AM
  - Consider bringing CICS down before Midnight and up after 4AM
- Of course, this is not always feasible.
- Requires planning, sufficient time, and appropriate products

# Tactics for Db2 Tuning

## Understand where your monthly peaks are and attack the peaks.

- The SCRT reports the peak for each LPAR
  - *As well as the next highest peak usage time*
- When Db2 contributes to the peak, use Db2 tuning techniques to reduce consumption
  - Do not have to focus on the absolute peak hour
  - Remember, it is a FOUR HOUR rolling average
  - So anything running at the peak hour and the trailing 4 hours is fair game for tuning to reduce cost!

## When you reduce one peak, find the next and attack that one.

- Let's look at some techniques that can be used.

# Tuning “Outside the Box”

## Remember, you are charged for each product based upon the LPARs in which the product runs.

- You may be able to tune Db2 usage and reduce the cost of other products that run in the same LPARs
  - z/OS, CICS, IMS, COBOL, etc.
- Likewise, you may be able to tune non-Db2 usage and reduce the cost of Db2
  - *Consider:* say you have a batch VSAM program that consumes a lot of resources during your monthly peak, tuning that program to be more efficient could reduce your peak **R4H**, and thereby reduce the cost for all products running in that LPAR.
    - For example, running a REPRO to reorganize the file and eliminate CI and CA splits can improve VSAM performance and reduce resource requirements
    - And if that job ran in the same LPAR as Db2, it could also reduce the cost of Db2

# Where is Your Peak?

## What if your monthly R4H peak is in the batch window?

- Are you meeting your batch SLAs?
- Consider what can be done to reduce MSU consumption during batch processing.
- COMMIT processing? Concurrency and locking? SQL tuning?

## What if your monthly R4H peak is during your daily transaction processing?

- Have you optimized your transactions?
- Is the workload appropriate for your system?
- Consider what can be done to reduce MSU consumption during OLTP.
- CICS tuning? RDO specifications? SQL tuning?



- **Tune the SQL/application**

- Most relational tuning experts agree that the majority of performance problems with relational database applications are caused by poorly coded application code and SQL.

*as high as 70% to 80%*

- **Other Db2 tuning options**

- Tune the database
- Tune the subsystem
- Tune the environment

# SQL/Application Tuning

**Identify the heaviest resource consuming SQL during the peak processing time(s)**

**Identify tuning opportunities**

**Review access paths and tune**

- **Indexing, eliminating scans**
- **Let SQL do the work (not the program)**
- **Avoid sorting when possible**
- **Favor Stage 1 and Indexable predicates**
- **Review locking considerations**
  - UOW, lock avoidance, COMMIT, isolation level (dirty read)
- **Review cursors**
  - Multi-row FETCH

# Be Aware of the Whack-a-Mole Effect








<https://cloudframe.com/the-whack-a-mole-effect-tuning-mainframe-usage-for-cost-savings/>

# Summary



- **Consider cost reduction along with performance improvement by taking peak R4H into account**
- **Find the big resource consumers in your peak R4H window**
- **Look for ways to minimize resource consumption**
- **Find areas where SQL tuning and MRF can improve Db2 performance**
- **Consider a proactive approach with tools:**
  - Find tuning opportunities
  - Analyze workloads
  - Convert and tune SQL



# This presentation was prepared by:

**Craig S. Mullins**

**Mullins Consulting, Inc.**

**President & Principal Consultant**

**IBM Champion for Data and AI**

**IBM Gold Consultant**

[mullc@craigsmullins.com](mailto:mullc@craigsmullins.com)

[www.mullinsconsulting.com](http://www.mullinsconsulting.com)

