


**CMG 2015 – San Antonio, TX**

# **Best Practices for Improved z/OS Performance and Lower TCO**

By Ivan Gelb



# TRADEMARKS

- The following are trade or service marks of the IBM Corporation: CICS, CICS TS , CICSplex, DB2, IBM, MVS, OS/390, z/OS, Sysplex, Parallel Sysplex. Any omissions are purely unintended.
- © 2015 Gelb Information Systems Corp.  
URL: [www.gelbis.com](http://www.gelbis.com)  
Phone: 732-303-1333  
No part of this material can be reproduced by any means without prior written permission from the author and with proper attribution displayed.



# MOTHER OF ALL DISCLAIMERS (MOAD )

- All of the information in this document is tried and true. However, this fact alone cannot guarantee that you can get the same results at your place and with your skills. In fact, some of this advice can be hurtful if it is misused and misunderstood. As with all kinds of analysis, anything you may hear or read can be understood and misunderstood in many ways that may seem contradictory to you. Gelb Information Systems Corporation, Ivan Gelb and any one found anywhere assume no responsibility for this information's accuracy, completeness or suitability for any purpose. Anyone attempting to adapt these techniques to their own environments anywhere do so completely at their own risk. 😊



# CMG-T Sessions Today + Panel

- Part 1 - Best Practices for Improved z/OS Performance and Lower TCO
- Part 2 - WLM Caused Pain and Pleasure
- Part 3 - Holistic z/OS Performance and Capacity Management
- PANEL – zEnterprise Q&A with Kathy Walsh and Norm Holander




# Agenda

- Best Practices
- VWLC
- zAAP & zIIP
- IRD
- WLM
- z Others...

**Objective Expert Finding:** Apples-to-Apples studies and projects show that z/OS TCO beats any LUW for applications of as little 8 MSUs and greater. Larger projects → Higher percentage, not just higher total, \$\$\$ savings.

**Note:** ➡ symbol used to flag Rules-of-Thumb (ROT) and recommendations



# Best Practices / Recommendations - 1

- ➡ Effective CP & PM activities will yield required service level quantity and quality for least total cost. Now go and convince your upper management of this!
- ➡ Design WLM service policy to mirror business activity – this enables the most effective CP & PM activities. You will know when you need to do something, and what you may need to pay closer attention to.
- ➡ Always communicate with management in common business terms: orders, deliveries, effect on total cost.
- ➡ Always avoid technical jargon and acronyms when communicating with business people.



# Best Practices / Recommendations – 2 (BP/R)

- ➡ Given less time and more data to analyze, choose your tools and techniques so **YOUR** effectiveness is improved
- ➡ Practice routine WLM service policy health checks. Ask: “Is it still working as we intended to?”
- ➡ Choose tools and techniques that enable analysis of each workload independently and in combination with present and future workloads
- ➡ For capacity planning studies, insure that you isolate workloads not just along business importance but also based on key attributes that affect scalability: IO intensity, virtual storage needs, use of z/OS services, total CPU time in applications code, etc...



# BP/R – 3: CICS Regions Consolidation

- Recall why each region was created, if you can, and then it is time to put “Humpty Dumpty” back together again.
- Very large potential savings when fewer regions are run.
- Study potential benefits of making CICS applications threadsafe. 80/20 rule works very well for this activity!





# BP/R – 4: PR/SM Considerations

- If total PR/SM overhead greater than about 1.5%, try to figure out what is causing it and is it worth it.
- Run minimum possible number of LPARS
- Define the minimum required logical processors



# Variable Workload License (VWLC) Charging Method for Software

## Audience Poll:

1. **Sub-capacity licensed now with IBM?**
2. **Sub-capacity licensed now with other z-Software Vendors?**



# VWLC Opportunity

- Can we save \$ without hurting anything?  
A(nswer): Yes
- How to obtain management support for activity?  
A: Demonstrate the estimates of savings possible under your software capacity planning scenarios
- When will ALL software vendors participate in VWLC?  
A: When **WE** “persuade” them that there is no other way to do business with us.
- What to do now?  
A: (1) Education; (2) Develop software capacity plans; (3) Prepare estimates of savings/costs



# VWLC Overview - 1

- Variable workload license (VWLC) charging method available in USA since March 2001 for selected IBM software products. Examples: z/OS, COBOL, CICS, DB2, CICS, IMS, MQSeries plus over 25 more.
- Started sub-capacity software licensing trend. Software license capacity can be dramatically lower than installed hardware capacity.
- Concept moving ahead very slowly in the independent software vendors' (ISV) world.



# VWLC Overview - 2

- Basis for sub-capacity of VWLC products is LPAR utilization
- Monthly charge based on highest rolling 4 hour average by product summed for LPARs w. software present in them
- Product isolation into LPARs for software capacity planning is a potentially cost saving activity
- 5 – 15% monthly software cost savings are possible
- LPAR's total capacity may be capped via PR/SM to comply with software license agreement



# VWLC Overview - 3

- Sub-capacity Reporting Tool (SCRT) assists the determination of LPAR utilization where VWLC products are running.
- ➡ For download and details, visit:
  - <http://www.ibm.com/zseries/swprice/scrt>
- ➡ Also obtain the WLC Tool from the above URL. It reports the rolling 4 hour average for LPARs.
- ➡ Reference:  
Using the Sub-Capacity Reporting Tool, SG24-6522



**z Application Assist Processor  
(zAAP)**

**z Integrated Information  
Processor (zIIP)**

**Audience Poll:**

- 1. Java programs already on your zSeries CPU?**
- 2. What is the vehicle for Java code? Batch work? CICS? WebSphere? Other?**
- 3. DB2 Version 8? z/OS 1.6 or higher?**
- 4. zIIP or zAAP processors installed?**



# zAAP Opportunity

- Can we save \$s?  
A: Yes if you have workloads with Java content on z/OS 1.6 running on z890 and z990 processors.
- When to do it?  
A: As soon as you have the required hardware and software levels to support it.
- How to guarantee management support?  
A: Just estimate the initial savings. Later, keep reporting the ongoing savings in hardware and ALL software costs;
- What to do?  
A: Use available tools to estimate potential savings.





# zAAP Overview - 1

- zSeries® Application Assist Processor (zAAP), available starting with the IBM eServer zSeries 990 (z990) and zSeries 890 (z890) servers, is a specialized processing unit that provides a z/OS Java™ execution environment
- Unlike standard CPs, ICFs and IFLs, zAAPs can do nothing on their own
- Like ICFs and IFLs, zAAPs do not add to software costs processing on the standard CPs.
- zAAPs are designed to operate asynchronously with the general processors to execute Java programming under control of the IBM Java Virtual Machine (JVM)



# zAAP Overview - 2

- Prerequisites:
- (zAAPs) may be purchased and installed on all z9s, z990 and z890 servers and future follow-on models only.
- The operating system must be migrated to:
  - z/OS 1.6 (or z/OS.e 1.6), the IBM Solution Developers Kit (SDK) for z/OS, Java 2 Technology Edition, V1.4 with PTF (or later) for APAR PQ86689 must be used and for WebSphere-based Java workloads, WebSphere Version 5.1 or above is required.



# zAAP Overview – 3


- zAAPs can offer a number of cost advantages:
  - Low cost of acquisition - \$125,000 USD each
  - Significantly lower maintenance costs than the general purpose CPs
  - No IBM software charges on zAAP capacity
  - Potential to reduce costs for sub-capacity eligible IBM software because of lower 4 hour rolling average utilization on general purpose CP based LPARs.
- ➔ For more information visit:  
<http://www-1.ibm.com/servers/eserver/zseries/zaap/>



# VWLS and zAAP Example

## BEFORE zAAP:


- Machine Type: 2084-B16
- Rated @ 647 MSUs
- Sub-Capacity Pricing based on;
  - LPAR A rolling 4hr avg @ 233 MSUs
  - LPAR B rolling 4hr avg @ 200 MSUs
  - LPAR C rolling 4hr avg @ 267 MSUs
- Rolling 4hr avg of Machine = 547 MSUs
- Average Prime Shift Machine Utilization = 80%



LPAR A	LPAR B	LPAR C
80%	70%	90%
233	200	267

## AFTER zAAP:

- Machine Type: 2084-B16
- Rated @ 647 MSUs
- Java cycles executed on zAAPs
- New Sub-Capacity Pricing on *reduced* rolling 4hr avg
  - LPAR A rolling 4hr avg @ 213 MSUs
  - LPAR B rolling 4hr avg @ 100 MSUs
  - LPAR C rolling 4hr avg @ 167 MSUs
- New Rolling 4hr avg of Machine = 480 MSUs
- Average Prime Shift Machine Utilization = 50%



LPAR A	LPAR B	LPAR C	zAAP
70%	30%	50%	
213	100	167	

Average LPAR Utilization	Sub-Capacity MSUs
80%	233
70%	200
90%	267
70%	213
30%	100
50%	167

Illustration Source: "z890 and z990 zAAP – What it Can Do for You," By Kathy Walsh, IBM Corp.



# zIIP Overview - 1

- z9 Integrated Information Processor (zIIP), available on z9 BC and EC mainframes, benefit DB2 V8 workloads:
  - **ERP or CRM application serving** - For applications, running on z/OS, UNIX, Intel, or Linux on System z that access DB2 for z/OS V8 on a System z9, via DRDA over a TCP/IP connection, DB2 gives z/OS the necessary information to have portions of these SQL requests directed to the zIIP.
  - **Data Warehousing applications** – Requests that utilize DB2 for z/OS V8 for long running parallel queries, including complex star schema parallel queries.
  - **Some DB2 for z/OS V8 utilities** – A portion of DB2 utility functions used to maintain index maintenance structures (LOAD, REORG, and REBUILD INDEX) that typically run during batch, can be redirected to the zIIP.



# zIIP Overview - 2

- (zIIPs) may be purchased and installed on z9 BC and EC servers and future follow-on models only.
- Additional requirements:
  - z/OS 1.6 (or z/OS.e 1.6)
  - DB2 Version 8
  - PR/SM enabled



# Intelligent Resource Director (IRD)

## Audience Poll:

1. IRD in production use?
2. Any negative IRD experiences or rumors?



# IRD Opportunity

- Is anyone using it now?

A: Yes, about 20 – 25% of the sites that can.

- Can IRD help us?

A: Yes – automates shifting the processor and IO capacity among LPARs within a SYSPLEX.

- Can IRD hurt us?

A: Not if you do your “homework”

- What to do now to use IRD later?

A: (1) Prepare an effective WLM service policy; (2) Test IRD; (3) Implement it





# IRD Overview

- Intelligent Resource Director (IRD) uses: Partition Resource/System Manager (PR/SM), Workload Manager (WLM), I/O Supervisor (IOS), Channel Subsystem, and Parallel Sysplex
- IRD dynamically moves resources to most important work not meeting service level goals
- IRD components:
  1. LPAR CPU management
  2. Dynamic Channel Path Management
  3. Channel Subsystem Priority Queuing



# IRD - LPAR CPU management

- LPAR CPU management adjusts LPAR weight and the number of online logical CPUs as required by WLM
- Works to meet goals of important work as determined from WLM service policy
- If weight of an LPAR is raised or decreased, its number of online logical CPUs will be changed to maintain a match between logical and physical CPU speed
- On various LPAR Activity reports, look for fractional number of CPUs as evidence of IRD's activities



# IRD – Dynamic Channel Path Management

- IRD DCM dynamically moves channel paths through the ESCON Director from one I/O control unit to another
- Must define channel paths as “managed”
- Improves performance by moving bandwidth to where it is needed the most
- Can improve availability – failed channel dynamically replaced
- **DCM will consider changes for each LPAR every 10 seconds, but only ONE LPAR will be changed per interval, AND same LPAR will be changed only once / minute**



# IRD – Channel Subsystem Priority Queuing

- Channel Subsystem Priority Queuing (CSSPQ) sets I/O request priorities
- WLM can change a workload's I/O priority if the work is not meeting goals and I/O activity delays are the reason
- With Enterprise Storage Server (ESS, a.k.a. Shark), WLM can pass a priority to the control unit
- WLM sets priorities (no external user controls) as follows:
  - System work is at highest priority (range is 1 – 16)
  - Highest importance work missing goal
  - Light weight I/O user have higher priority than heavy weights
  - Discretionary work is lowest priority



# Workload Manager (WLM)

## Audience Poll:

- 1. Does WLM Goal mode deliver the service levels you hoped for?**
- 2. Do you schedule WLM service policy “checkups” at regular interval, or you wait until.....you must!**




# WLM Opportunity

- Can WLM help us?  
A: Yes
- Can WLM hurt us?  
A: Yes, and it is very likely doing it right now!
- What to do now to use WLM better?  
A: (1) Insure that the policy reflects the business priorities; (2) Create effective report classes for low level analysis; (3) Monitor and improve the service policy continuously – especially with velocity goals.
- “Our WLM service policy is perfect because Guru IM-so-Perfect developed it.”  
A: ? ? ? ? OK! Hello! I am your new Guru...Perfection passed on while you were looking elsewhere.



# WLM Advice

-  **Recommendation:** Create resource groups for any workload you wish to control regardless of processor utilization level.

**Example 1:** A service class to keep discretionary work from raising R4HA (rolling 4 hr. average)

**Example 2:** service class can be limited to maximum of 1 service unit / second rate



# What Does WLM Manage?

WLM managed resources:

- Processor (CPU)
- Storage controls
- Multi-programming levels (MPL)
- Physical I/O priority
- Dynamic parallel I/O access volumes (PAVs)
- JES batch initiators
- DB2 stored procedure address spaces
- Websphere scalable address spaces





# WLM Role in CP & PM

- WLM is single most critical success factor (CSF) for CP & PM
- Insure that critical business workloads are captured in service policy so they are easy to observe and analyze.
- WLM exercises control on following:

WHAT?	HOW?
CPU access priority	Task dispatch priority guided by importance and service level goal
CPU time limits	Defined via resource groups
I/O performance	Priority propagations & PAV-s (parallel access volumes)
Enclaves for DDF, stored procedures, etc...	Coded min/max service level definitions
Dynamic batch initiators	Goal and resource driven controls
Storage paging	Isolation to protect working set size



# Z Others

## Audience Poll:

1. z/VM,
2. IFL,
3. zBX, and
4. Specialty IBM database processors



# In Closing...(alphabetical order!)

- IRD – the final piece required for automation of on demand computing's performance
- VWLC – opportunity to save 5 – 15% of annual and potentially large one-time software costs
- WLM – the critical success factor for **ALL** performance management (PM) and capacity planning (CP) activities for lowest TCO
- zAAP & zIIP– opportunity: improve performance of Java and DB2 workloads while you reduce hardware and **ALL** software costs
- Z Others: z/VM, IFL, zBX, and specialty IBM database processors are industry leading hardware and software alternatives which outperform and under-price LUW options.



# Z Others...

- Consider the industry-unique investment protection IBM offers with the z Series specialty processors.
- z/VM, IFL, zBX, and specialty IBM database processors produce massive \$ savings over the competition with:
  - Better performing hardware and software solutions
  - Reduced hardware and software costs due to the efficiencies
- Any questions about these offerings?



# Need / Want to Know More - 1

- **Start at:**  
[www.ibm.com/servers/eserver/zseries/](http://www.ibm.com/servers/eserver/zseries/)
- **Large Systems Performance Reference:**  
<http://www-1.ibm.com/servers/eserver/zseries/lpr/>
- **“HOT TOPICS” a z/OS newsletter:**  
[www.ibm.com/servers/s390/os390/bkserv/hot\\_topics.html](http://www.ibm.com/servers/s390/os390/bkserv/hot_topics.html)
- **Computer Measurement Group (CMG):**  
[www.cmg.org](http://www.cmg.org) and [www.cmg2006.com](http://www.cmg2006.com)
- **SHARE:** [www.share.org](http://www.share.org)



# Need / Want to Know More - 2

- **IRD Information:**  
[www.ibm.com/servers/eserver/zseries/ird](http://www.ibm.com/servers/eserver/zseries/ird)
- **Washington Systems Center (WSC) documents:**  
[www.ibm.com/support/techdocs](http://www.ibm.com/support/techdocs)
- **VWLC Information:**  
[www.ibm.com/servers/eserver/zseries/swprice](http://www.ibm.com/servers/eserver/zseries/swprice)
- **WLM Information:**  
[www.ibm.com/servers/eserver/zseries/zos/wlm](http://www.ibm.com/servers/eserver/zseries/zos/wlm)
- **zAAP:**  
<http://www-1.ibm.com/servers/eserver/zseries/zaap/gettingstarted/>



# Time for...

