z/OS Workload Manager
Revisiting Goals Over Time

z/OS Performance
Education, Software, and
Managed Service Providers

PIVOTOR®

Creators of Pivotor®
www.pivotor.com

Instructor: Peter Enrico

Email: Peter.Enrico@EPStrategies.com

Enterprise Performance Strategies, Inc.
3457-53rd Avenue North, #145
Bradenton, FL 34210
http://www.epstrategies.com
http://www.pivotor.com

Voice: 813-435-2297
Mobile: 941-685-6789

Performance Workshops Available

During these workshops you will be analyzing your own data!

- **WLM Performance and Re-evaluating of Goals**
    - Instructor: Peter Enrico
    - September 15 – 19, 2014   - Kansas City, Missouri, USA
    - October 20 - 24, 2014   - Munich, Germany

- **Parallel Sysplex and z/OS Performance Tuning**
  (Web / Internet Based!)
    - Instructor: Peter Enrico
    - July 29 – 31, 2014 (Web)
    - Considering a fall class...

- **Essential z/OS Performance Tuning Workshop**
    - Instructors: Peter Enrico and Tom Beretvas

- **z/OS Capacity Planning and Performance Analysis**
    - Instructor: Ray Wicks

Contact, Copyright, and Trademark Notices

Questions?

Send email to Peter at Peter.Enrico@EPStrategies.com, or visit our website at http://www.epstrategies.com or http://www.pivotor.com.

PIVOTOR®

www.pivotor.com

Copyright Notice:

© Enterprise Performance Strategies, Inc. All rights reserved. No part of this material may be reproduced, distributed, stored in a retrieval system, transmitted, displayed, published or broadcast in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Enterprise Performance Strategies. To obtain written permission please contact Enterprise Performance Strategies, Inc. Contact information can be obtained by visiting http://www.epstrategies.com.

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check®, Reductions®, Pivotor®**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM®, z/OS®, zSeries® WebSphere®, CICS®, DB2®, S390®, WebSphere Application Server®, and many others.

Other trademarks and registered trademarks may exist in this presentation

---

WLM Reports Processing/Discussion Offer !!!

- ☐ Special Reports Offer!
  - ■ See your WLM SMF 72.3 records in chart and table format

  - ■ Please contact me, Peter Enrico for instructions for sending raw SMF data
    - ☐ Send an email to peter.enrico@epstrategies.com

  - ■ Deliverable: Dozens of coupling facility based reports (charts and tables)
    - ☐ WLM SMF30 Address Space Analysis
    - ☐ WLM - Period Setup Analysis
    - ☐ WLM - Importance Level Analysis
    - ☐ WLM - PI Analysis
    - ☐ WLM - Velocity Goal Analysis
    - ☐ WLM - Response Time Goal Analysis
    - ☐ WLM - Discretionary Goal Analysis
    - ☐ WLM - Enclave Analysis
    - ☐ WLM - Multiple Period Analysis
    - ☐ WLM - CPU Analysis
    - ☐ WLM - Storage Analysis
    - ☐ WLM - DASD IO Analysis
  - ■ One-on-one phone call to explain your coupling facility measurements

Presentation Overview

- What does it mean to 'revisit' and 'reevaluate' goals?
- Why revisit and reevaluate goals?
  - Scenario 1: Improperly set goals or controls
  - Scenario 2: Over time, goals are now regularly being missed
  - Scenario 3: Planning for environmental changes
  - Scenario 4: Exploitation of additional WLM functions
  - Scenario 5: Changes to WLM, system problem, improper tuning
  - Scenario 6: Changes to business priorities and objectives
  - Scenario 7: Inaccuracy of reported measurements
  - Scenario 8: Exploitation of non-WLM functions that influence performance
  - Scenario 9: Occasionally 'something happens'

Note: Some causes for each can overlap

Note: This one hour presentation is a subset of a much longer comprehensive presentation and workshop materials on this subject

---

Revisiting and Reevaluating Goals

- Revisiting Goals
  - The act of determining if a current set of goals needs to change
    - Usually done when a change is being planned to the environment
    - Done to preempt possible WLM goal mode problems

- Reevaluating Goals
  - The act of determining if a new value of a goal or WLM control
    - Usually done when goals are being missed or workloads are not performing as well as expected
    - Done to fine tune goals or to resolve problems

- The goal of this presentation
  - To help you think about the types of situations that should cause you to take another look and rethink your WLM goals and settings

  - Provides just some of the many examples I've seen and worked on

Scenario 1: Improperly Set Goals or Controls

- ☐ You inherited the responsibility for WLM, and your systems have been in WLM goal mode for many years now
  - ■ But some goals are never been regularly met
  - ■ Why can some goals never be met?

- ☐ Many common causes
  - ■ Improper velocity or response time goals (too aggressive or too easy)
  - ■ Using average response time goals instead of percentile response time goals
    - ☐ 'Outlier' transactions may be skewing the calculated average
  - ■ Improper WLM importance level
  - ■ Improper period duration
  - ■ Unlike work in period
  - ■ Not enough work in Discretionary
 - ■ Improper use of resource group minimums or maximums
  - ■ etc...

Example: Average response time goals should not be used. Instead... convert all Average RT goals to Percentile RT goals

- ☐ Average Response Time Goal:
  - ■ The average response time desired for a given set of ended transactions

$$\text{Average Response Time} = \frac{\text{Sum of Elapsed Time for Ended Transactions}}{\text{Number of Ended Transactions}}$$

- ☐ Average response time goals are not recommended since they can be easily influenced by 'outlier' transactions
  - ■ Example
    - ☐ Average response time goal set to 1 second
    - ☐ 99 transactions complete in 1 second, but 1 transaction completes in 2 minutes
    - ☐ Average response time achieved is 2.2 seconds
    - ☐ Goal missed even though 99% of transactions completed within 1 second
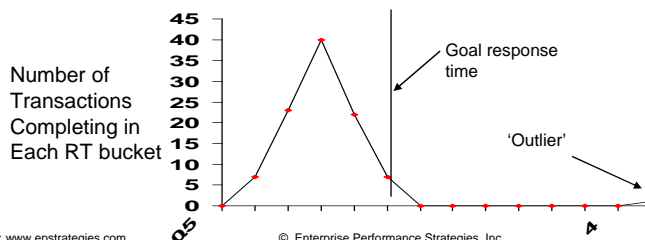
Number of Transactions Completing in Each RT bucket
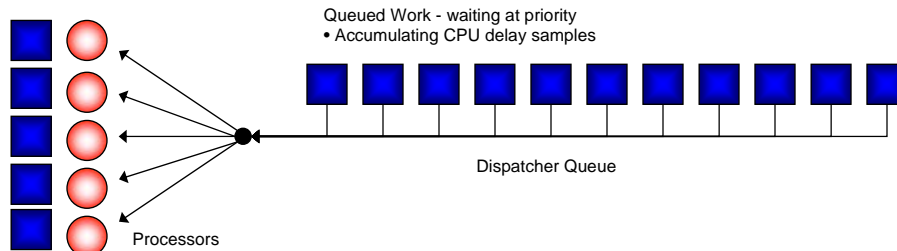
Goal response time

'Outlier'

Example: Your velocity goals may be too aggressive

☐ Periods with more address spaces or enclaves (dispatchable units) tend to have more CPU queuing
- Since delays are inherent such environments usually result in lower velocities
- Example:
  - ☐ Have a period with 100 active IMS MPRs
  - ☐ All want CPU but we are only running on a 5-way processor
  - ☐ Velocity of 50 impossible
  - ☐ Velocity goal of 10 more typical

Dispatched Work
• Accumulating CPU Using Samples

Queued Work - waiting at priority
• Accumulating CPU delay samples

Dispatcher Queue

Processors

Example: Do not use Velocity to prioritize work

☐ How many WLM Service Definitions have you seen with something like this?

- Which is the more aggressive goal?
  - ☐ That is... which period will WLM try to treat better than the other?

Service Class	Period	Imp	Goals
CICSHIGH	1	1	Velocity 50
CICSMED	1	1	Velocity 20

Scenario 2: Over Time, Goals Are Now Regularly Being Missed

- We tuned our WLM goals settings a number of year ago...
  - But time has passed and now goals are being missed
  - Why are goals being now missed?
  - Why are the workloads not being treated as well as they should be?

- Common causes
  - Workload growth
    - Application workloads
    - Workload growth in SYSSTC
    - Growth to system address spaces and monitors
  - Workload reduction
  - New work introduced into the system or sysplex
  - Improperly tuned system
  - Changes in the software environment
  - Changes in the hardware environment

Scenario 2: Caused by Workload Growth

- When workloads grow they tend to consume more resources
  - Growing workloads require a larger share of the (static) capacity
  - WLM policy adjustment algorithms will allocate resource to meet goals
    - Discretionary and low importance periods should start suffering before higher importance workloads

- Indicators from measurements and monitors
  - Higher PIs of lower importance workloads
  - Increase in transaction activity
  - Increase consumption of processor, storage, etc. for periods with workload growth
  - Less available resources for delays causing higher PIs
    - Fewer available frames
    - High CPU Busy Time percentages
  - etc..

- Recommended actions:
  - Ease lower importance goals
  - Obtain more capacity
  - Accept that goal is sometimes missed, but be able to explain why

Example: Growth in SYSTEM and SYSSTC

- □ Too much work in SYSSTC allows less CPU to goal periods
  - ■ Remember: CPU Available to goal periods and discretionary =
    100% CPU – ((CPU Consumed by SYSTEM) –
  - ■ (CPU Consumed by SYSSTC))

- □ Too little work in discretionary may not provide donor resources
  - ■ Remember: WLM first looks for free resources. If not enough free resources than takes first from discretionary before taking from goal periods
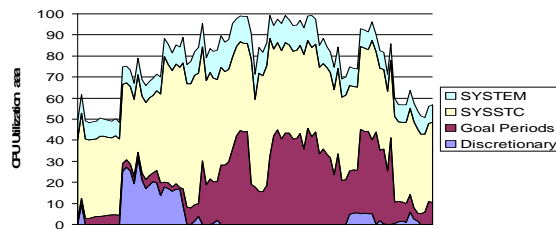
255	SYSTEM
254	SYSSTC
253	
	'Unused'
249	(SYSSTC1-5)
248	Small Consumer
247	
	Priorities Used for RT or Velocity Periods (i.e. Imp 1 – 5)
203	
202	Unused
201	
	Discretionary (MTTW)
192	
191	Quiesce

Only classify to SYSTEM and SYSSTC work that needs to be there

Keep CPU consumed by SYSSTC and SYSTEM to less than 10%.

Actually, less than 10% on MP LPARs
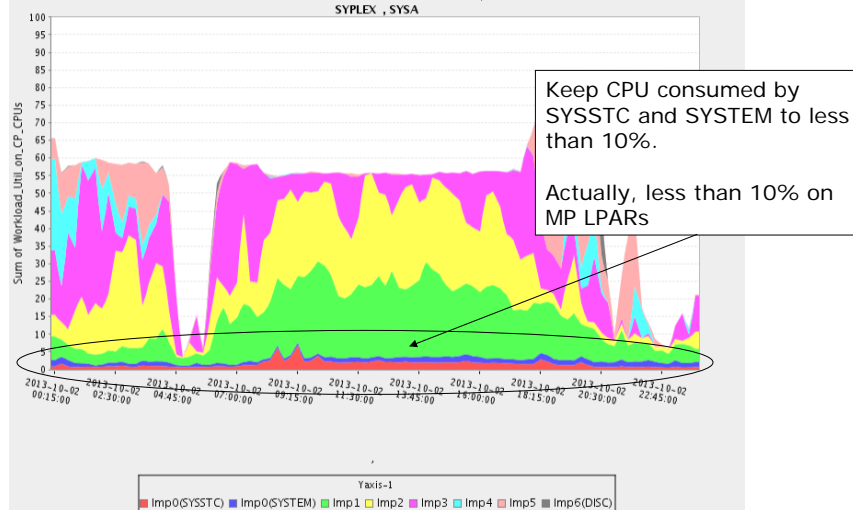
Scenario 2: Caused by Workload Reduction

- If a workload shrinks it is possible for goals to be missed
  - Fewer using and delay samples calculate a less meaningful velocity
  - Fewer ended transactions calculate a less meaningful response time
  - Less discretionary work could mean less discretionary resource donors

- Indicators:
  - Velocity goals
    - Fewer address spaces or enclave transactions contributing samples (see next foil)
  - Response time goals:
    - Fewer than 10 ended transactions within a 20 minute period of time
  - Discretionary goals:
    - Shortage of available resources and little discretionary work to steal from

- Recommended Actions:
  - Investigate consolidation of periods
  - Investigate easing of goal and/or importance
  - Accept that goal is sometimes missed, but be able to explain why

---

Example: Velocity May Be Affected By Workload Reduction

- Amount of work contributing to samples affects calculated velocity
  - Velocities are calculated on a service class (goal) period basis

  - Number of address spaces or enclaves in period will influence:
    - Number of using and delay samples collected
      - Fewer address spaces or enclaves mean fewer contributing samples
    - Calculated velocities (will be more sensitive)

$$\text{Velocity based on } \frac{5000}{5000+10000} \text{ same as Velocity based on } \frac{5}{5+10} = 33\%$$
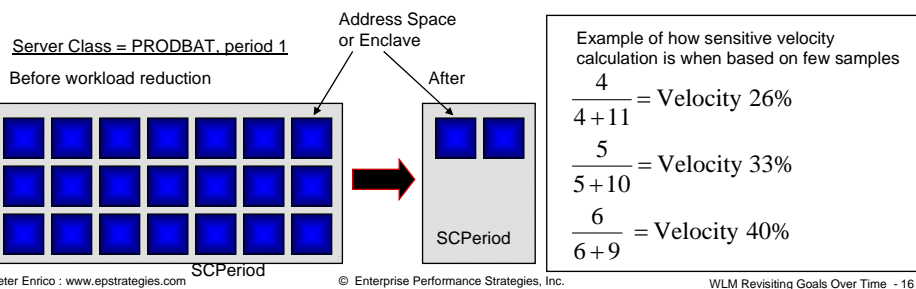
Server Class = PRODBAT, period 1

Before workload reduction

Address Space or Enclave

After

SCPeriod

Example of how sensitive velocity calculation is when based on few samples

$$\frac{4}{4+11} = \text{Velocity } 26\%$$

$$\frac{5}{5+10} = \text{Velocity } 33\%$$

$$\frac{6}{6+9} = \text{Velocity } 40\%$$

Scenario 2: Other Possible Causes

- New work introduced into the system or sysplex
  - New work may take needed resources from lower importance work

- New workload types
  - More vendors are using WLM services
  - Workload for new subsystem type that is not properly classified
    - For example, minor workload types like TCP, NETV, and others

- Changes in the software environment
  - New / changed software may require more resources

- Changes in the hardware environment
  - Changes to processor environment
    - (Slower or faster) and/or (fewer or more) processors and/or HiperDispatch
  - Less available storage due to some other system facility
    - Example: Very large WAS heaps
  - Changes to I/O subsystem cache or device clusters

Scenario 3: Planning for Environmental Changes

- All is running great in goal mode
  - But environmental changes are planned
  - Do we need to consider making any WLM changes?

- Typical changes to be concerned with
  - Changes in hardware
    - Changes to processor environment
      - Introduction of new processors
      - Changes to speed of processors (faster or slower)
      - Changes to number of processors (more or less)
      - Changes to both number and speed of processors
      - Changes to the LPAR definitions affecting share of processor
      - Changes to capacity controls such as HiperDispatch

  - Changes to the sysplex configuration
    - Introduction or removal of systems
    - Redistribution of workloads across the sysplex

Example: How Processors Can Affect Velocity

☐ **CPU delay could be inherent to many periods**
- ■ Few engines than work that wants to run
  - ☐ Periods with more ready address spaces than processors will have delay
- ■ Faster engines
  - ☐ May cause using samples of quick running work to be missed
- ■ Other causes of delay include
  - ☐ Reduced preemption
  - ☐ Fair share dispatching

Dispatched Work
• Accumulating CPU Using Samples
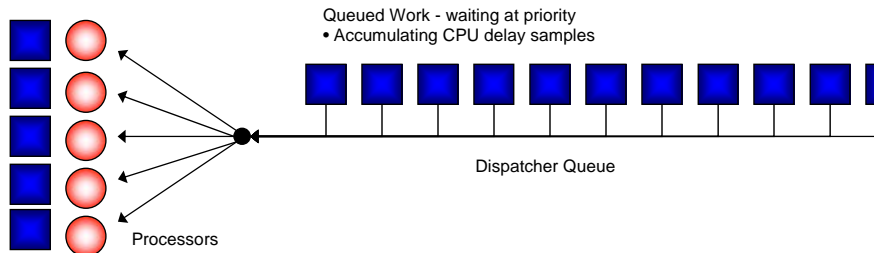
Queued Work - waiting at priority
• Accumulating CPU delay samples

Dispatcher Queue

Processors

Example: Remember that goals are Sysplex wide

☐ **Sysplex of a little of everything**
- ■ Asymmetrical capacity local systems in same Sysplex where some workloads with mutually exclusive and disparate workloads
  - ☐ Note: Probably all installations have this type of Sysplex in one form or another
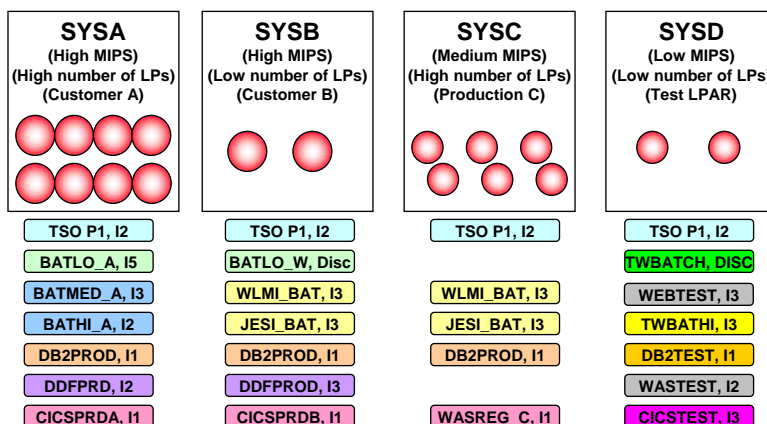
SYSA (High MIPS) (High number of LPs) (Customer A)	**SYSB** (High MIPS) (Low number of LPs) (Customer B)	**SYSC** (Medium MIPS) (High number of LPs) (Production C)	**SYSD** (Low MIPS) (Low number of LPs) (Test LPAR)
TSO P1, I2	TSO P1, I2	TSO P1, I2	TSO P1, I2
BATLO_A, I5	BATLO_W, Disc		TWBATCH, DISC
BATMED_A, I3	WLMI_BAT, I3	WLMI_BAT, I3	WEBTEST, I3
BATHI_A, I2	JESI_BAT, I3	JESI_BAT, I3	TWBATHI, I3
DB2PROD, I1	DB2PROD, I1	DB2PROD, I1	DB2TEST, I1
DDFPRD, I2	DDFPROD, I3		WASTEST, I2
CICSPRDA, I1	CICSPRDB, I1	WASREG_C, I1	CICSTEST, I3

Example: Remember that goals are Sysplex wide

- How do you assign response time or velocity goals in an asymmetrical environment?
  - A goal could be too easy for SYSA, but too hard for SYSB
- Sometimes classification by system is necessary (Qualifier SY)
  - If work is running on SYSA but it to a service class period with a different goal than if the work is running on SYSB
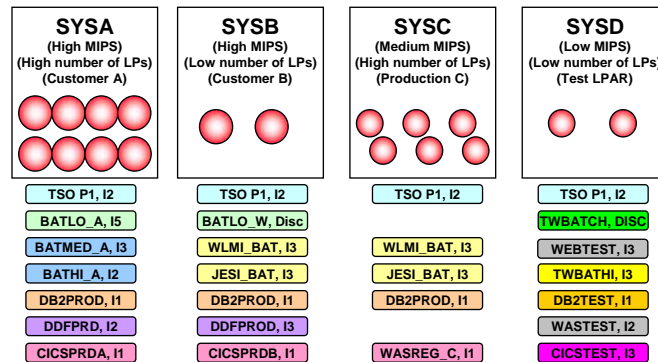
SYSA (High MIPS) (High number of LPs) (Customer A)	**SYSB** (High MIPS) (Low number of LPs) (Customer B)	**SYSC** (Medium MIPS) (High number of LPs) (Production C)	**SYSD** (Low MIPS) (Low number of LPs) (Test LPAR)
TSO P1, I2	TSO P1, I2	TSO P1, I2	TSO P1, I2
BATLO_A, I5	BATLO_W, Disc		TWBATCH, DISC
BATMED_A, I3	WLMI_BAT, I3	WLMI_BAT, I3	WEBTEST, I3
BATHI_A, I2	JESI_BAT, I3	JESI_BAT, I3	TWBATHI, I3
DB2PROD, I1	DB2PROD, I1	DB2PROD, I1	DB2TEST, I1
DDFPRD, I2	DDFPROD, I3		WASTEST, I2
CICSPRDA, I1	CICSPRDB, I1	WASREG_C, I1	CICSTEST, I3

Scenario 3: Planning for Environmental Changes

- Typical changes to be concerned with cont…
  - Changes to processor environment or PR/SM settings
    - Weights, HiperDispatch, soft or hard capping controls

  - Changes in software
    - Upgrades to new level of the operating system
    - Different product levels of software on systems in sysplex
    - Upgrades to levels of software products
    - Changes to WLM functionality of workload management

  - Merging of data centers, or Sysplexes, or Systems

  - Workload changes
    - Introduction of new workload
    - Removal of a workload
    - Moving a workload from one system to another

  - Any changes that affects capacity

Scenario 4: Exploitation of Additional WLM Functions

- **All is running great in goal mode**
  - But we are interested in exploiting some WLM functionality that is currently not being used
  - Do we need to consider goals when making these WLM changes?

- **New functions waiting to be exploited**
  - New subsystem types
  - New classification qualifiers and group options
  - WLM managed initiators
  - Scheduling environments
  - I/O priority management if it is still turned off
  - CPU critical control
  - I/O priority groups (Normal or High) (like CPU critical but for I/O)
  - CICS or IMS transaction management
  - Storage critical controls
  - Resource groups
  - WLM-Managed DB2 Buffer Pool Adjustment

Long Term CPU Protection - Background

- **CPU problem inhibiting installations from migrating to goal mode**
  - Some installations are concerned that WLM will not react fast enough for high priority work

- Importance 1 (circle)
- Importance 2 (triangle)
- Importance 3 (square)
- Importance 4 (diamond)
- Importance 5 (pentagon)

DP	
250	●
245	●
240	● ● ▲ ▲
235	▲ ▲ R
230	▲ ■ ■
225	◆ ◆ ◆
220	⬠ ⬠ ⬠
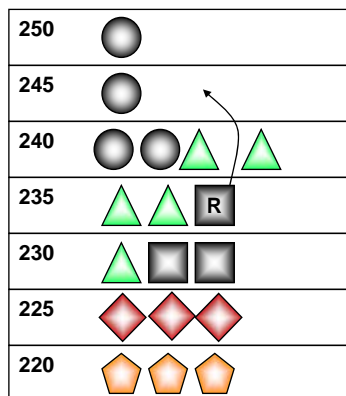
Note: To make the point, just a few priorities between DP 208 and DP253 are shown.
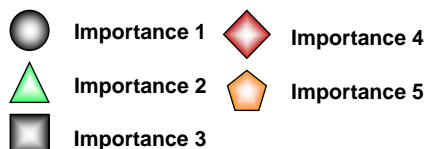
- In goal mode, DPs tend to be ordered by importance

- If work is missing its goal WLM may decide to put it at a DP at or above a higher importance period

- The problem occurs when this lower importance period starts to consume more CPU and causes the higher importance period to miss its goal

- WLM will recognize this condition and fix it … but it can be slow to react

Example: CPU Critical Control

- Long Term CPU Protection Control
  - Objective is to help ensure that critical work will have a higher CPU DP than lower importance work
  - For CPU sensitive work

- Original Problem:
  - Some installations are concerned that WLM will not react fast enough for high priority work

- When new 'CPU Critical' indicator = 'Yes'
  - Lower importance work will 'generally' receive a lower CPU DP than work marked 'CPU Critical'
  - Some exceptions are made for enqueue promotions

Example: Manage Region Using Goals of ….

- Installation's have several choices for how WLM will manage their CICS and IMS workloads
  1. Region Management:
     - Allow all regions to be managed towards velocity goals and importance levels
     - Response time goals of transactions have no influence

  2. Transaction Management:
     - Allow all regions to be managed to meet the response time goals of the transactions they are serving.
     - The goals and importance levels of the regions are ignored.

  3. Combination of Region and Transaction Management
     - Allows certain workloads to manage with 'Region Management' and other workloads to be managed with 'Transaction Management'
     - Usually used when Transaction Management is not effective for certain workloads

  4. Transaction Management while still honoring goal and importance of select regions
     - Allow identified regions to have their goal honored.
     - So full Transaction Management, but select region goal and importance level is used.

Option 4 - Transaction Management with Region Importance (keyword both)

- ☐ WLM will allow the goal and importance level of select regions to be honored when managing them to meet the goals of the transactions they are serving
  - ■ Sometimes you may want WLM to favor CICS TORs or IMS Control regions more

- ☐ Example 1: All regions assigned the same service class
  - ■ Some regions in service class have goal honored, and others have goal ignored
  - ■ Confuses the evaluation of region service class measurements

    - ☐ Service class CICSREGS   Velocity 60, Importance 1

- ☐ Example 2: Separate certain regions into different service classes
  - ■ Goal for regions specified as 'BOTH' will be honored
  - ■ Different importance levels allow WLM to favor management of high importance regions over lower important regions

    - ☐ Service class CICSTORS   Velocity 60, Importance 1, BOTH
    - ☐ Service class CICSAORS   Velocity 60, Importance 2

IMP 1, Velocity 60

IMP 2, Velocity 60

Additional Considerations

- ☐ When Importance level assigned to regions is considered, WLM will consider this importance level when managing the regions in the $SMFSxxx periods

Processors

| A | A | A | A | A | A | T | A | A | T |

Although a terrible / ill thought out IBM design, BOTH is meant to help this situation

| T | T | A | A | A | A | A | A | A | A |

Processors

Additional Considerations

☐ When Importance level assigned to regions is considered, WLM will consider this importance level when managing the regions in the $SMFSxxx periods

CICSTRX1

CICSTRX2

$SRMS001

CICSTRX3

CICSTRX4

CICSTORs

CICSTORS
Velocity 60
Importance 1
BOTH

CICSAORS
Velocity 60
Importance 2

| T | T | A | A | A | A | A | A | A | A |
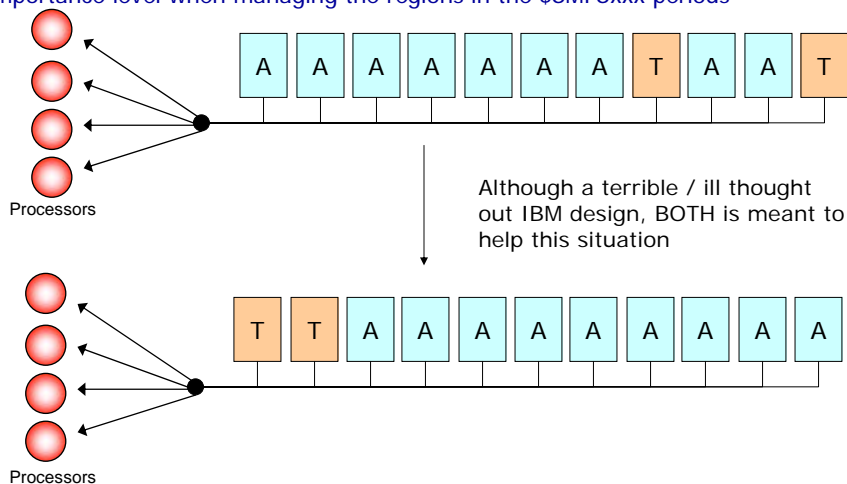
Processors

Example of a Topology – with BOTH option

AORs

$SRM001

CICSHIGH — Imp1
RT 90% within X

CICSMED — Imp2
RT 90% within Y

$SRM002

CICSOTHR — Imp3
RT 75% within Z

AORs

CICSTORS — Imp1
Velocity 60

TORs

FORs

All other transactions here

CICSAORS — Imp2
Velocity 60
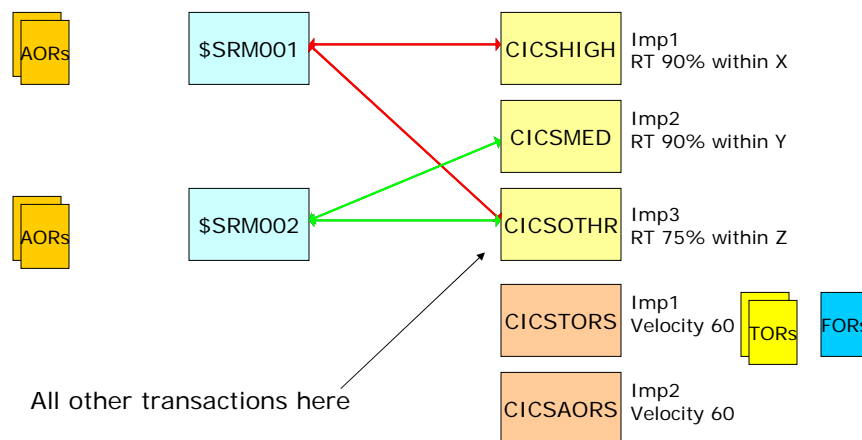
Manage Regions Using Goals of....

☐ ' BOTH' classification rule control

- ■ Only applies to CICS and IMS regions

- ■ Causes region's goal and importance level to be honored during transaction management
  - ☐ Otherwise both importance level and goal are ignored

```
Modify Rules for the Subsystem Type       Row 1 to 2 of Command
 ===> _____ SCROLL ===> PAG

  Subsystem Type . : STC        Fold qualifier names?   Y  (Y or N)
  Description  . . . IBM-defined subsystem type
  Action codes:  A=After    C=Copy        M=Move     I=Insert rule
                 B=Before   D=Delete row  R=Repeat   IS=Insert Sub-rule

            -------Qualifier---------    ----Class-------   Storage   Manage Region
   Action    Type      Name    Start     Service  Report   Critical  Using Goals Of
                                Default:  STCLOW
   ____   1  TNG       CICSTEST ___       ONLTEST  _____    NO     REGION
   ____   1  TN        CICST*   ___       CICSTORS _____    YES    BOTH
   ____   1  TN        CICSA*   ___       CICSAORS _____    YES
```

Scenario 5: Changes to WLM / Problems / Tuning

☐ All is running OK in goal mode
- ■ Some goals are missed due to high transaction response times or poor velocities
- ■ There appears to be enough available resources
- ■ WLM appears to be doing the best it can

☐ Typical causes
- ■ WLM APARs and problems
- ■ Looping jobs or subsystem problems
  - ☐ Cause performance degradation
- ■ Reduction in processor capacity
  - ☐ Decrease in number of available processors
  - ☐ Reduction in processor due to LPAR weights (i.e. short engines)
- ■ Improper tuning of non-WLM resources or facilities
  - ☐ Example: Improperly tuned XCF, CF, CICS, IMS, DB2, etc.. Causes elongated transaction response times
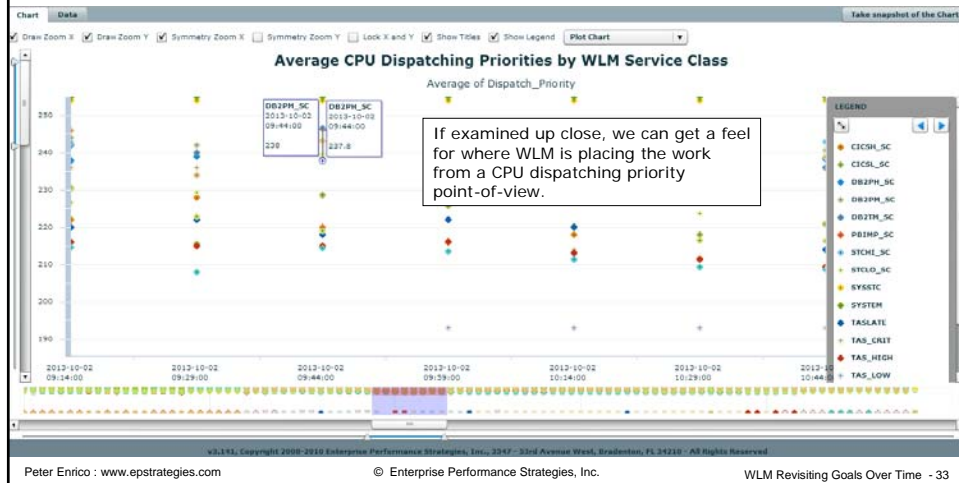
Example: Examine CPU Dispatching Priorities

- ☐ If the CPU dispatch priority is high, and the workload has the CPU and storage it needs….
  - ■ Then WLM is probably doing the best it can and the workload needs to be tuned

If examined up close, we can get a feel for where WLM is placing the work from a CPU dispatching priority point-of-view.

Scenario 6: Changes to Business Priorities and Objectives

- ☐ All is running great in goal mode
  - ■ But the installation's priorities and objectives are changing
  - ■ These changes may need to be reflected in the WLM service definition

- ☐ Typical reasons workload priorities and objectives change
  - ■ The business objectives change
  - ■ Merging of two companies / data centers
    - ☐ And you thought moving from compat mode to goal mode was tough?
  - ■ System or data center consolidation
    - ☐ Consolidating workloads from multiple images to fewer images
    - ☐ You need to reconsider both goals and assigned importance
  - ■ Introduction of new workload
  - ■ Server consolidation
  - ■ Clean up of Service Definition
    - ☐ Many service definition I look at are woefully under commented

Exercise Overview

- Most performance analysts modify their WLM service definition to
    - Influence the performance of their systems
    - Influence the performance of their workloads
    - Take advantage of new WLM functions
    - Influence WLM algorithms to manage goals and resources

- However, there are many changes you can make to your WLM Service Definition which will have little to no influence to any of these

- Growing up, you were always told to clean your room

- Now that you are an adult, you need to clean up your WLM Service Definition
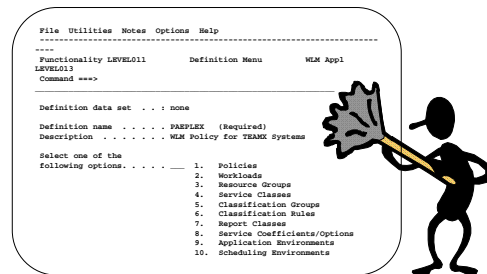
```
File  Utilities  Notes  Options  Help
-------------------------------------------------------------------------
----
   Functionality LEVEL011        Definition Menu        WLM Appl
LEVEL013
   Command ===>

Definition data set  . . : none

Definition name . . . . . PAEPLEX   (Required)
Description  . . . . . . . WLM Policy for TEAMX Systems

Select one of the
following options. . . . . __  1.   Policies
                               2.   Workloads
                               3.   Resource Groups
                               4.   Service Classes
                               5.   Classification Groups
                               6.   Classification Rules
                               7.   Report Classes
                               8.   Service Coefficients/Options
                               9.   Application Environments
                              10.   Scheduling Environments
```

Converting Your Service Definition to HTML

- When cleaning up your service definition it may help to have it in an easy to read and analyze format

- Convert your WLM service definition to HTML for readability and reference

    1. Using the WLM ISPF application: `GML`   `HTML`
        - 'File' option at top of screen
        - Sub option 'Print as GML'
        - Creates a flat file of service definition with GML formatting tags

    2. Download GML version of WLM Service Definition to workstation as text file

    2. Go to www.epstrategies.com and select WLM Tool button to convert

    5. Follow instructions - select file and fill in email address

    6. Presto! HTML file will be emailed to you within minutes

Scenario 7: Inaccuracy of Reported Measurements

- All is running great in goal mode
  - But performance measurements indicate that goals are being missed
  - Could the measurements be wrong?

- Typical cases where the reported measurements are inaccurate
  - WLM CICS and IMS server management support
    - All (or some) regions are being managed to transaction response time goals
    - But performance reports report PI as if regions are being managed to velocity goal
  - Service class contains a mixture of regions
    - Those being managed by transaction goals
    - Those being managed by velocity goals
    - Performance reports just report on service class period with no regard to this
  - Report classes contain mixture of work with unlike goals
    - Has no effect on WLM management and goals
    - But more complete / accurate data when report classes are homogeneous
  - Running scalable and non-scalable web servers together

Keep Service Classes Homogeneous

- Keep work in each service class relatively homogeneous
  - WLM takes different actions for different types of work to meet goals
  - Reports may not always reflect reality
  - Example: Don't mix CICS transactions in same service class as TSO

- Separate unlike work
  - Don't mix enclave work with non-enclave work
  - Don't mix interactive work with non-interactive work
  - Don't mix participants with non-participants
  - Don't mix server with non-servers
  - Don't mix regions managed towards region goal with regions managed towards transaction goals
  - Don't mix Batch in WLM inits with batch in JES inits
  - Don't assign goals to spaces that should truly be in SYSTEM and SYSSTC
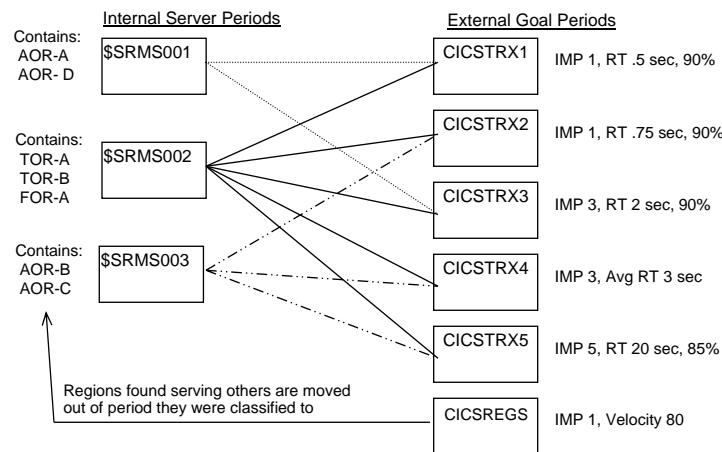 - Don't put stuff into SYSTEM and SYSSTC that should not be there

WLM Server Topology CICS Example

□ Server Topology - CICS Example
  ■ Performance monitors report achieved velocity and PI of CICSREGS
  ■ However, no regions are running in CICSREGS, they are running in $SRMxxxx

Internal Server Periods   External Goal Periods

Contains:
AOR-A
AOR- D $SRMS001   CICSTRX1   IMP 1, RT .5 sec, 90%

                              CICSTRX2   IMP 1, RT .75 sec, 90%
Contains:
TOR-A
TOR-B $SRMS002   CICSTRX3   IMP 3, RT 2 sec, 90%
FOR-A

                              CICSTRX4   IMP 3, Avg RT 3 sec
Contains:
AOR-B $SRMS003
AOR-C   CICSTRX5   IMP 5, RT 20 sec, 85%

Regions found serving others are moved
out of period they were classified to   CICSREGS   IMP 1, Velocity 80

Example of junk measurements : CICS and IMS Report

```
                              W O R K L O A D   A C T I V I T Y
                                                                              PAGE  16
     z/OS V1R12          SYSPLEX EPSPLEX          DATE 07/09/2012      INTERVAL 09.59.886   MODE = GOAL
                         RPT VERSION V1R12 RMF    TIME 09.20.00

                         POLICY ACTIVATION DATE/TIME 06/28/2012 06.15.08
-------------------------------------------------------------------------------------- SERVICE CLASS PERIODS

REPORT BY: POLICY=WLMEPS      WORKLOAD=CICS      SERVICE CLASS=CICSHIGH RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=2
                                                CRITICAL      =NONE

  -TRANSACTIONS-   TRANS-TIME HHH.MM.SS.TTT
AVG      0.00    ACTUAL          48
MPL      0.00    EXECUTION        0
ENDED    3841    QUEUED           0
END/S    6.40    R/S AFFIN        0
#SWAPS      0    INELIGIBLE       0
EXCTD       0    CONVERSION       0
AVG ENC  0.00    STD DEV        293
REM ENC  0.00
MS ENC   0.00

          RESP --------------------------- STATE SAMPLES BREAKDOWN (%) --------------------------- ------STATE------
SUB    P  TIME --ACTIVE-- READY IDLE ----------------------------WAITING FOR---------------------------- SWITCHED SAMPL(%)
TYPE      (%)  SUB  APPL               I/O PROD MISC LOCK CONV                                           LOCAL SYSPL REMOT
CICS BTE 20.1K 0.2  0.0   0.0  0.0   0.0 62.2 37.4  0.0  0.2                                              0.2   0.0   0.0
CICS EXE 88.8 11.3  0.0   0.2 85.6   2.9  0.0  0.0  0.0  0.0                                              3.1   0.0   0.0
DB2  BTE 11.3  8.4  0.0   0.0  0.0  56.6  0.0  9.6 25.3  0.0                                              0.0   0.0   0.0
DB2  EXE  3.8 17.9  0.0   0.0  0.0  75.0  0.0  3.6  3.6  0.0                                              0.0   0.0   0.0

GOAL: RESPONSE TIME 000.00.01.000 FOR  95%

         RESPONSE TIME EX   PERF
SYSTEM     ACTUAL%  VEL% INDX

SYS1        100     N/A  0.5

                         ----------RESPONSE TIME DISTRIBUTION----------
    ----TIME----     --NUMBER OF TRANSACTIONS--  -------PERCENT------  0   10  20  30  40  50  60  70  80  90  100
    HH.MM.SS.TTT     CUM TOTAL       IN BUCKET   CUM TOTAL   IN BUCKET |....|....|....|....|....|....|....|....|....|....|
<  00.00.00.500        3812            3812        99.2        99.2    >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
<= 00.00.00.600        3821               9        99.5         0.2    >
```

Good data, but no resource data

State Data usually garbage and useless... Ignore!

Valid Response Time Distribution Data

Scenario 8: Exploitation of non-WLM Functions That Influence Performance

- ☐ All is running great in goal mode
  - ■ But our installation is planning on taking advantage of system facilities that will affect the performance of the workloads
  - ■ How do we manage what may be conflicting objectives?

- ☐ New functions waiting to be exploited
  - ■ HiperDispatch

  - ■ Capacity On Demand (COD)

  - ■ Intelligent Resource Director (IRD) (pre-HiperDispatch)
    - ☐ What will be the affect on goals when WLM dynamically manages LPAR weights and number of logical processors?

  - ■ Workload License Charges (WLC)
    - ☐ What will be the affect on goals when product is capped due to exceeding defined capacity?

 - ■ How will HiperDispatch, WLC and COD all work together?

Vertical versus Horizontal CPU Management (HiperDispatch off, HiperDispatch on)

- ☐ Vertical

- ☐ Horizontal

Arranged into High, med, low Pools by PR/SM, and affinity nodes by z/OS
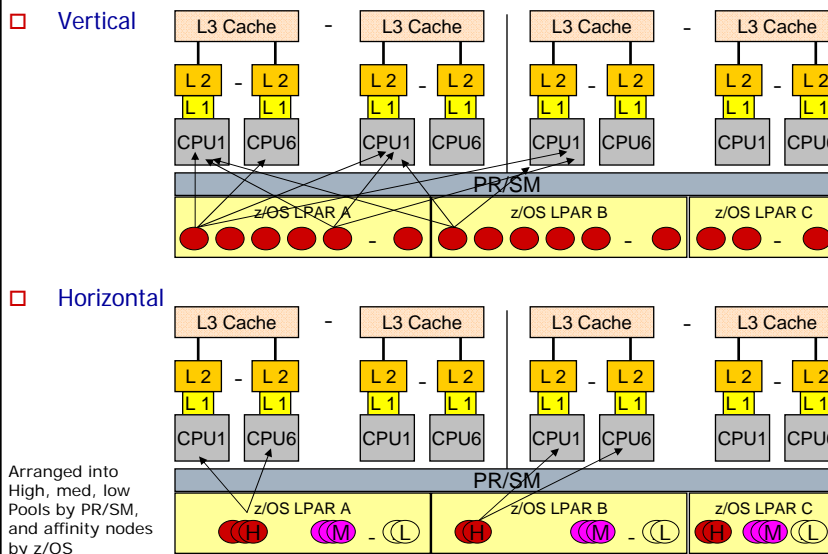
Example: WLC - Will Affect Goals

☐ Workload License Charges (WLC)
  ■ What will be the affect on goals when product is capped due to exceeding defined capacity?
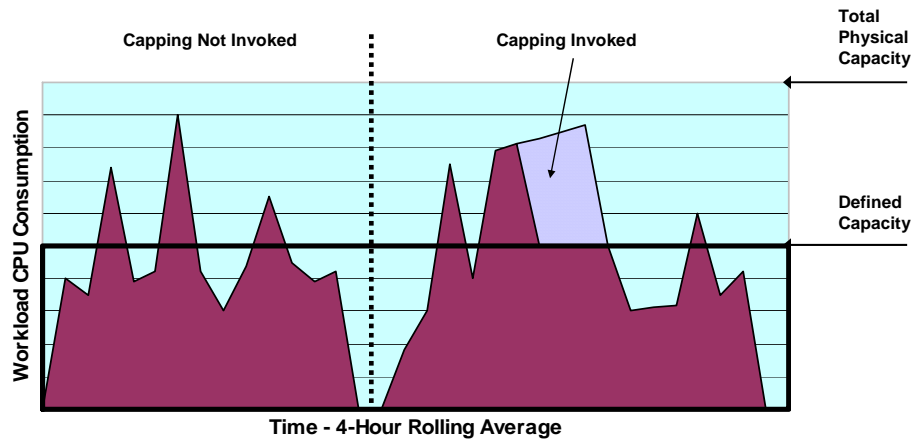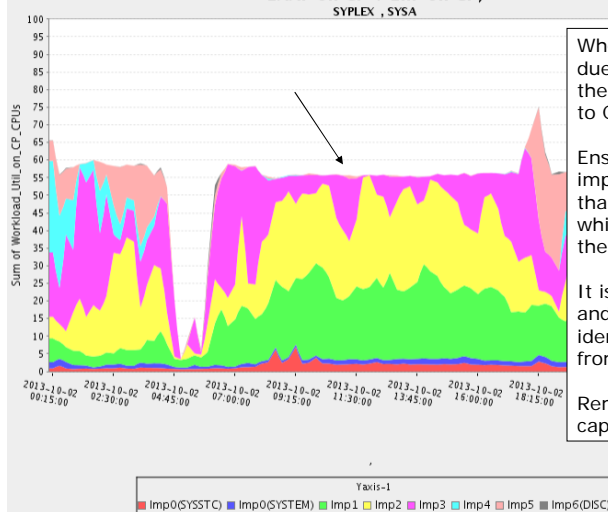
Only classify to SYSTEM and SYSSTC work that needs to be there

When CPU to an LPAR is limited due to either soft or hard capping the workloads will have less access to CPU capacity.

Ensure that your goals and importance levels are setup such that you have identified to WLM which workloads should be affected the most.

It is important to use discretionary and importance 5 to help clearly identify to WLM work to be stolen from first.

Remember, even during periods of capping, WLM tries to meet goals.

Scenario 9: Occasionally 'Something Happens'

- All is running great in goal mode
  - But every once in a while 'something happens' that causes
    - Goals to be missed
    - Workloads or system to act 'funky'
  - Most times all eventually goes back to normal
    - Some situations are repeatable, others are not
  - What is going on?

- When this usually occurs
  - Peak periods
  - Startup of a new workload or big batch job
  - Influx of workload activity
  - System problems / system dumping / subsystem ABENDs / Application problems
  - Storage shortages
  - ?????

Policy Adjustment Actions - CPU

- Dispatching priority adjustments
  - Objective: Increase Receiver's CPU using, or decrease Receiver's CPU delay
  - Interesting concepts:
    - Wait-to-Using ratio - ratio of CPU delay samples to CPU using samples (change in ratio used to determine change in CPU delay)
    - Maximum demand
      - Theoretical maximum percentage of total processor time a period can consume if it had no CPU delay
    - Achievable maximum demand
      - Percentage of total processor time a service period is projected to consume, taking into account demand of all higher work

  - Some possible actions
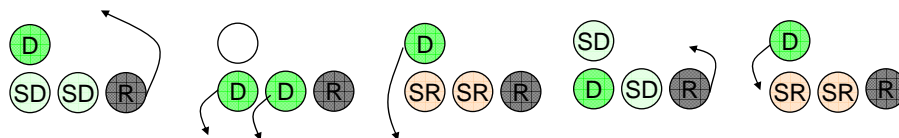
Summary

- There are many reason to revisit and reevaluate goals and WLM settings
  - Scenario 1: Improperly set goals
  - Scenario 2: Lately, goals are regularly being missed
  - Scenario 3: Planning for environmental changes
  - Scenario 4: Exploitation of additional WLM functions
  - Scenario 5: Changes to WLM, system problem, improper tuning
  - Scenario 6: Changes to business priorities and objectives
  - Scenario 7: Inaccuracy of reported measurements
  - Scenario 8: Exploitation of non-WLM functions that influence performance
  - Scenario 9: Occasionally 'something happens'

- Although some of the analysis is standard, the complete analysis for each is unique to the situation at hand
  - Requires analysis of performance measurement
  - An understanding of the sysplex, system, and workloads
  - An understanding of WLM

WLM Reports Processing/Discussion Offer !!!

- Special Reports Offer!
  - See your WLM SMF 72.3 records in chart and table format

  - Please contact me, Peter Enrico for instructions for sending raw SMF data
    - Send an email to peter.enrico@epstrategies.com

  - Deliverable: Dozens of coupling facility based reports (charts and tables)
    - WLM SMF30 Address Space Analysis
    - WLM - Period Setup Analysis
    - WLM - Importance Level Analysis
    - WLM - PI Analysis
    - WLM - Velocity Goal Analysis
    - WLM - Response Time Goal Analysis
    - WLM - Discretionary Goal Analysis
    - WLM - Enclave Analysis
    - WLM - Multiple Period Analysis
    - WLM - CPU Analysis
    - WLM - Storage Analysis
    - WLM - DASD IO Analysis
  - One-on-one phone call to explain your coupling facility measurements

Performance Workshops Available

During these workshops you will be analyzing your own data!

- ☐ WLM Performance and Re-evaluating of Goals
    - ☐ Instructor: Peter Enrico
    - ☐ September 15 – 19, 2014   - Kansas City, Missouri, USA
    - ☐ October 20 - 24, 2014   - Munich, Germany

- ☐ Parallel Sysplex and z/OS Performance Tuning
  (Web / Internet Based!)
    - ☐ Instructor: Peter Enrico
    - ☐ July 29 – 31, 2014 (Web)
    - ☐ Considering a fall class...

- ☐ Essential z/OS Performance Tuning Workshop   www.pivotor.com
    - ☐ Instructors: Peter Enrico and Tom Beretvas

- ☐ z/OS Capacity Planning and Performance Analysis
    - ■ Instructor: Ray Wicks