



z/OS Workload Manager Revisiting Goals Over Time

(2023 Edition)

Peter Enrico



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



Creators of Pivotor®

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



z/OS Workload Manager Revisiting Goals Over Time

(2023 Edition)

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



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



Creators of Pivotor®



Contact, Copyright, and Trademarks



Questions?

Send email to performance.questions@EPStrategies.com, or visit our website at <https://www.epstrategies.com> or <http://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

Abstract



- There are many reasons to ‘open up’ a WLM service definition. Some changes need to be made for tuning; others for clean-up; and still others because the goals are out date. Currency examination needs to be performed for planned environmental changes, or changes to business priorities. The point is that there are many reasons to open up the WLM service definition for examination and potential changes.
- During this presentation Peter Enrico will review the different reasons to roll up your sleeve and analyze your WLM service definition. During this presentation you are sure to learn a lot about a wide variety of areas related to WLM management on z/OS.

EPS: We do z/OS performance...



- We are z/OS performance!
- Pivotor
 - Performance reporting and analysis of your z/OS measurements
 - Example: SMF, DCOLLECT, other, etc.
 - Not just reporting, but cost-effective analysis-based reporting based on our expertise
- Performance Educational Workshops (while analyzing your own data)
 - Essential z/OS Performance Tuning
 - Parallel Sysplex and z/OS Performance Tuning
 - WLM Performance and Re-evaluating Goals
- Performance War Rooms
 - Concentrated, highly productive group discussions and analysis
- MSU reductions
 - Application and MSU reduction

z/OS Performance workshops available



During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
 - March 20-24, 2023
- Parallel Sysplex and z/OS Performance Tuning
 - May 2-3, 2023
- WLM Performance and Re-evaluating Goals
 - October 2-6, 2023
- Also... please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Like what you see?



- Free z/OS Performance Educational webinars!
 - The titles for our Fall 2022-2023 webinars are as follows:
 - ✓ *Key Reports to Evaluate z16 Processor Caches*
 - ✓ *Understanding System Recovery Boost's Impact on Performance and Performance Reporting*
 - ✓ *WLM Management of DDF Work: What can you do and what has changed?*
 - ✓ *Intensity! Understanding the Concepts and Usage of Intensity Measurements*
 - ✓ *High, Medium, Low: Understanding how HiperDispatch influences performance in z/OS*
 - ✓ *How and why Pivotor is different than other performance management reporters*
 - *Putting a lid on XCF*
 - *Key Reports to Evaluate Usage of Parallel Access Volumes*
 - *Key Reports to Evaluate Coupling Facility CPU Utilization*
 - *Understanding how memory management has evolved in z/OS*
 - Let me know if you want to be on our mailing list for these webinars
- If you want a free cursory review of your environment, let us know!
 - We're always happy to process a day's worth of data and show you the results
 - See also: <http://pivotor.com/cursoryReview.html>

EPS presentations this week

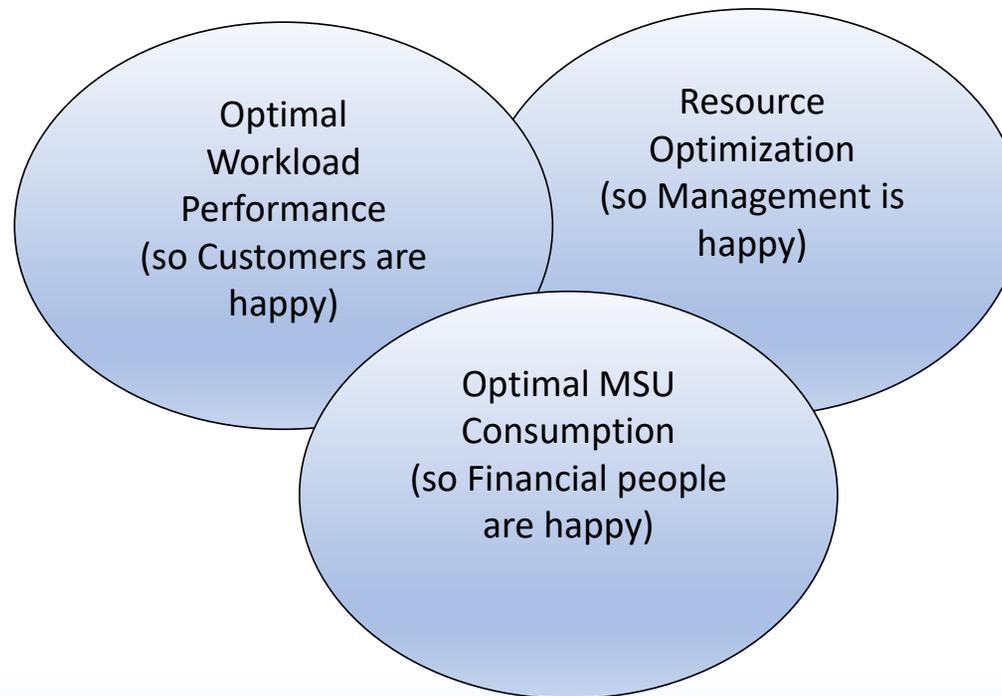


What	Who	When	Where
PSP: z/OS Performance Tuning – Some Top Things You May Not Know	Peter Enrico Scott Chapman	Tue 13:15	Strand 12A
z/OS WLM – Revisiting Goals Over Time	Peter Enrico	Mon 16:00	Empire C
Sharing CPUs: How z/OS & PR/SM Manage Logical & Physical Processors	Scott Chapman	Wed 08:00	Empire C
Observability Shootout	Scott & other ISVs	Wed 16:00	Empire C
I/O, I/O It's Home to Memory We (Should) Go	Scott Chapman	Fri 09:15	Strand 12A

The Performance Balancing Act



- Performance on z/OS is about finding an optimal balance of the following:



Presentation Overview



- What does it mean to ‘revisit’ and ‘reevaluate’ goals?
- Reasons to Re-evaluate Goals have not changed much in last 27 years
 - 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: Use Percentile Response Time goals over Average Response Time goals



- Average Response Time Goal:

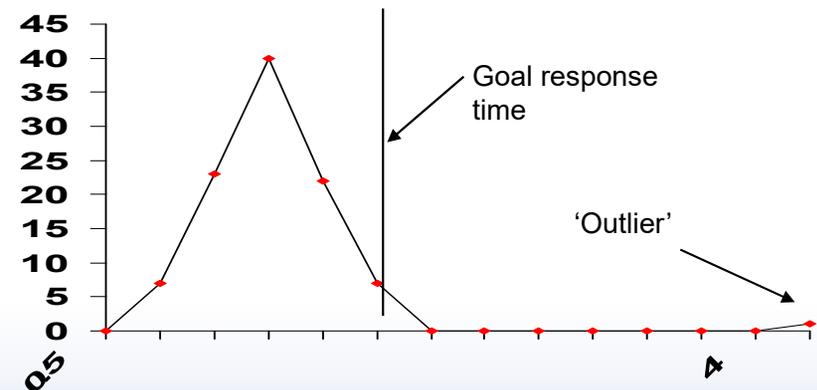
- The average response time desired for a given set of ended transactions

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

- Average response time goals are not recommended since they can be easily influenced by 'outlier' transactions

- Average response time goal set to 1 second
- 99 transactions complete in 1 second, but 1 transaction completes in 2 minutes
- Average RT achieved is 2.2 seconds
- Goal missed!
 - even though 99% of transactions completed within 1 second

Number of Transactions Completing in Each RT bucket

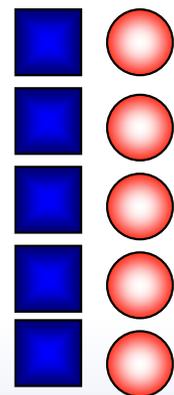


Example: Your velocity goals may be too aggressive



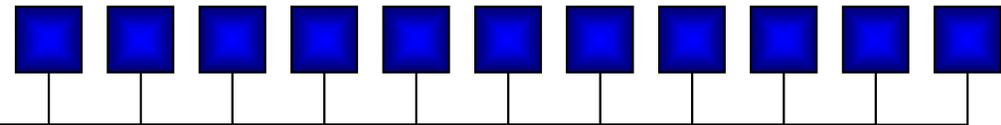
- Periods with more address spaces or enclaves (dispatchable units) tend to have more CPU queuing
 - Since delays are inherent such environments usually result in lower velocities
 - Example: Have a period with 100 active IMS MPRs
 - All want CPU but we are only running on a 5-way processor
 - Velocity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

Dispatched Work
• Accumulating CPU Using Samples



Processors

Queued Work - waiting at priority
• Accumulating CPU delay samples



Dispatcher Queue

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



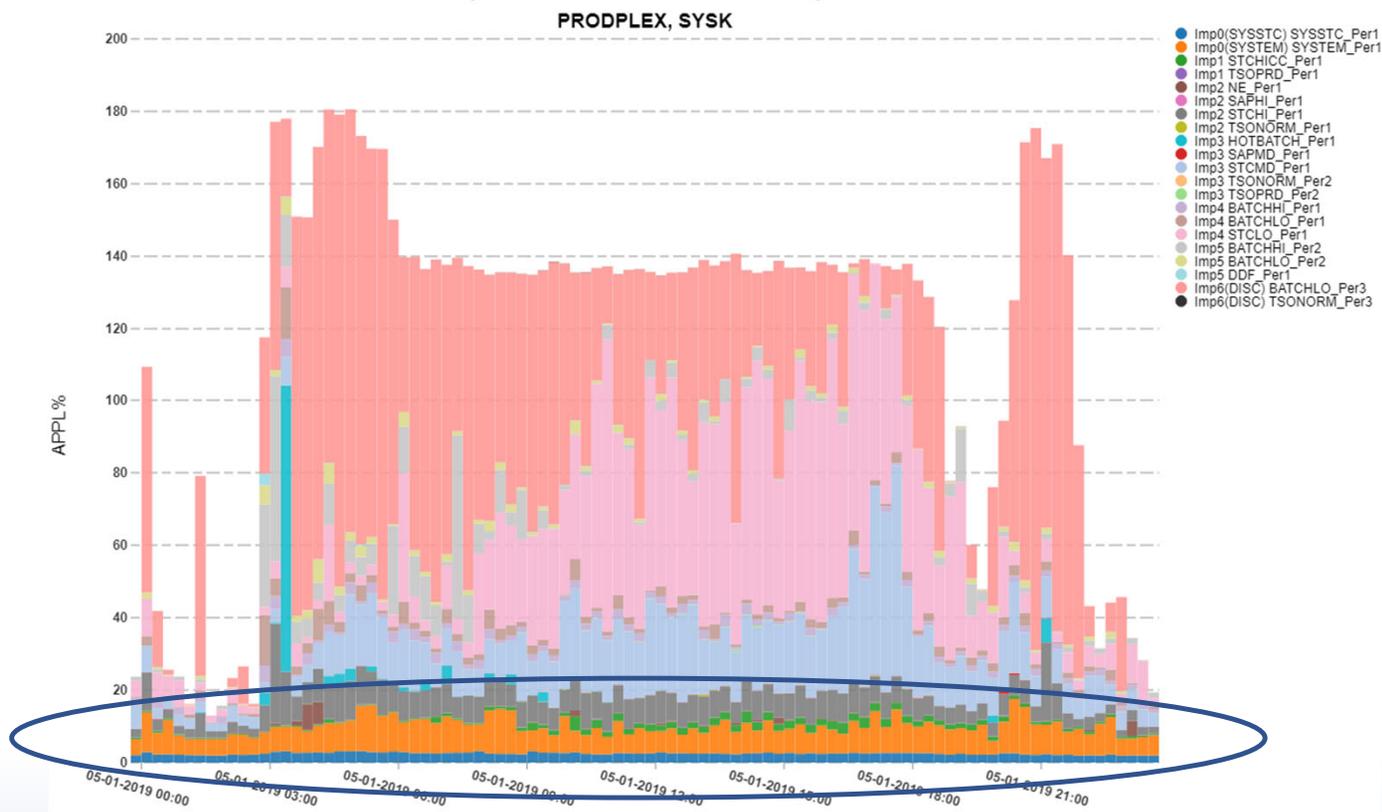
- Too much work in SYSSTC allows less CPU to goal periods
 - Remember:
CPU Available to goal periods and discretionary is CPU not consumed by SYSTEM and SYSSTC
- Too little work in discretionary may not provide donor resources
 - This is OK, but remember: WLM first looks for free resources. If not enough free resources than takes first from discretionary before taking from goal periods

255	SYSTEM
254	SYSSTC
253	'Unused' (SYSSTC1-5)
249	
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

Examine CPU consumed by SYSTEM and SYSSTC



WLM Trx CPU - CP APPL% by Service Class
(Transaction CP + zIIP on CP)



Scenario 2: Caused by Workload Reduction



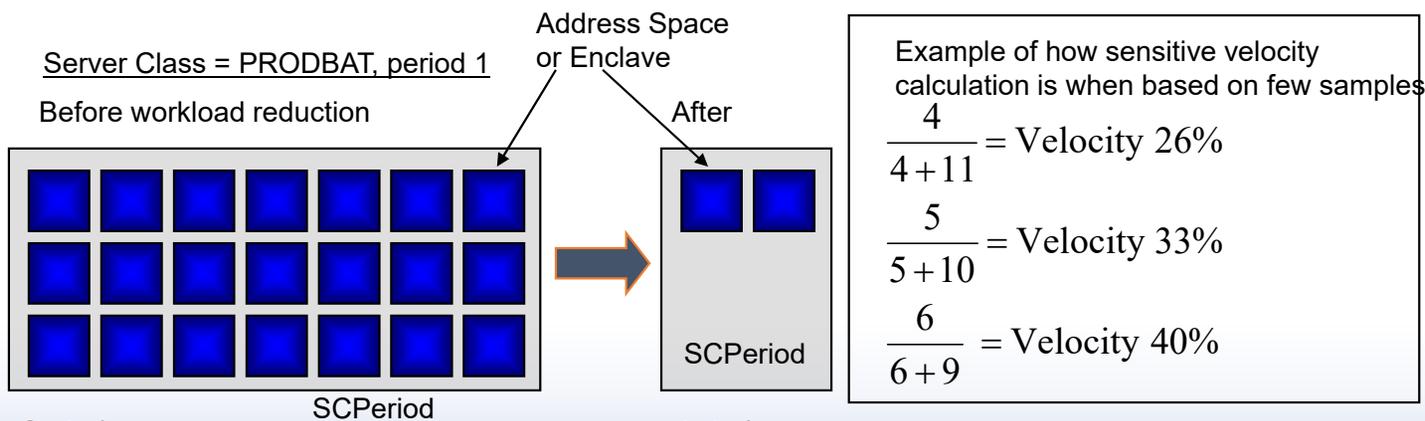
- If a workload shrinks it is possible for goals to be missed
 - Fewer using and delay samples calculate a less meaningful velocity
 - Fewer ended transactions calculate a less meaningful response time
 - Less discretionary work could mean less discretionary resource donors
- Indicators:
 - Velocity goals
 - Fewer address spaces or enclave transactions contributing samples
 - 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
 - Number of address spaces or enclaves in period will influence a period's velocity:
 - Number of using and delay samples collected
 - Fewer address spaces or enclaves mean fewer contributing samples
 - Calculated velocities (will be more sensitive)

Velocity based on $\frac{5000}{5000 + 10000}$ same as Velocity based on $\frac{5}{5 + 10} = 33\%$



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



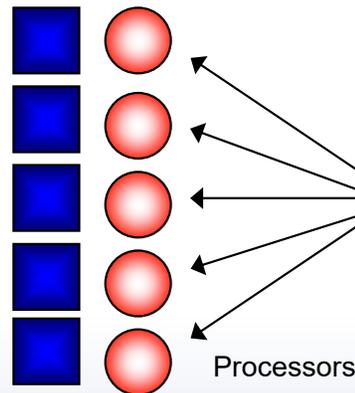
- All is running great in goal mode
 - But environmental changes are planned
 - Do we need to consider making any WLM changes?
- Typical changes to be concerned with
 - Changes in hardware
- Changes to processor environment
 - Introduction of new processors
 - Changes to speed of processors (faster or 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

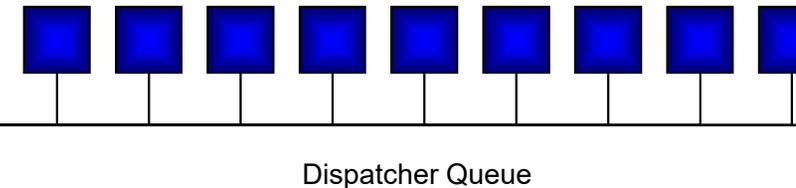


- CPU delay could be inherent to many periods
 - Few engines than work that wants to run
 - Periods with more ready address spaces than processors will have delay
 - Faster engines
 - May cause using samples of quick running work to be missed
 - Other causes of delay include
 - Reduced preemption
 -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

Dispatched Work
• Accumulating CPU Using Samples



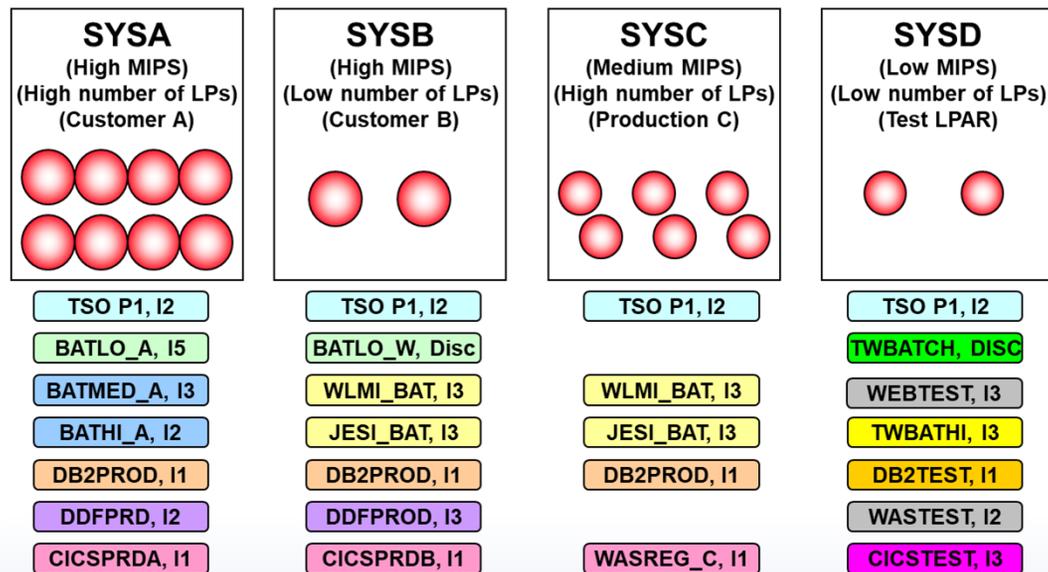
Queued Work - waiting at priority
• Accumulating CPU delay samples



Example: Remember that goals are Sysplex wide



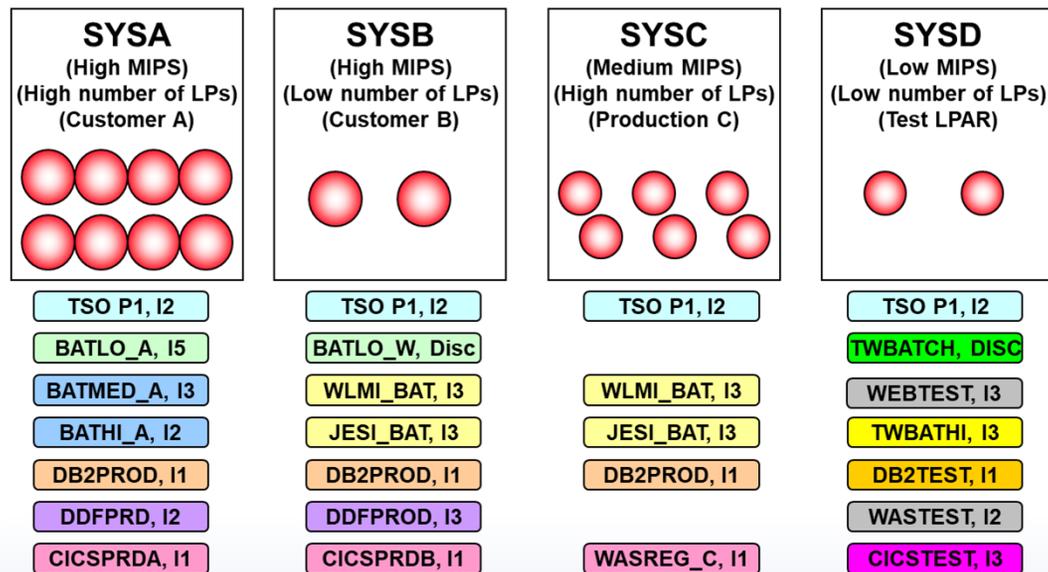
- Sysplex of a little of everything
 - Asymmetrical capacity local systems in same Sysplex where some workloads with mutually exclusive and disparate workloads
 - Note: Probably all installations have this type of Sysplex in one form or another



Example: Remember that goals are Sysplex wide



- How do you assign 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



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 (or turning off exploitation) some WLM functionality that is currently not being used
 - Do we need to consider goals when making these WLM changes?

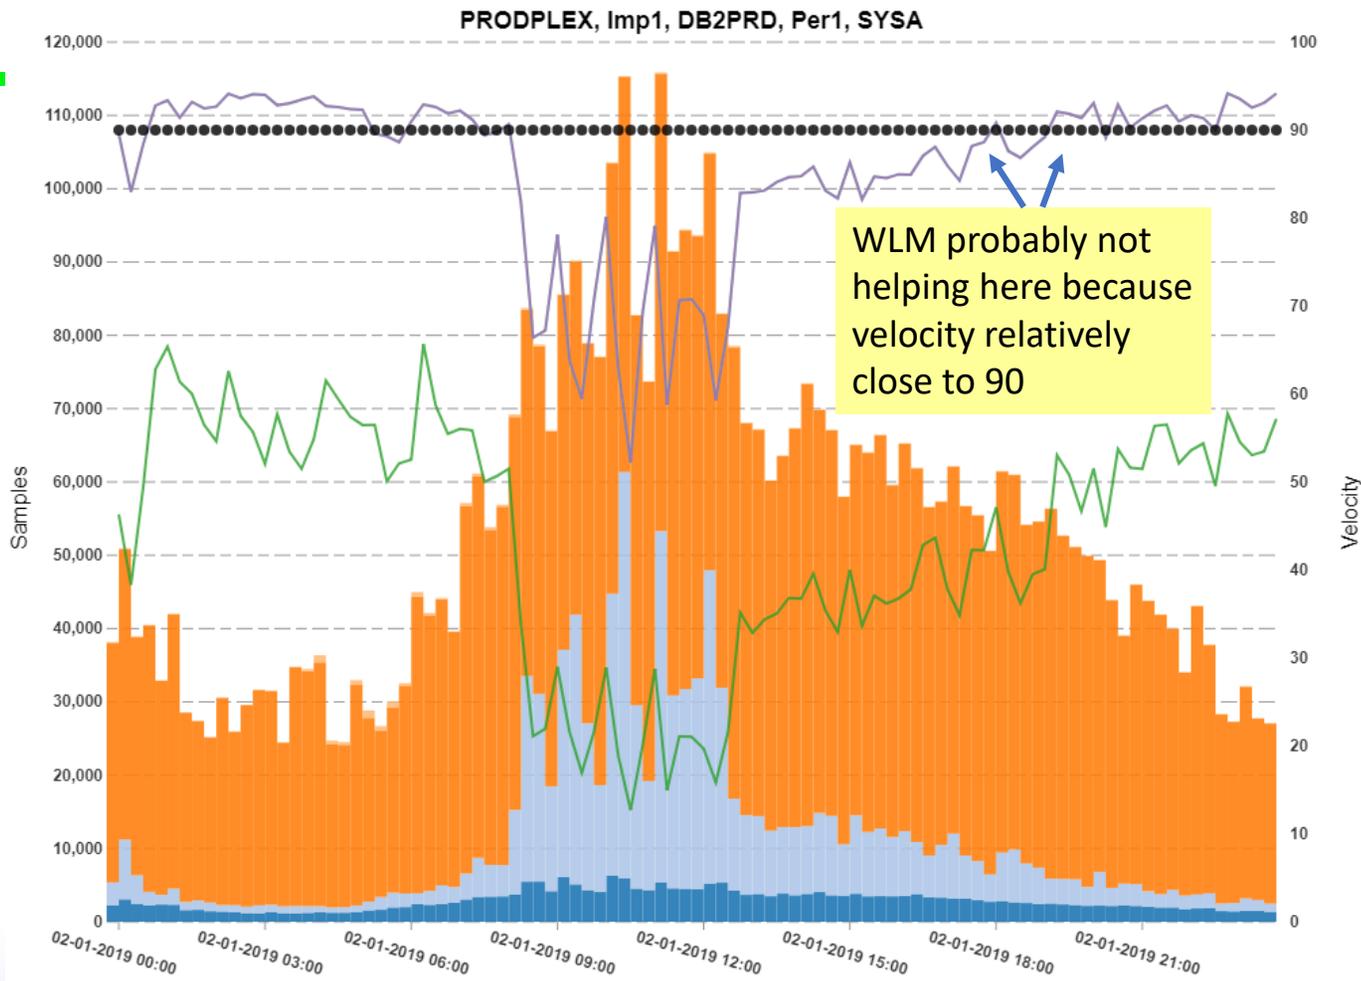
- Any functions waiting to be exploited
 - New subsystem types
 - New classification qualifiers and group options
 - WLM managed initiators
 - Scheduling environments
 - I/O priority management if it is still turned on
 - CPU critical control
 - CICS or IMS transaction management
 - Storage critical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

I/O Priority Management Introduction



- I/O Priority Management was added in 1997
- I/O queueing was a significant issue!
 - ESCON channel could only be servicing a single I/O at a time
 - Without PAVs, a volume could only be servicing a single I/O at a time
- Goal was to manage I/O priority separately from CPU priority
 - **Without I/O Priority Management, I/O priority = CPU priority**
- Higher priority I/Os could move to the front of the I/O queue
 - Help I/O-limited work without impacting the work's CPU dispatching priority
- I/O priority also passed to the DASD controller
 - So can influence queueing within the controller as well

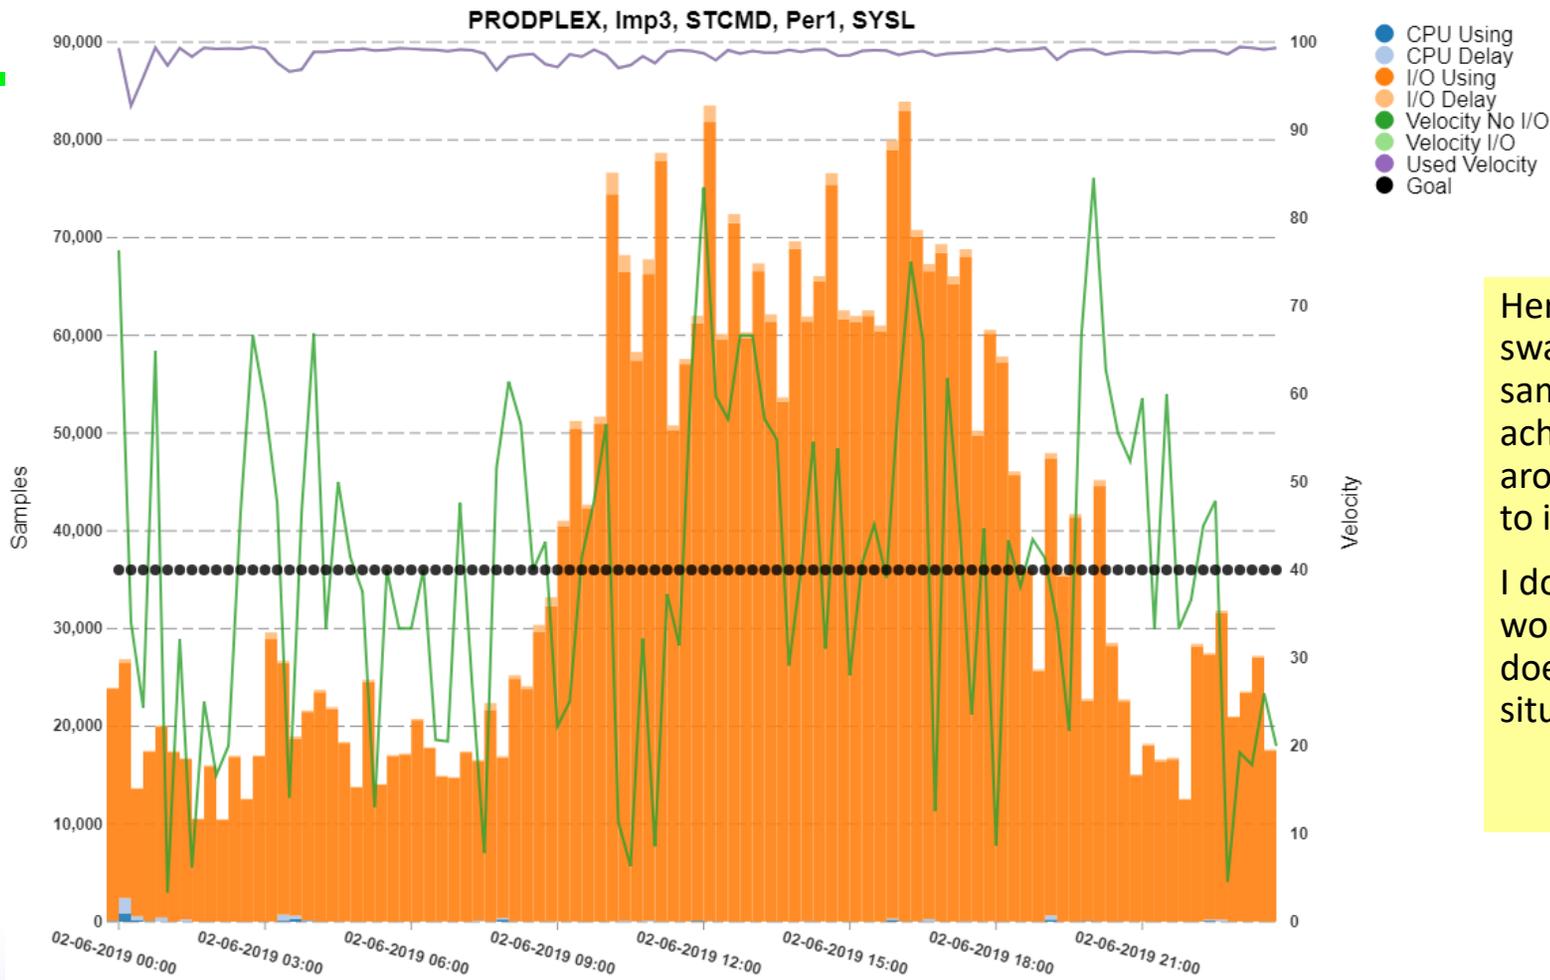
CPU & I/O Contribution to Velocity



In this case, this importance 1 workload needs to suffer significant CPU delays before those overcome the I/O using to significantly alter the velocity.

This workload might be better protected **without** I/O Priority management and with a goal of ~60 instead of 90.

CPU & I/O Contribution to Velocity



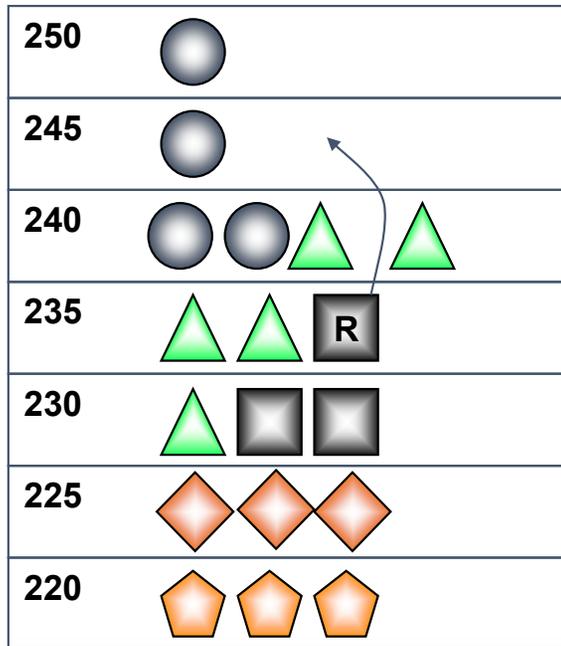
Here I/O has completely swamped STCMD's CPU samples. With an achieved velocity of around 97, WLM is sure to ignore this work.

I don't know what this work actually is, but it doesn't seem like a good situation.

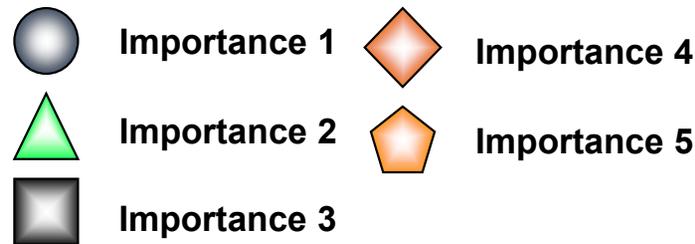
Long Term CPU Protection - Background



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



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

z/OS 3.1 WARNING!



- As of z/OS 3.1, IBM made an incompatible change to WLM
- By default, all Importance level 1 work will be managed as CPU critical
 - IBM is making a huge assumption that Importance levels are tuned
 - IBM is also assuming that all Importance 1 is has predictable behavior 24/7
 - Essentially, IBM seems to be winding down WLM management, and is starting to reintroduce fixed CPU dispatching priorities
- Keep the new default of Importance 1 as CPU critical if...
 - If you believe worked marked as importance level 1 is truly correct
 - If the work at importance level 1 is predicable in its resource consumption and workload volume

Example: Manage Region Using Goals of

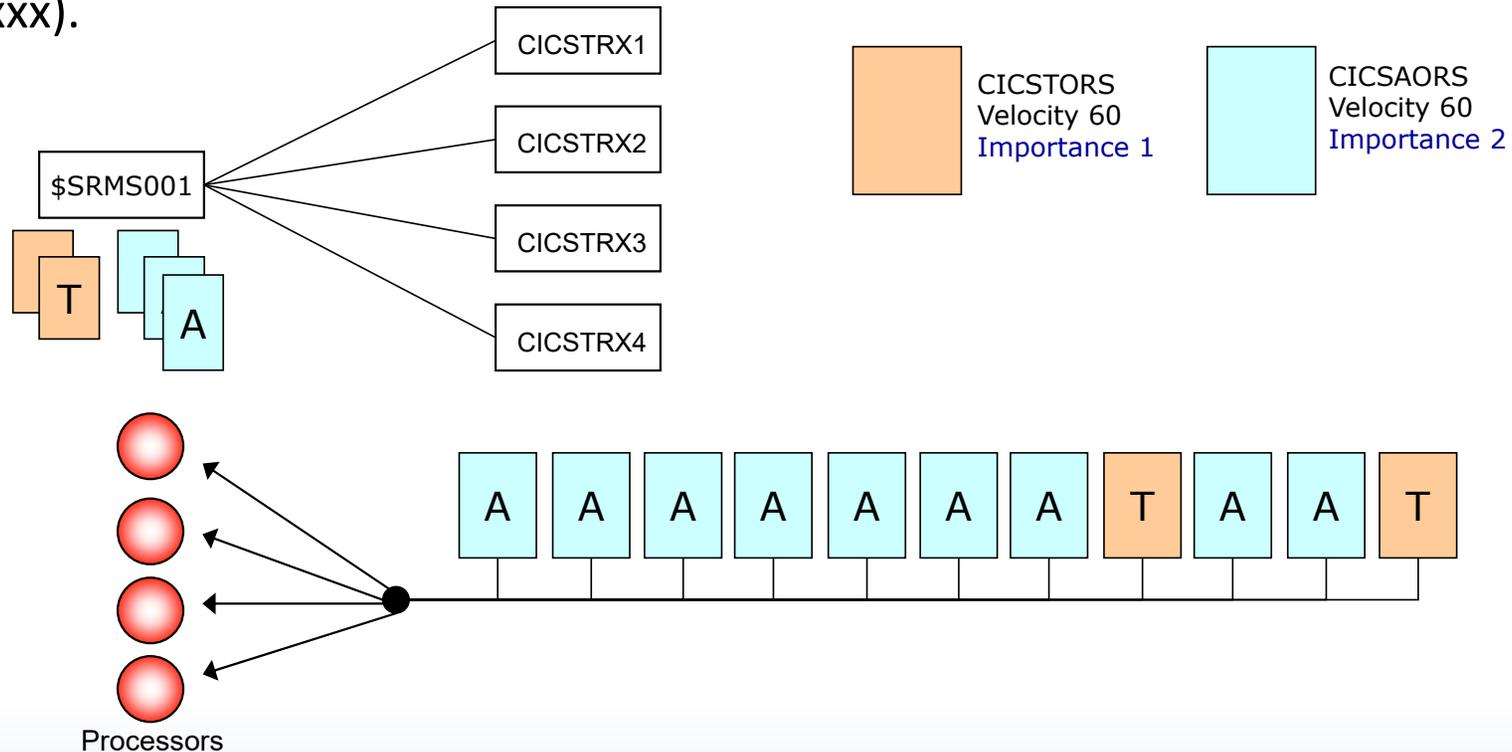


- Installation's have several choices for how WLM will manage their CICS and IMS workloads
- Region Management:
 - Allow all regions to be managed towards velocity goals and importance levels
 - Response time goals of transactions have no influence
- Transaction Management:
 - Allow all regions to be managed to meet the response time goals of the transactions they are serving.
 - The goals and importance levels of the regions are ignored.
- 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
- 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.

Additional Considerations



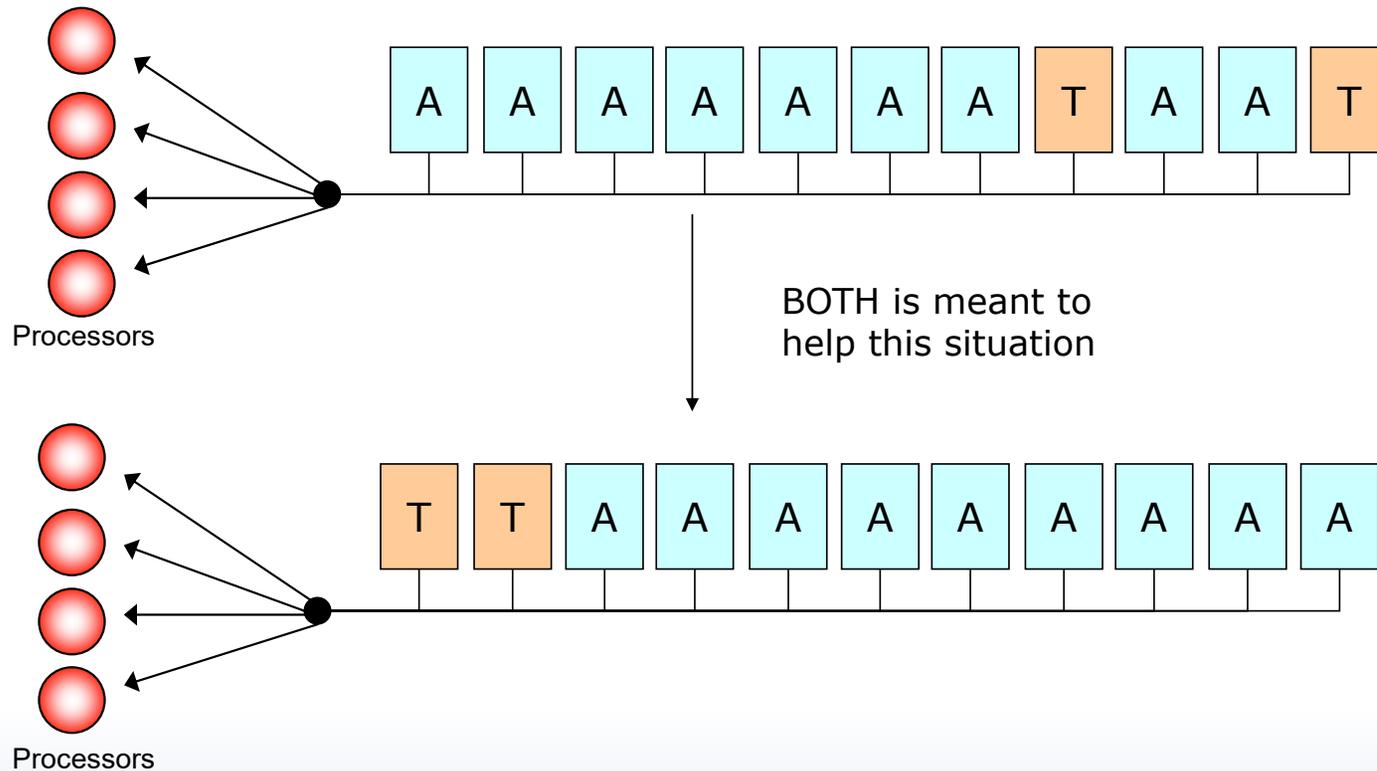
- By default, WLM will only consider the fact all regions that are processing a common set of transactions should be managed together in a common internal period (\$SMFSxxx).



Additional Considerations



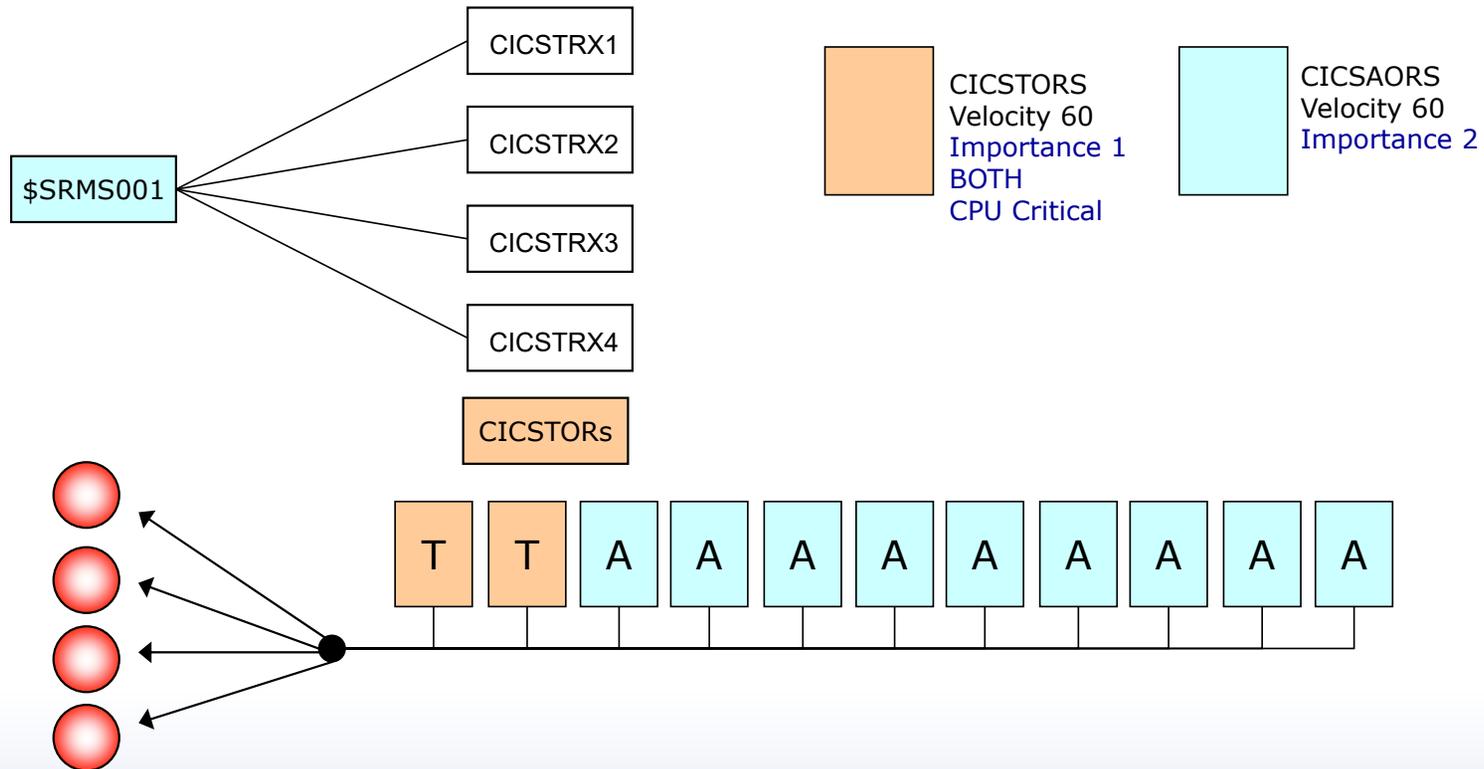
- When Importance level assigned to regions is considered, WLM will consider this importance level when managing the regions in the \$SMFSxxx periods



Additional Considerations



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



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, CPU Critical
- Service class **CICSAORS** Velocity 60, Importance 2



IMP 1, Velocity 60



IMP 2, 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-----
Action  Type      Name      Start      Class-----  Storage  Manage Region
                Default:  Service  Report  Critical  Using Goals Of
-----
_____ 1  TNG      CICSTEST  _____  ONLTST  _____  NO      REGION
_____ 1  TN       CICST*   _____  CICSTORS _____  YES     BOTH
_____ 1  TN       CICS*A   _____  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

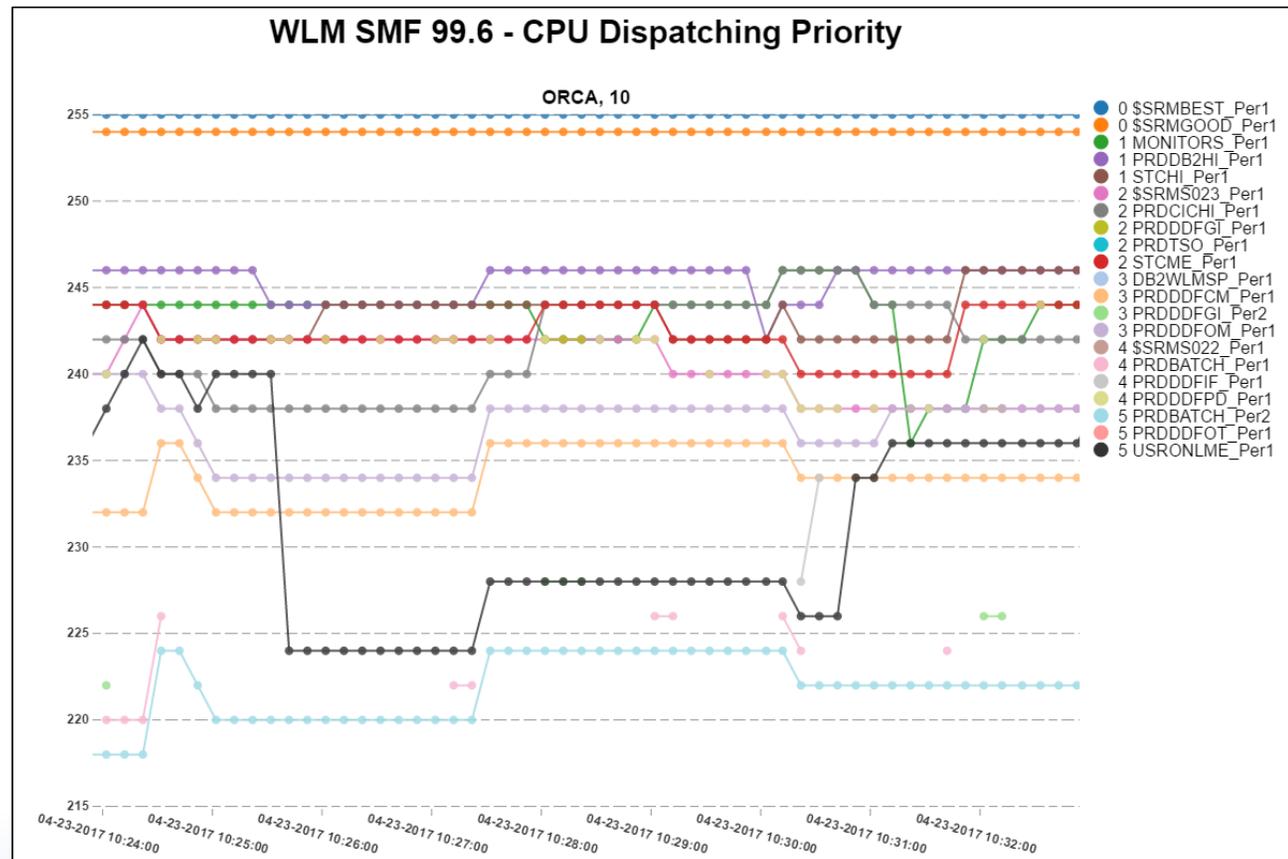
- Typical causes
 - WLM APARs and problems
 - Looping jobs or subsystem problems
 - Cause performance degradation
 - Reduction in processor capacity
 - Decrease in number of available processors
 - Reduction in processor due to LPAR weights (i.e. short engines)
 - Improper tuning of non-WLM resources 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



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
- 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

Clean Up Your Service Definition!



- 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

A screenshot of a terminal window showing the 'File Utilities Notes Options Help' menu. The window title is 'LEVEL025 Definition Menu WLM Appl LEVEL025'. The terminal text includes: 'Command ==> _____', 'Definition data set . . . : none', 'Definition name _____ (Required)', 'Description _____', and a list of 13 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, 11. Guest Platform Management Provider, 12. Tenant Resource Group, 13. Tenant Report Classes. To the right of the terminal window is a black stick figure holding a large grey leaf, symbolizing cleaning.

WLM Service Definition in XML Format



- WLM service definition can be save in two formats
 - ISPF Tables or XML format (*the XML is relatively new function*)
- Today, it is recommended to save the WLM service definition in XML format
 - The ISPF tables are automatically updated when they are touched by new APARs or z/OS releases. This then makes then ineligible to be updated if the APARs or z/OS releases are rolled back, or if an older release needs to edit or access.

- Select

- File
- -> Save as

Use Save as to save the currently displayed service definition in a PDS as ISPF tables or in a PS as XML

```
File Utilities Notes Options Help
-----
Functionality LEVEL025 Definition Menu WLM Appl LEVEL035
Command ===> _____

Definition data set . . . : none

Definition name . . . . . _____ (Required)
Description . . . . . _____

Select one of the following options.
___ 1. Policies                               12. Tenant Resource Groups
    2. Workloads                             13. Tenant Report Classes
    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
   11. Guest Platform Management Provider
```

After you save WLM Service Definition to XML file...



- The XML file will look crazy!

```
Classification>Ú<SubsystemType>IWEB</SubsystemType>Ú<Description>N/A</Descriptio
n>Ú<CreationDate>1900/01/01 00:00:00</CreationDate>Ú<ModificationDate>2000/02/17
18:04:47</ModificationDate>Ú<ModificationUser>U873</ModificationUser>Ú<DefaultS
erviceClassName>NEWWKL</DefaultServiceClassName>Ú<DefaultReportClassName>NEWIWEB
</DefaultReportClassName>Ú</Classification>Ú<Classification>Ú<SubsystemType>JES<
/SubsystemType>Ú<Description>JES2 Rules</Description>Ú<CreationDate>1900/01/01 0
0:00:00</CreationDate>Ú<ModificationDate>2011/05/16 07:12:56</ModificationDate>Ú
<ModificationUser>I014350</ModificationUser>Ú<DefaultServiceClassName>BATNORM</D
efaultServiceClassName>Ú<DefaultReportClassName>BATDEF</DefaultReportClassName>Ú
<ClassificationRules>Ú<ClassificationRule>Ú<QualifierType>UseridGroup</Qualifier
Type>Ú<QualifierValue>OPCUSER</QualifierValue>Ú<ServiceClassName>BATNORM</Servic
eClassName>Ú<ReportClassName>BATPROD</ReportClassName>Ú<StorageCritical>No</Stor
ageCritical>Ú<RegionGoal>No</RegionGoal>Ú<ClassificationRule>Ú<QualifierType>Per
form</QualifierType>Ú<QualifierValue>1</QualifierValue>Ú<ServiceClassName>BATNOR
M</ServiceClassName>Ú<ReportClassName>BATPROD</ReportClassName>Ú<StorageCritical
>No</StorageCritical>Ú<RegionGoal>No</RegionGoal>Ú<ClassificationRule>Ú<Classif
icationRule>Ú<QualifierType>Perform</QualifierType>Ú<QualifierValue>9</Qualifier
Value>Ú<ServiceClassName>BATEXT</ServiceClassName>Ú<ReportClassName>BATPROD</Rep
ortClassName>Ú<StorageCritical>No</StorageCritical>Ú<RegionGoal>No</RegionGoal>Ú
</ClassificationRule>Ú<ClassificationRule>Ú<QualifierType>TransactionName</Quali
fierType>Ú<QualifierValue>DD*</QualifierValue>Ú<ServiceClassName>BATNORM</Servic
eClassName>Ú<ReportClassName>DB2UTIL</ReportClassName>Ú<StorageCritical>No</Stor
ageCritical>Ú<RegionGoal>No</RegionGoal>Ú</ClassificationRule>Ú<ClassificationRu
le>Ú<QualifierType>TransactionName</QualifierType>Ú<QualifierValue>KS*</Qualifie
rValue>Ú<ServiceClassName>BATHIGH</ServiceClassName>Ú<ReportClassName>BATKS</Rep
ortClassName>Ú<StorageCritical>No</StorageCritical>Ú<RegionGoal>No</RegionGoal>Ú
</ClassificationRule>Ú<ClassificationRule>Ú<QualifierType>TransactionName</Quali
fierType>Ú<QualifierValue>MD*</QualifierValue>Ú<ServiceClassName>BATHIGH</Servic
eClassName>Ú<ReportClassName>BATMDHI</ReportClassName>Ú<StorageCritical>No</Stor
ageCritical>Ú<RegionGoal>No</RegionGoal>Ú</ClassificationRule>Ú<ClassificationRu
le>Ú<QualifierType>TransactionName</QualifierType>Ú<QualifierValue>M5*</Qualifie
rValue>Ú<ServiceClassName>BATHIGH</ServiceClassName>Ú<ReportClassName>BATM5</Rep
ortClassName>Ú<StorageCritical>No</StorageCritical>Ú<RegionGoal>No</RegionGoal>Ú
</ClassificationRule>Ú<ClassificationRule>Ú<QualifierType>TransactionName</Quali
```

Added Benefit!



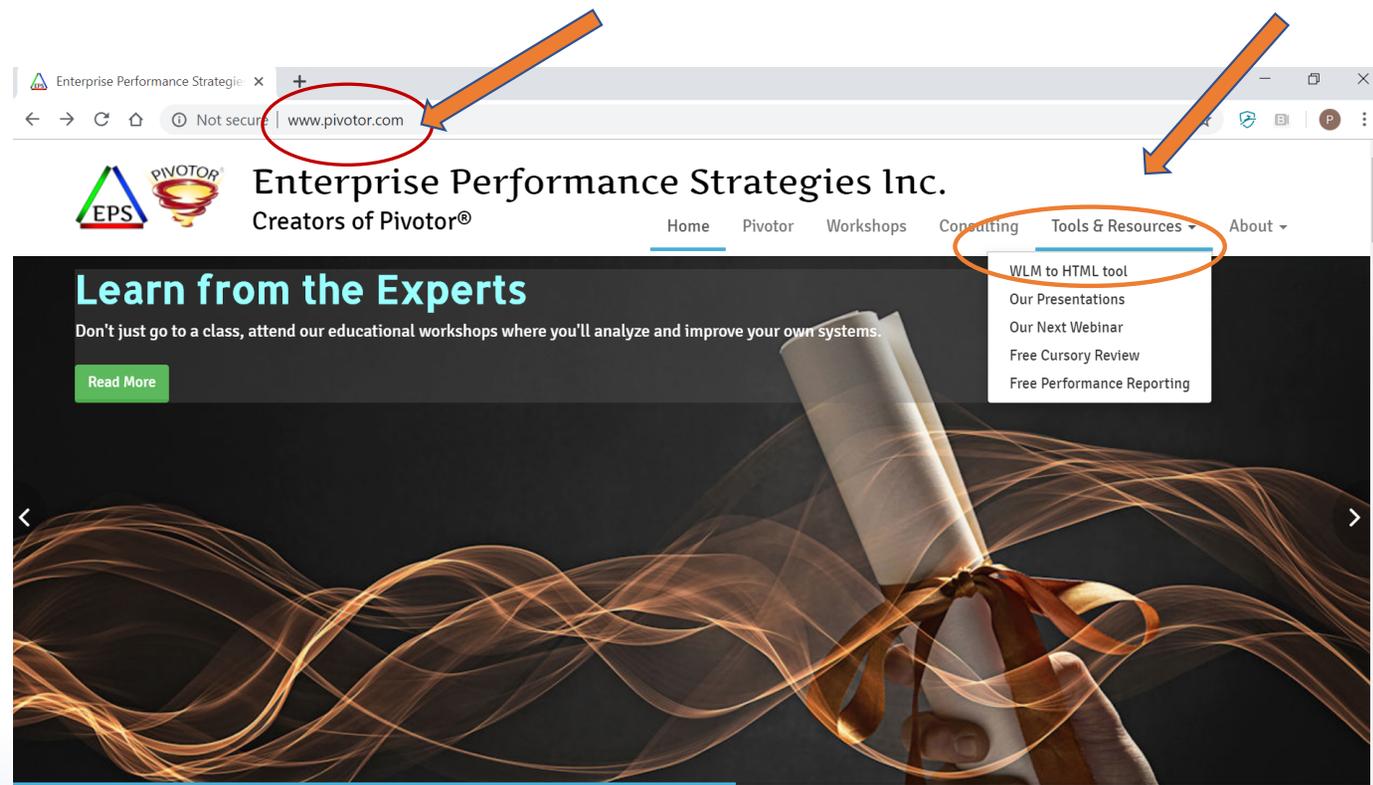
- Once saved as XML, a tool is available to nicely format the XML file into a easy to read format to assist during your WLM analysis

- Visit www.pivotor.com
or www.epstrategies.com

- Select 'Tools & Resources' option

- Select WLM TO HTML

- Provide your XML file and email address
 - HTML formatted WLM service definition emailed to you in seconds!



Example of HTML Formatted Service Definition



- Not only is the XML file nicely formatted

But there is some analysis built into the file to help you with your service definition cleanup.

[Click Here](#)

WLM Service Policy

file:///C:/Users/Peter/Enterprise%20Performance%20Strategies/EPS.Inc.Team.Site%20-%20Do...

Apps D3.js - Data-Driven D Filter: All open cases OwlCo Local Owl Pivot Reporting Se Office 365 Other bookmarks

EPS PIVOTOR ABCWLM Brought to you by Enterprise Performance Strategies at www.Pivotor.com Submit Again About Us

Introduction 13.02.2018 / 254

The service coefficients are defined as:

CPU	1.0
IOC	0.1
MSO	0.0000
SRB	1.0

The service options are:

I/O Priority Management	Yes
Dynamic Alias Management	Yes
I/O Priority Groups Enabled	No

Counts for this service definition:

Service Classes	20
Service Class Periods	23
Report Classes	330
Workloads	7
Classification Groups	0
Application Environments	63
Scheduling Environments	0
Resources	0

The first line of the notes reads:
Folgende Aenderungen wurden durchgefuehrt:

This service definition is at functionality level 029

The Proplid string is: WLM AA zOS V1 HBB7790 LEVEL031

The Replid string is: C8E4D2E8D3D44040D3E28A56A10E3000E2E8F0F4F5404040E2E8C3

Service Policy Overrides

Service Classes

Classification Rules

Classification Groups

Report Classes

Resource Groups

Application Environments

Resources

Scheduling Environments

Workloads

Notes

Subsystem - SC - RC Cross reference

SC - Subsystem Cross reference

Default Classifications

Change History

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

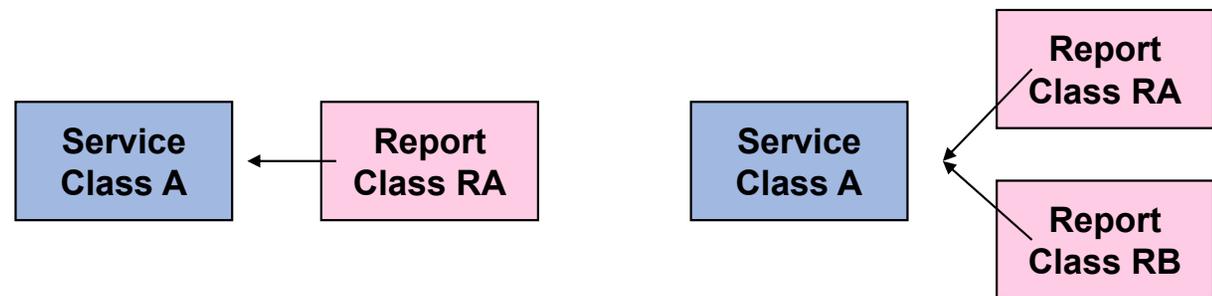
- Separate unlike work
 - Don't mix enclave work with non-enclave work
 - Don't mix interactive work with non-interactive work
 - Don't mix participants with non-participants
 - Don't mix server with non-servers
 - Don't mix regions managed towards region goal with regions managed towards transaction 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

Heterogeneous versus Homogeneous report classes

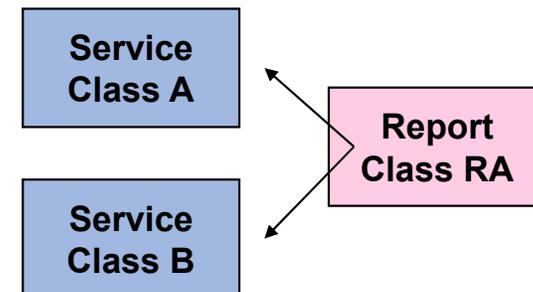


- Report classes are either heterogeneous or homogenous relative to the service class where the work in the report class is classified

- Homogenous Report Classes
- Report class is made up of work from a single service class



- Heterogeneous Report Classes
- Report class is made up of work from multiple service classes
- Even if goal of the different service class is the same

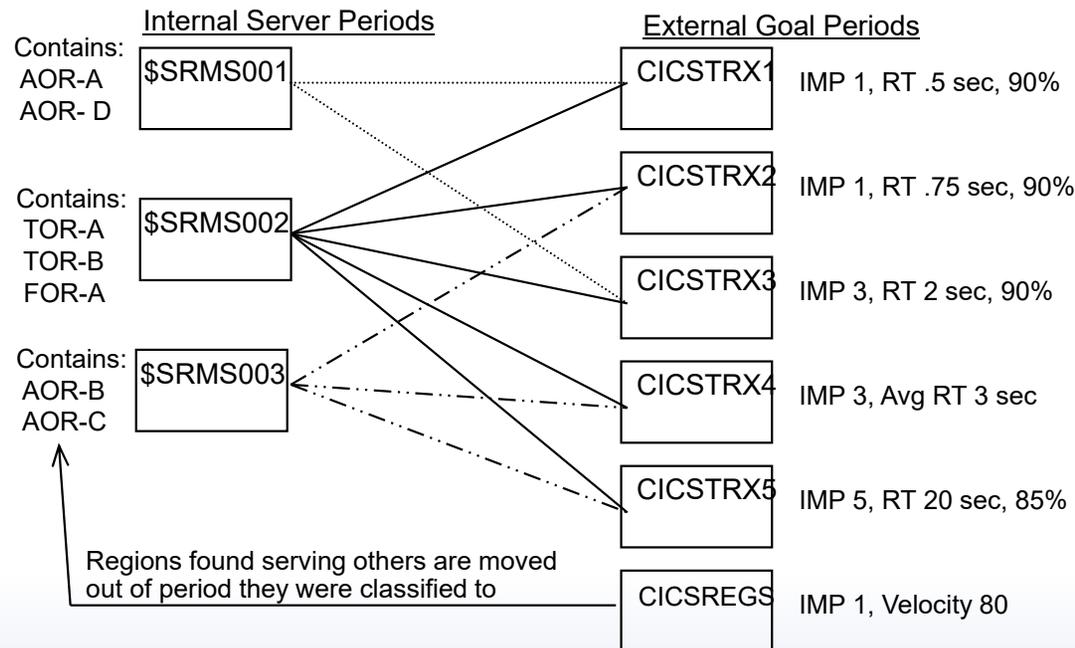


WLM Server Topology CICS Example



- 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



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



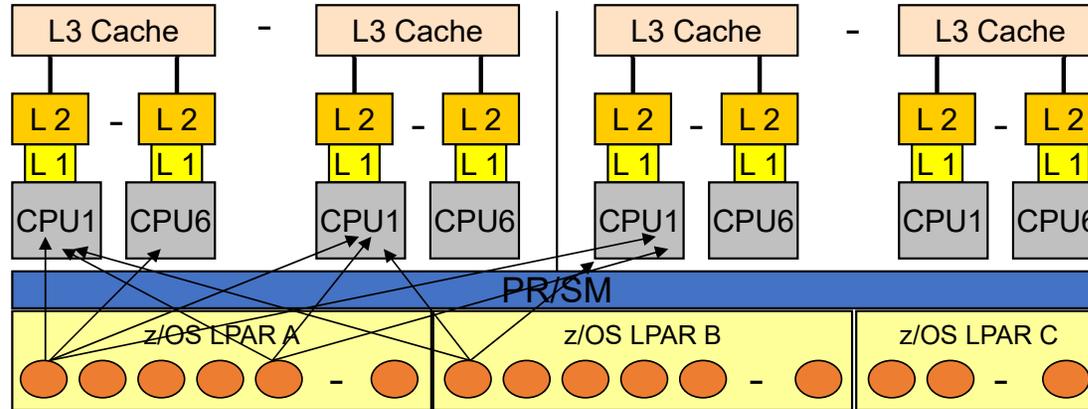
- All is running great in goal mode
 - But our installation is planning on taking advantage of system facilities that will affect the performance of the workloads
 - How do we manage what may be conflicting objectives?

- Automated Capacity controls such as HiperDispatch

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

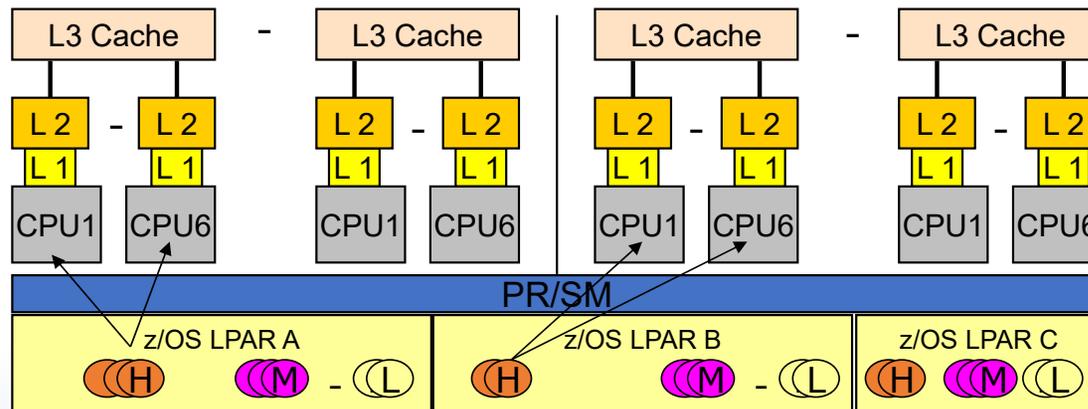


- Horizontal

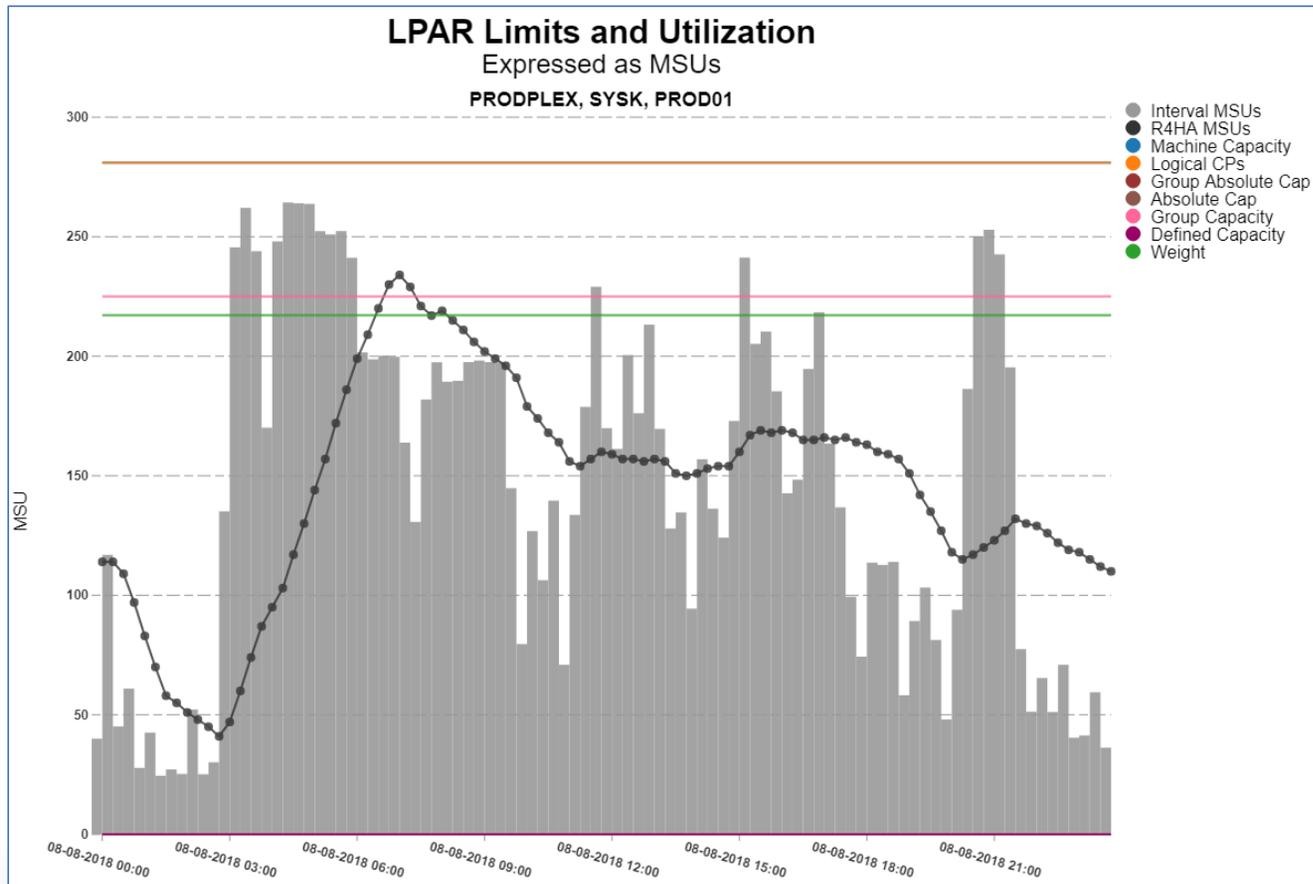


- Vertical

Arranged into High, med, low Pools by PR/SM, and affinity nodes by z/OS



LPAR Limits

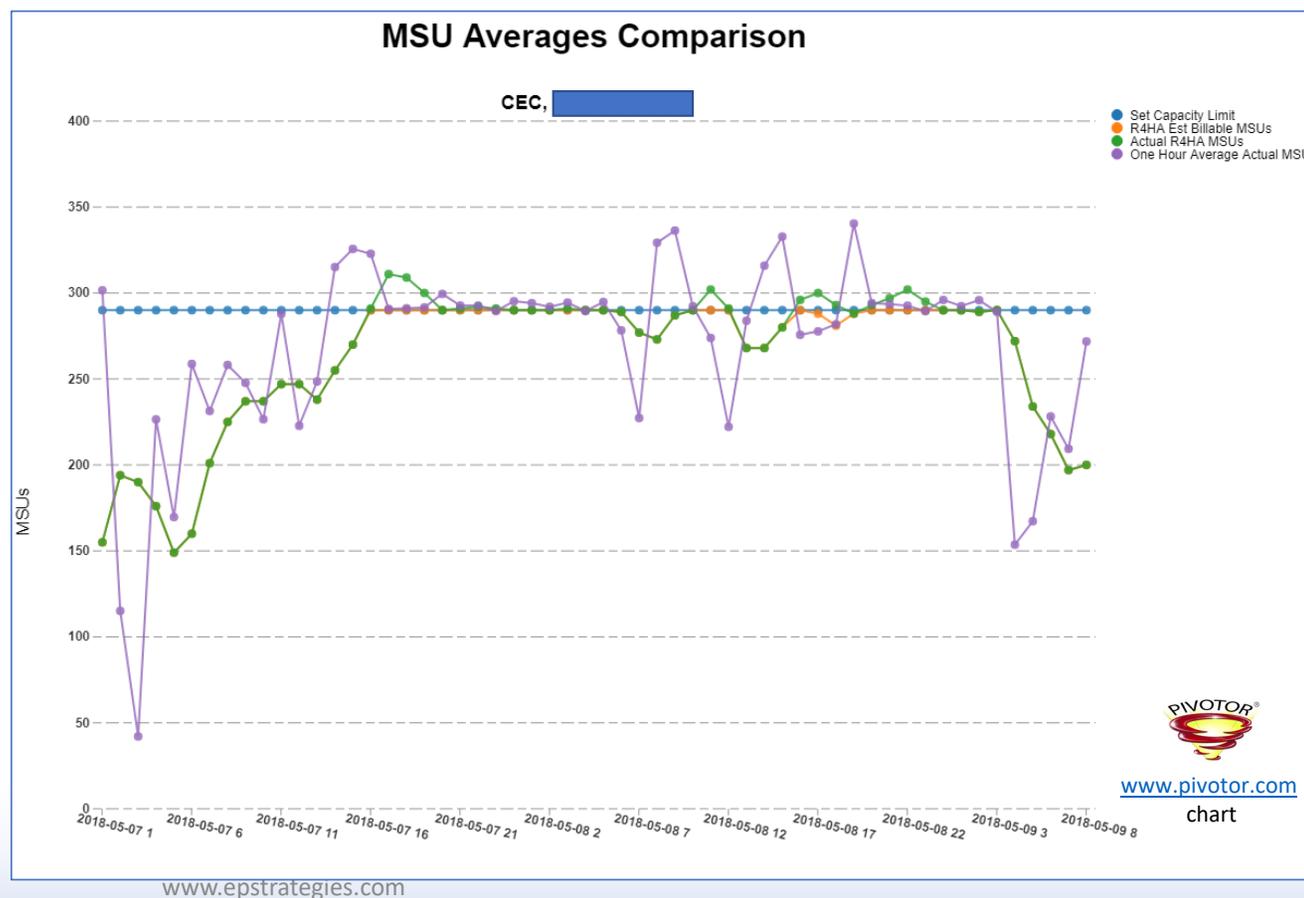


www.pivotor.com
chart

Looking for capping due to defined capacity limits



- Zooming into this one area for the machine
- Note on May 7 and 8 there are some interesting periods of time when it appears capacity limits are reached
 - Capping is most likely occurring



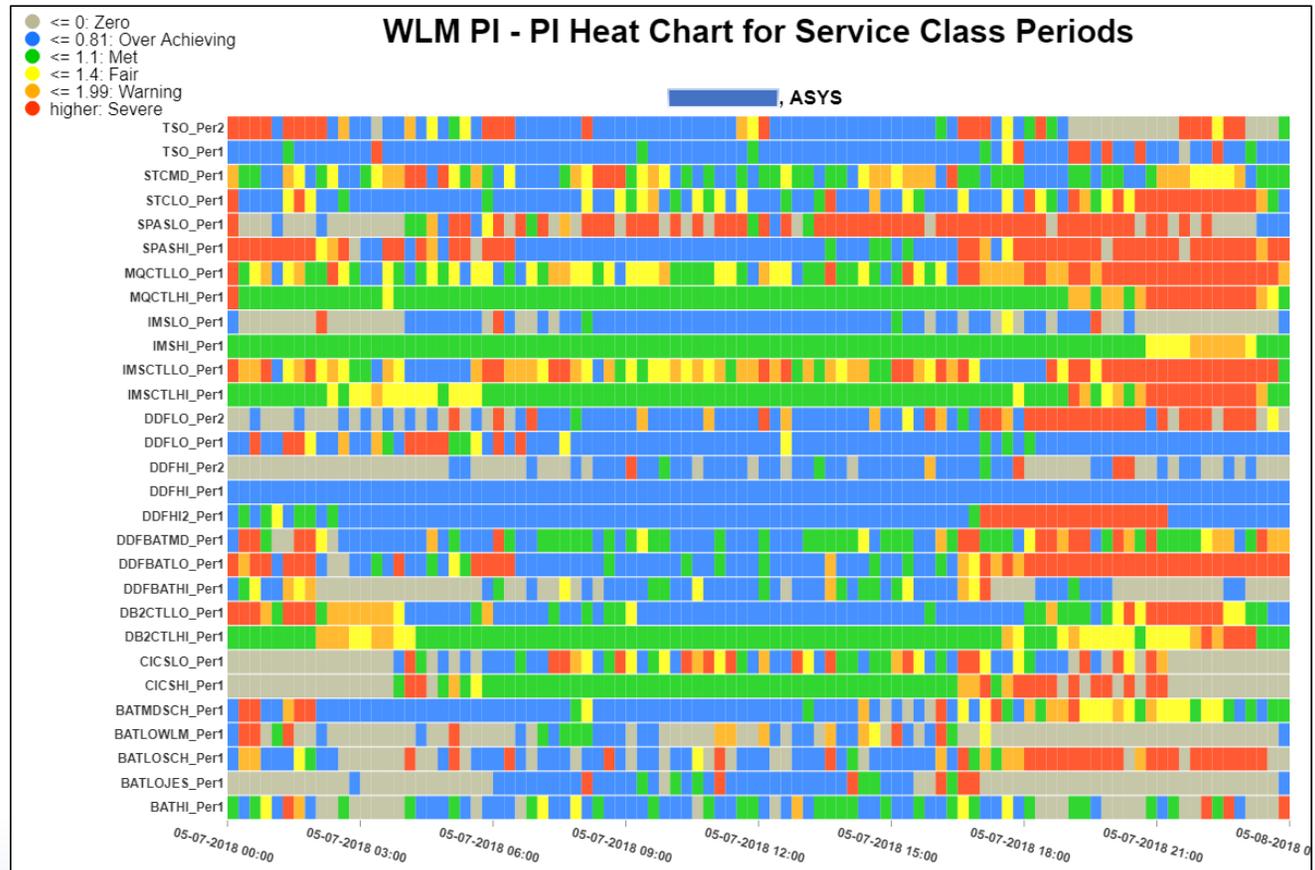
Capping affects Goals WLM Performance Indexes



- This WLM PI chart shows that when capping is enforced goals are affected
- The question is, are the lower importance workloads being hurt more
- Never assume goals and importance level are correct



www.pivotor.com
chart



CPU APPL% Consumption – By Importance, By Service Class Period

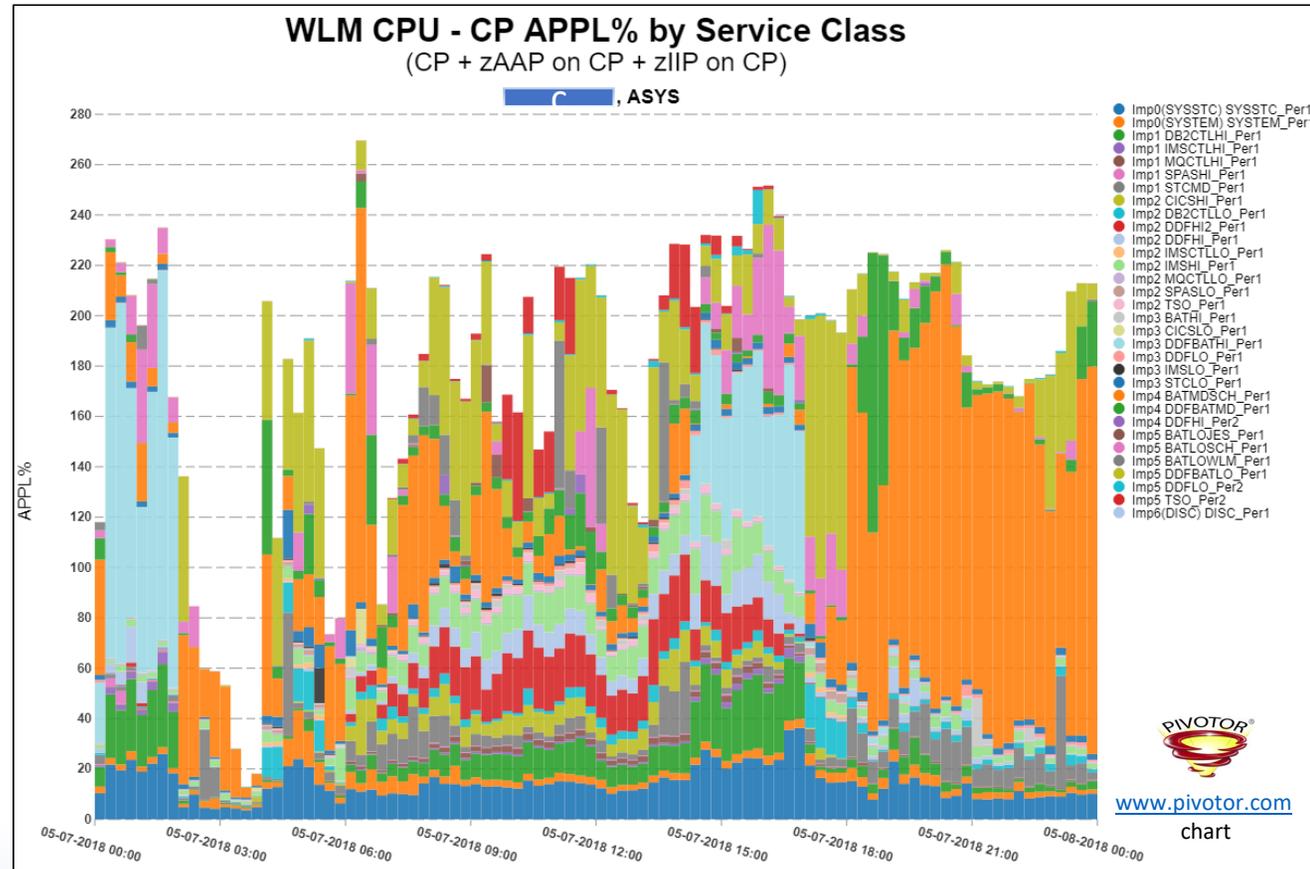


When CPU to an LPAR is limited due to either soft or hard capping the workloads will have less access to CPU capacity.

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.

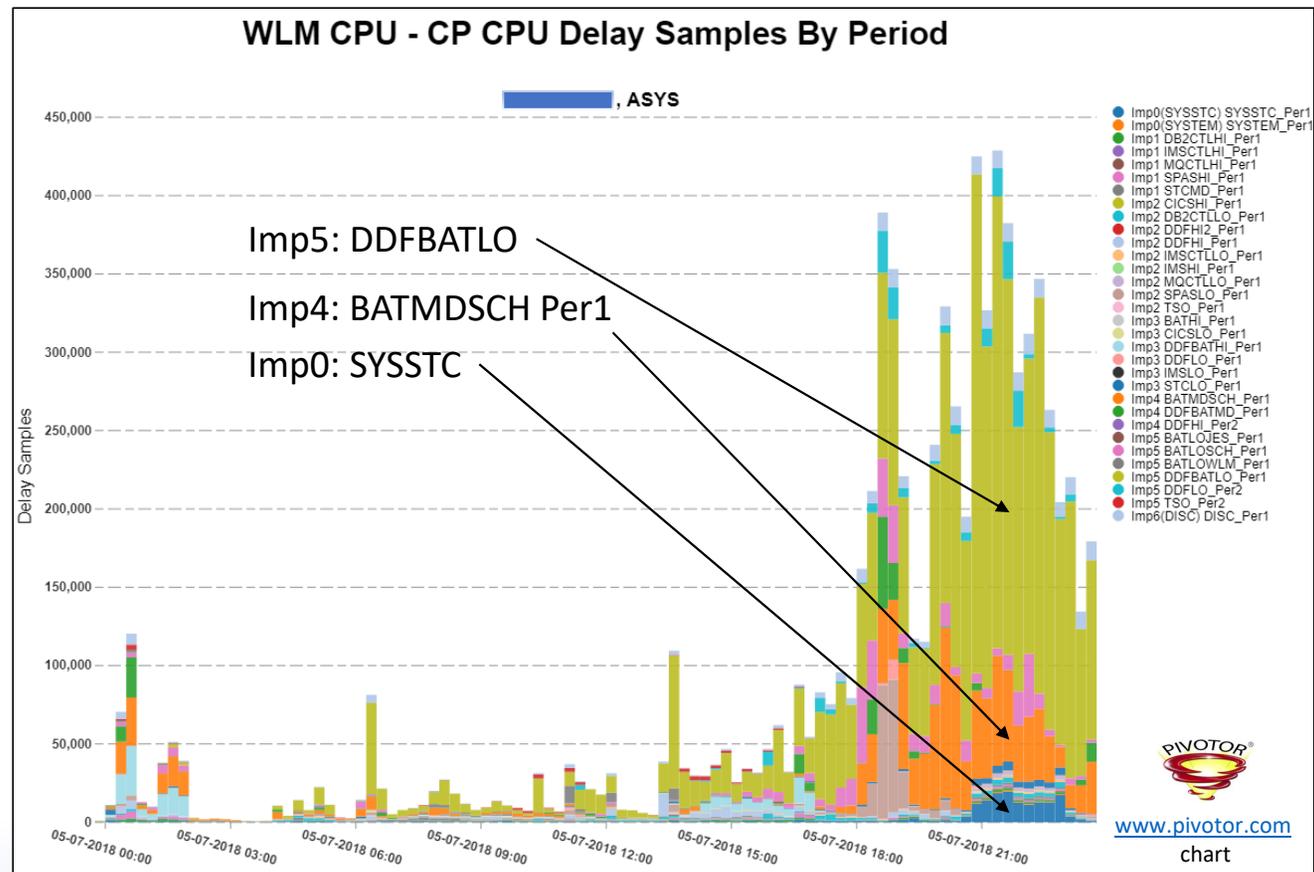
Remember, even during periods of capping, WLM tries to meet goals.



CPU delay samples – By Importance, By Service Class Period



- Look at CPU delays
 - Other delays types will be of interest, but for capping CPU delay will be the most interesting
- What work is delayed?
- Is the right work delayed?
- Is delay proportional to the work?



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?

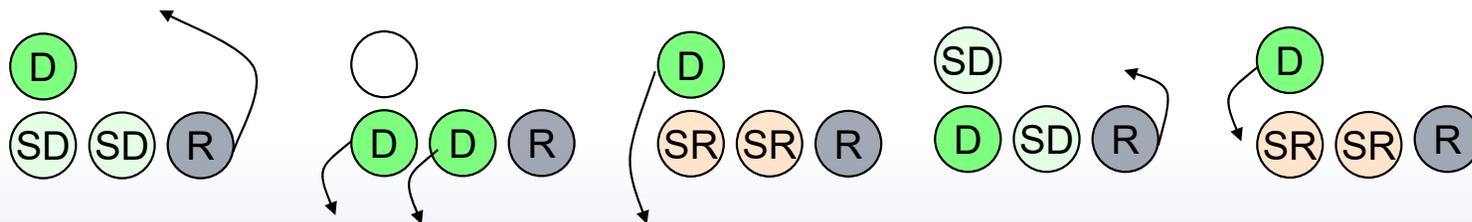
- When this usually occurs
 - Peak periods
 - Startup of a new workload or big batch job
 - Influx of workload activity
 - System problems / system dumping / subsystem ABENDs / Application problems
 - Storage shortages
 - ??????

Policy Adjustment Actions - CPU

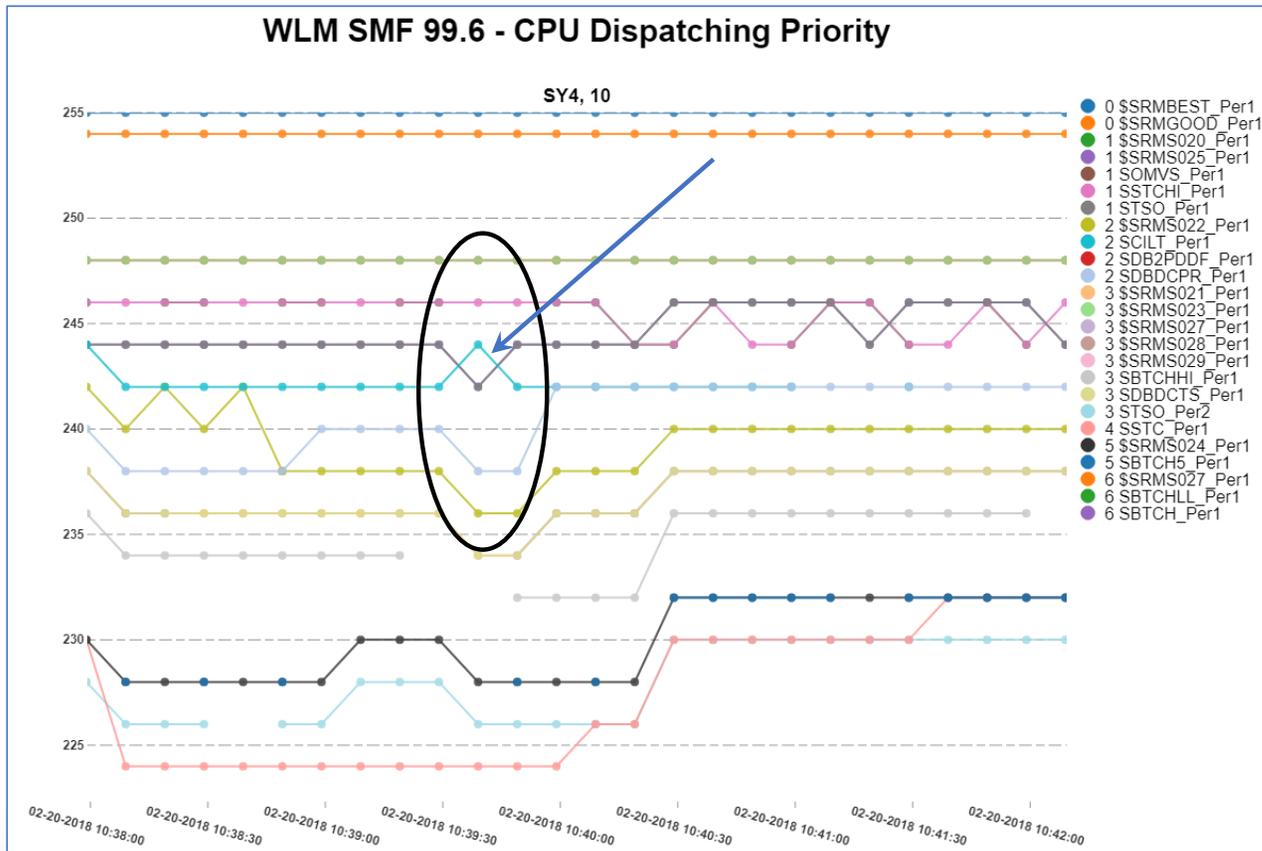


- Dispatching priority adjustments

- 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



SMF 99.6 CPU Dispatching Priority - Every 10 Seconds



SMF 99.1

- Example of WLM Actions Trace



system	Time	PA Int	RA Int	Code	Code	Explain	Local P	Sysple	Period	Service Class
SY4	10:39:39	142	173	270	PA_REC_CAND	Policy adjustment, receiver candidate selected.	50	50	1	SCILT
SY4	10:39:39	142	173	308	PA_DONOR_PERIOD	Policy adjustment, donor period.	0.64	0.94	1	STSO
SY4	10:39:39	142	173	308	PA_DONOR_PERIOD	Policy adjustment, donor period.	0.78	1.09	1	SSTCHI
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.01	0.14	1	\$SRMS024
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.01	0.14	1	SBTCH5
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.5	5.5	1	SOMVS
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.54	0.9	2	STSO
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.64	0.94	1	\$SRMS025
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.64	0.94	1	STSO
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.74	1.09	1	\$SRMS021
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.74	1.09	1	SDBDCTS
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	0.76	0.97	1	SBTCHHI
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	1.1	0.8	1	\$SRMS022
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	50	50	1	SCILT
SY4	10:39:39	142	173	525	HSK_UNBUNCH_PRTY	Housekeeping, unbunch priorities.	60	0.85	1	SDBDCPR
SY4	10:39:39	142	173	530	PA_PMDO_DON	Policy adjustment, assess moving primary processor donor down to occupied priority.	0.64	0.94	1	STSO
SY4	10:39:39	142	173	530	PA_PMDO_DON	Policy adjustment, assess moving primary processor donor down to occupied priority.	0.78	1.09	1	SSTCHI
SY4	10:39:39	142	173	531	PA_PCC_DON_VIOLTN	Policy adjustment, moving the donor to the receivers priority violates CPU critical rules.	0.78	1.09	1	SSTCHI
SY4	10:39:39	142	173	532	PA_PCC_BLK_R_IS_DON	Policy adjustment, cannot move the blocker up because it is the donor.	0.78	1.09	1	SSTCHI
SY4	10:39:39	142	173	580	PA_PMD_SEC_DON	Policy adjustment, assess moving secondary processor donor down.	0.64	0.94	1	\$SRMS025
SY4	10:39:39	142	173	580	PA_PMD_SEC_DON	Policy adjustment, assess moving secondary processor donor down.	0.78	1.09	1	\$SRMS020
SY4	10:39:39	142	173	620	PA_PMUO_REC	Policy adjustment, assess moving primary processor receiver up to occupied priority.	50	50	1	SCILT
SY4	10:39:39	142	173	635	PA_PMUUB_REC	Policy adjustment, assess moving primary processor receiver up to unoccupied priority between donor and receiver's current priorities.	50	50	1	SCILT
SY4	10:39:39	142	173	750	PA_PRO_INCP_REC	Policy adjustment, increase priority for receiver.	50	50	1	SCILT

SMF 99 Data Processing Offer



- Ever wonder what your SMF 99 records
 - The SMF 99s are a series of SMF records cut by WLM to help with performance debugging and better understanding of WLM settings and decisions.

- If you would like to see what some of your SMF 99 measurements contain
Send an email to Peter.Enrico@EPStrategies.com

- You can send Peter some raw SMF 99 data.
- Peter will process it
- Then you can have a web meeting with Scott and Peter to discuss your data.

Presentation Summary



- 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