# EPV Performance University 2021

- z/OS Performance Analysis
- WLM Update
- Db2 Performance Analysis

15th June 2021

# WLM basics - Part 2

# Agenda

- Service Class period goals
- Performance Index
- Dispatching priority
- How many Service Classes?

# Service Class period goals

# Service Class period goals

- One of these goals can be assigned to each period

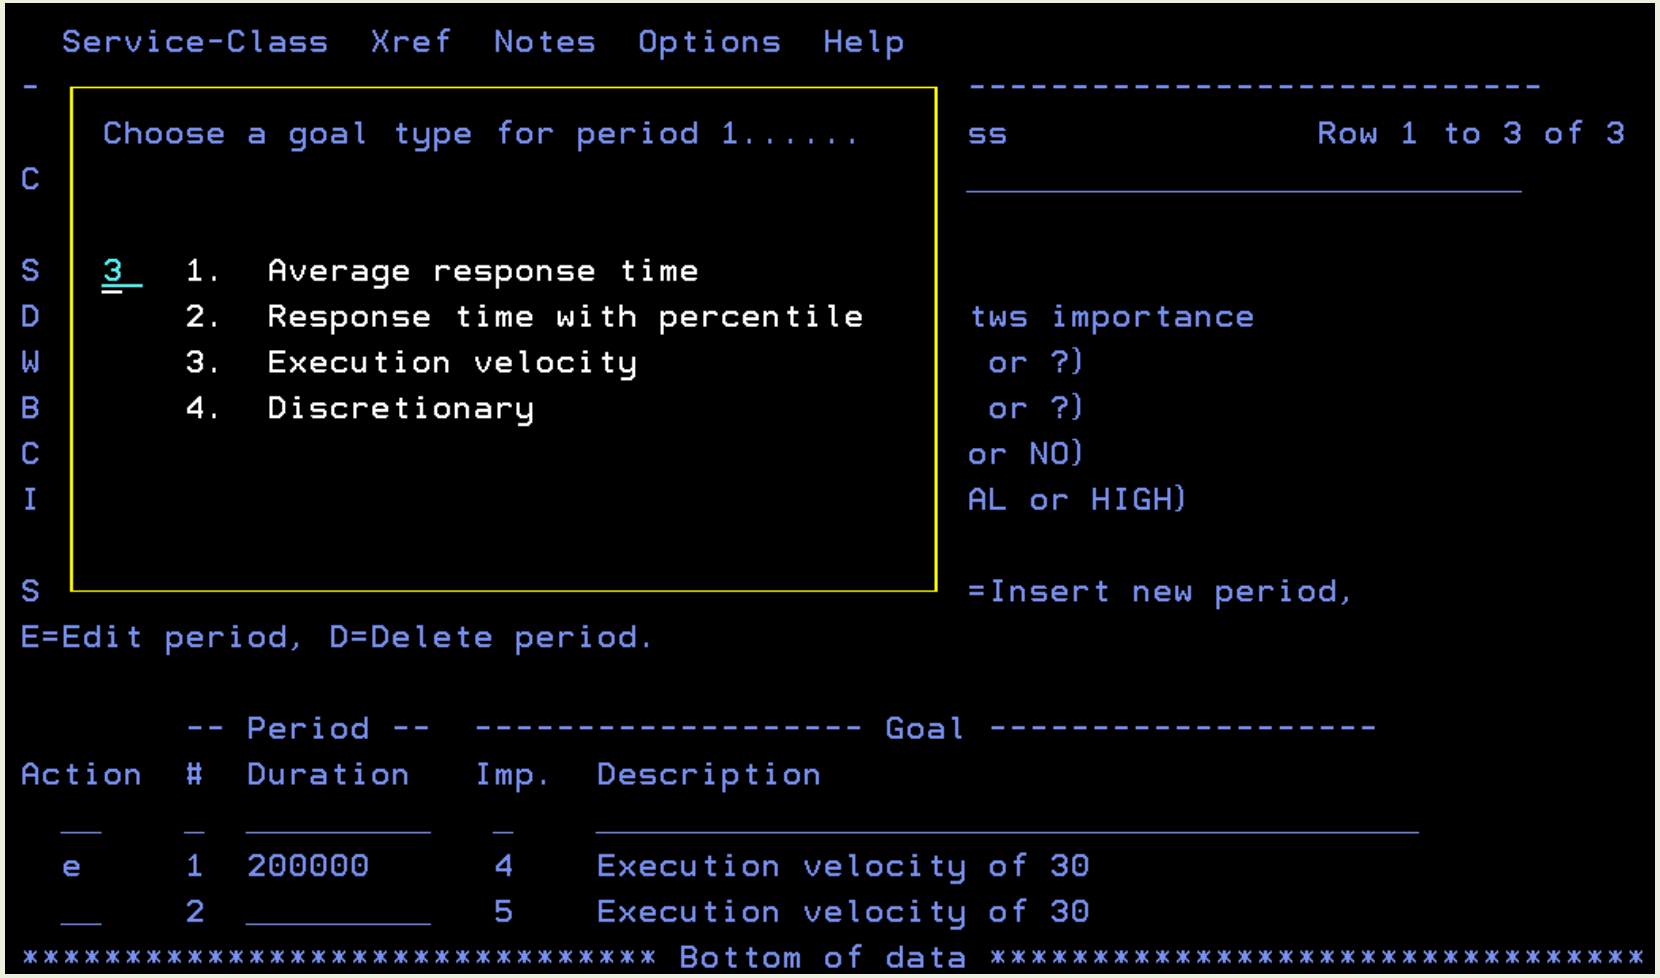

# Service Class period goals

- Average Response Time
  - Average response time for transactions to complete
  - Response time values from 15 milliseconds to 24 hours
  - Response time is in hours, minutes, or seconds with 3 decimals
- Host response time
- This goal is mostly for transaction oriented work but it can be used also for batch oriented work
- The key issue is the number of completions; if less than 30 in 1 hour response goal is not a good choice

# Service Class period goals

- If work inside the service class period is extremely variable, then the average is not meaningful; use a Response Time and Percentile goal
- Response Time and Percentile
  - You need to specify a percentile of transactions to be completed in a specified response time
  - Percentile value from 1 to 99
  - Response time values as for the average response
- Host response time
- If less than hundreds by hour it’s not a good choice

# Service Class period goals

- An Execution Velocity goal defines how fast work should run when ready, without being delayed for WLM managed resources
- Values from 1 to 99 are accepted
- It has been invented for work for which response time goals are not appropriate, such as started tasks, or long running batch work
- Velocity is not importance; don’t confuse them

# Service Class period goals

"Execution velocity is an abstract mathematical description with no objectively measurable metric"

# Service Class period goals

- That sentence is telling you that velocity is not an immediate and easy to understand concept
- There are many limitations on what can be measured and how it is measured
- Velocity is not very precise
- However we need a user goal for long running tasks; you can’t put all of them in discretionary or SYSSTC

# Service Class period goals

- WLM needs to know the amount of time spent by each workload at the different resources (using and waiting for them)
- This is the main reason why WLM periodically samples the state of each dispatchable unit (TCB and SRB) associated to an Address Space or an Enclave
- These samples are used to:
  - determine what resource is the primary bottleneck
  - evaluate the impact of WLM possible action
  - take policy adjustment decisions

# Service Class period goals

- WLM collected samples are also used for velocity calculation
- Unknown and idle states are excluded from calculation
- Only using and delays for WLM managed resources are included
- The percentage of samples contributing to velocity calculation is essential to make it a meaningful index of the level of service of the workload

# Service Class period goals

Velocity calculation (I/O priority management=Yes)

$$Vel = \frac{\text{CPU using + I/O using}}{\text{CPU using + I/O using + WLM managed delays}}$$

IO using = Connect time

managed delays = CPU, paging, swapping, MPL, IOSQ, pending, server AS, initiator, capping …

# Service Class period goals

| | |
|---|---|
| UNKNOWN DELAY | Percentage of samples the A.S. and Enclaves in the period (see Note 1) have been delayed for unknown reasons (resources not managed by WLM). |
| IDLE | Percentage of samples the A.S. and Enclaves in the period were in a IDLE situation (includes STIMER wait, TSO terminal input wait, TSO terminal output wait, APPC wait, waiting for batch initiators,etc) |
| TOTAL USING | Percentage of samples the A.S. and Enclaves in the period used system resources; this field is used to calculate the Performance Index for periods with a Velocity Goal; using samples include CPU (standard processors) usage, AAP usage, IIP usage, I/O connect time. |
| TOTAL DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a WLM managed resource; this field is used to calculate the Performance Index for periods with a Velocity Goal. |

- Using and delay samples will be used to calculate the Performance Index for periods running with a velocity goal
- Using samples are only for CPU, zIIP and I/O connect time



# Service Class period goals

| | |
|---|---|
| CPU DELAY | Percentage of samples the A.S. and Enclaves in the period. were waiting to use standard processors. |
| DASD DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IOSQ and PENDING during DASD I/O operations; this value can be included in Velocity calculations depending on the value assigned to the "I/O priority management" parameter. |
| AAP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for AAPs. |
| IIP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IIPs. |
| CPU CAPPING DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting to use standard processors for Resource Group Capping or Discretionary Goal Management; Resource Group Capping is activated by WLM when the limit of the Resource Group, to which the Service Class is assigned, has been reached (see Note 3); Discretionary Goal Management is activated when a service class period is over performing (is reaching an objective better then expected); in this situation WLM starts "capping" this period and divert resources to "discretionary" workloads; only service class periods with a velocity goal less equal 30% or those with an average response time goal greater than 1 minute can be capped in this way. |

- CPU delay includes soft capping delays

# Service Class period goals

| | |
|---|---|
| CPU DELAY | Percentage of samples the A.S. and Enclaves in the period. were waiting to use standard processors. |
| DASD DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IOSQ and PENDING during DASD I/O operations; this value can be included in Velocity calculations depending on the value assigned to the "I/O priority management" parameter. |
| AAP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for AAPs. |
| IIP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IIPs. |
| CPU CAPPING DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting to use standard processors for Resource Group Capping or Discretionary Goal Management; Resource Group Capping is activated by WLM when the limit of the Resource Group, to which the Service Class is assigned, has been reached (see Note 3); Discretionary Goal Management is activated when a service class period is over performing (is reaching an objective better then expected); in this situation WLM starts "capping" this period and divert resources to "discretionary" workloads; only service class periods with a velocity goal less equal 30% or those with an average response time goal greater than 1 minute can be capped in this way. |

- What about I/O disconnect time?

# Service Class period goals

| | |
|---|---|
| CPU DELAY | Percentage of samples the A.S. and Enclaves in the period. were waiting to use standard processors. |
| DASD DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IOSQ and PENDING during DASD I/O operations; this value can be included in Velocity calculations depending on the value assigned to the "I/O priority management" parameter. |
| AAP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for AAPs. |
| IIP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IIPs. |
| CPU CAPPING DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting to use standard processors for Resource Group Capping or Discretionary Goal Management; Resource Group Capping is activated by WLM when the limit of the Resource Group, to which the Service Class is assigned, has been reached (see Note 3); Discretionary Goal Management is activated when a service class period is over performing (is reaching an objective better then expected); in this situation WLM starts "capping" this period and divert resources to "discretionary" workloads; only service class periods with a velocity goal less equal 30% or those with an average response time goal greater than 1 minute can be capped in this way. |

- zAAP not supported anymore in modern machines (z13, z14 and z15)

# Service Class period goals

| | |
|---|---|
| CPU DELAY | Percentage of samples the A.S. and Enclaves in the period. were waiting to use standard processors. |
| DASD DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IOSQ and PENDING during DASD I/O operations; this value can be included in Velocity calculations depending on the value assigned to the "I/O priority management" parameter. |
| AAP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for AAPs. |
| IIP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for IIPs. |
| CPU CAPPING DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting to use standard processors for Resource Group Capping or Discretionary Goal Management; Resource Group Capping is activated by WLM when the limit of the Resource Group, to which the Service Class is assigned, has been reached (see Note 3); Discretionary Goal Management is activated when a service class period is over performing (is reaching an objective better then expected); in this situation WLM starts "capping" this period and divert resources to "discretionary" workloads; only service class periods with a velocity goal less equal 30% or those with an average response time goal greater than 1 minute can be capped in this way. |

- CPU capping delay only includes resource capping delays

# Service Class period goals

| | |
|---|---|
| MPL DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for the SWAP-IN process to begin, even if they were READY, due to the MPL target set by WLM. |
| AUX PRIV DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a private area PAGE FAULT operation. |
| AUX COMM DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a common area PAGE FAULT operation. |
| AUX XMEM DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a cross memory PAGE FAULT operation. |
| AUX VIO DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a VIO PAGE FAULT operation. |
| AUX SHSP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a Hiperspace Standard PAGE FAULT operation. |
| AUX EHSP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a Hiperspace ESO PAGE FAULT operation. |

- Storage related delays (very uncommon these days)



# Service Class period goals

| | |
|---|---|
| SERV PRIV DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a server A.S. private area PAGE FAULT operation. |
| SERV VIO DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a server A.S. VIO PAGE FAULT operation. |
| SERV HSP DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a server A.S. Hiperspace PAGE FAULT operation. |
| SERV SWAP-IN DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a server A.S. SWAP-IN process to begin; a high value generally indicates an auxiliary storage subsystem problem. |
| SERV MPL DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a server A.S. SWAP-IN process to begin due to the MPL targer set by WLM. |

- Storage related delays due to server address spaces (very uncommon these days)



# Service Class period goals

| | |
|---|---|
| QMPL DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for a "server" A.S. availability; this delay is tipical with Application Environments when WLM can activate more Address space to answer to incoming requests based on the importance of the A.S.; Application Environments are generally used for DB2 Stored Procedures, HTTP pages requests and for WEBSPHERE (servlet, JSP, EJB); this delay includes also WLM managed initiators delays. |
| JES INIT DELAY | Percentage of samples the A.S. and Enclaves in the period were waiting for an initiator; this metric is available only for JES managed initiators. These samples are not included in PI calculation. |

- Initiator related delays
- JES init delays are not included in Performance Index calculation because JES initiators are not managed by WLM



# Service Class period goals

- Each workload has its own velocity, depending on the type of work and the hardware configuration
- You should set velocity goals based on what you measure in not constrained peak hours
- It is normally useless to define several service class periods with the same importance and similar velocity goals
- Velocity values within the same importance should have a difference of 10% to 15%

# Service Class period goals

- Velocity goals are more sensitive to configuration changes than response time goals and should be  adjusted when required after configuration changes
- These configuration changes include:
  - Change of the physical machine characteristics
  - Changes to the capacity available to the system
  - Changes to the number of online processors
  - Implementation of HiperDispatch
  - Implementation of SMT

# Service Class period goals

*Table 1. Recommended velocity goal based on the number of central processors per LPAR*

| Number of CPs defined for the LPAR | Recommended velocity goal |
|---|---|
| 1-5 | 50-70 |
| 6-15 | 60-80 |
| More than 15 | 70-90 |

From DB2 12 for z/OS - Managing Performance

# Service Class period goals

- A discretionary goal means “do the best that you can”
- It is the goal of the SYSOTHER service class (the  default service class with the exception of STC)
- Default STC service class is SYSSTC; you should create a default service class to manage unknown STCs
- Discretionary is usually applied to batch jobs
- A discretionary goal can be assigned only to the last period of a service class

# Service Class period goals

- A z/OS system can run at 100% CPU busy without problems if you have enough discretionary work
- WLM knows immediately which work must donate resources when an important workload spikes
- The percentage of discretionary work should be high enough to allow the fluctuations of the non discretionary work
- Not useful if nondiscretionary work may drive the system to 100% busy for long periods of time



# Service Class period goals

- WLM processes work with discretionary goal using resources not required to meet the goals of other service classes
- Discretionary work always wait for the zIIP; it doesn't overflow to CPU
- In the past many problems of discretionary work not running for a long time were experimented
- Many enhancements have been done which makes discretionary goal more usable now

# Service Class period goals

- Reduced Preemption
  - to reduce the overhead of re-dispatching many times less important work
  - to permit Discretionary work to run
- z/OS  usually avoids interrupting work execution as soon as a new work becomes ready to execute
- Rather it relies on the normal behavior of work to release the CPU voluntarily, and backs this up with timed preemptions

# Service Class period goals

- Blocked workload support
  - If the CPU utilization of a system is at 100%, workloads with low importance might not get dispatched for a long time
  - This could cause problems if the work holds a resource delaying more important workloads
- WLM checks if some work does not get CPU service for a certain time interval due to its low dispatch priority and temporarily promotes it to a higher dispatch priority
- This helps to complete low priority work in a finite time period, without permanently delaying high priority work

# Service Class period goals

| | |
|---|---|
| AVG_BLOCKED | Average number of address spaces and enclaves found blocked during the reporting interval. |
| MAX_BLOCKED | Maximum number of address spaces and enclaves found blocked during the reporting interval. |
| PROMOTE_RATE | Blocked dispatchable work units promoted (increasing their dispatching priority) per second. If this value is close to the MAX PROMOTE RATE you should evaluate the opportunity to increase the BLWLTRPCT OPT parameter which specifies how much of the CPU capacity is to be used to promote blocked workloads. |
| MAX_PROMOTE_RATE | Max number of blocked dispatchable work units which may get promoted (increasing their dispatching priority) per second. |

- The statistics provided above can be used to check blocked workload support

# Service Class period goals

- Max promote rate reduced in this example
- Any idea of the reason?

SYSTEM STATES SWITCH

SYSTEM STATES - Wed, 1 Feb 2017 - PROD

| STATE | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| IN | 309,0 | 318,4 | 318,6 | 318,7 | 316,4 | 488,1 | 532,7 | 550,3 | 568,7 | 579,0 | 585,1 | 584,9 | 580,0 |
| IN_READY | 5,1 | 5,0 | 3,9 | 5,3 | 4,1 | 4,3 | 3,5 | 4,8 | 6,3 | 6,5 | 6,8 | 7,3 | 14,8 |
| OUT_WAIT | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 |
| OUT_READY | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,1 |
| LOGICAL_OUT_WAIT | 99,8 | 115,3 | 103,5 | 98,1 | 95,8 | 92,2 | 113,3 | 123,3 | 171,4 | 199,3 | 241,3 | 244,8 | 241,6 |
| LOGICAL_OUT_READY | 0,1 | 0,0 | 0,0 | 0,0 | 0,0 | 0,1 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,3 |
| ASCH_ACTIVE | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 |
| BATCH_ACTIVE | 67,7 | 73,4 | 69,9 | 72,2 | 70,6 | 237,7 | 284,2 | 298,5 | 316,5 | 323,6 | 324,3 | 323,9 | 321,0 |
| OMVS_ACTIVE | 55,3 | 70,3 | 61,1 | 65,0 | 60,0 | 60,0 | 60,0 | 62,7 | 90,3 | 94,9 | 102,5 | 100,5 | 102,0 |
| STC_ACTIVE | 278,5 | 283,9 | 285,2 | 277,5 | 279,5 | 280,6 | 297,7 | 296,8 | 303,9 | 302,1 | 319,1 | 316,8 | 312,7 |
| TSO_ACTIVE | 7,4 | 6,2 | 6,0 | 2,1 | 2,1 | 2,1 | 4,0 | 15,6 | 29,5 | 57,8 | 80,6 | 88,5 | 86,2 |
| AVG_BLOCKED | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 |
| MAX_BLOCKED | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 |
| PROMOTE_RATE | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,1 | 0,0 | 0,1 |
| MAX_PROMOTE_RATE | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 159,0 | 148,2 | 122,3 |

# Service Class period goals

SYSTEM WLC USAGE SWITCH

SYSTEM WLC USAGE - Wed, 1 Feb 2017 - PROD

| METRIC | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| CEC MSU | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 | 7.205,0 |
| IMAGE MSU | 650,0 | 650,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 700,0 |
| GROUP MSU | 650,0 | 650,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 700,0 |
| DEF MSU | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 |
| MIN ENT | 497,9 | 497,9 | 490,2 | 490,2 | 490,2 | 490,2 | 490,2 | 490,2 | 490,2 | 490,2 | 490,2 | 490,2 | 536,2 |
| ROLLING 4*HOUR | 462,7 | 438,5 | 412,0 | 424,2 | 385,3 | 332,8 | 293,0 | 232,5 | 248,3 | 287,7 | 369,7 | 531,3 | 602,7 |
| MAX ENT | 650,0 | 650,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 640,0 | 700,0 |
| % ACT SOFTCAPP | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 27,6 | 92,3 |
| % WLM SOFTCAPP | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 0,0 | 27,8 | 92,5 |
| GROUP MSU AVA | 111,2 | 128,8 | 150,0 | 156,5 | 199,2 | 253,3 | 297,3 | 349,8 | 329,7 | 281,3 | 189,0 | -23,0 | -29,7 |

# Performance Index

# Performance Index

- WLM maintains a Performance Index (PI) for each Service Class period, to measure how the actual performance varies from the goal
- PI is a “simple” number allowing to compare work with different goal types:

| | |
|---|---|
| <1 | performance better than the goal |
| =1 | performance equal to the goal |
| >1 | performance worse than the goal |

- Sysplex PI and System PI
- Updated every 10 seconds



# Performance Index

- Calculating the PI for a service class period with Average Response goal

$$\text{PI AVG response} = \frac{\text{Actual value}}{\text{Goal}}$$

- Let's suppose the goal is 3 seconds and the measured value is 6 seconds

**PI = 6 / 3 = 2**

- Goal is missed so PI is GT 1

# Performance Index

- Calculating the PI for a service class period with Execution Velocity goal

$$\text{PI Velocity} = \frac{\text{Goal}}{\text{Actual value}}$$

- Let's suppose the goal is 60% and the measured value is 20%

**PI = 60% / 20% = 3**

- Goal is missed so PI is GT 1

# Performance Index

- Calculating the PI for a service class period with a Response and Percentile goal is a bit more complex
- The goal is composed of two values: a response time goal and a percentage goal (i.e. 90% of the ended transactions completed in less than 10 seconds)
- WLM splits the response time goal you defined into 14 response time buckets, to accumulate the transactions depending on their response time

# Performance Index

- Bucket upper limits
  1) 50% of the goal
  2) 60% of the goal
  3) 70% of the goal
  4) 80% of the goal
  5) 90% of the goal
  6) 100% of the goal
  7) 110% of the goal
  8) 120% of the goal
  9) 130% of the goal
  10) 140% of the goal
  11) 150% of the goal
  12) 200% of the goal
  13) 400% of the goal
  14) >400% of the goal

# Performance Index

- These considerations apply:
  - If **N** transactions complete in bucket 1 the PI is 0,5; it's the minimal possible value; over achieved goal;
  - if **N** transactions complete in bucket 1 to 6 the PI is 1,0; goal exactly met;
  - if **N** transactions complete in buckets from 7 to 13 the goal is missed; the PI goes from 1,1 to 4,0; 4,0 is the maximum possible value and indicates an heavily missed goal;
  - if the sum of all the transactions belonging to buckets 1 to 13 is still below **N**, the PI is so bad that it can't be calculated; it is shown as '**' in EPV



# Performance Index

- Exercise 1

Goal is 90% before 0,030 sec; calculate N and PI at 9 and 10

| SRV | | BUCKET | 9 | 10 | | | 9 | 10 |
|---|---|---|---|---|---|---|---|---|
| MYSRVCLS | <= | 0,015 | 249.048 | 393.933 | | | 249.048 | 393.933 |
| MYSRVCLS | <= | 0,018 | 9.077 | 9.677 | | | 258.125 | 403.610 |
| MYSRVCLS | <= | 0,021 | 6.508 | 7.107 | | | 264.633 | 410.717 |
| MYSRVCLS | <= | 0,024 | 6.821 | 8.876 | | | 271.454 | 419.593 |
| MYSRVCLS | <= | 0,027 | 6.510 | 7.528 | | | 277.964 | 427.121 |
| MYSRVCLS | <= | 0,030 | 5.244 | 5.580 | | | 283.208 | 432.701 |
| MYSRVCLS | <= | 0,033 | 3.030 | 3.375 | | | 286.238 | 436.076 |
| MYSRVCLS | <= | 0,036 | 4.239 | 5.005 | | | 290.477 | 441.081 |
| MYSRVCLS | <= | 0,039 | 3.892 | 4.453 | | | 294.369 | 445.534 |
| MYSRVCLS | <= | 0,042 | 3.510 | 3.903 | | | 297.879 | 449.437 |
| MYSRVCLS | <= | 0,045 | 3.096 | 3.538 | | | 300.975 | 452.975 |
| MYSRVCLS | <= | 0,060 | 11.675 | 13.620 | | | 312.650 | 466.595 |
| MYSRVCLS | <= | 0,120 | 15.012 | 18.953 | | | 327.662 | 485.548 |
| MYSRVCLS | > | 0,120 | 7.972 | 9.916 | | | 335.634 | 495.464 |
| | | Total | 335.634 | 495.464 | | | | |

# Performance Index

- Exercise 1

At 9, N is 335.634 * ,9 = 302.071

At 10, N is 495.464 * ,9 = 445.918

At 9, 302.071 transactions completed in bucket 12 so PI is 2

At 10, 445.918 transactions completed in bucket 10 so PI is 1,4

| PCT RESPONSE GOAL | | 0,030 | | | | | | |
|---|---|---|---|---|---|---|---|---|
| PCT GOAL | | 90% | | | | | | |
| | | | | | | | | |
| **SRV** | | **BUCKET** | **9** | **10** | | | **9** | **10** |
| MYSRVCLS | <= | 0,015 | 249.048 | 393.933 | | | 249.048 | 393.933 |
| MYSRVCLS | <= | 0,018 | 9.077 | 9.677 | | | 258.125 | 403.610 |
| MYSRVCLS | <= | 0,021 | 6.508 | 7.107 | | | 264.633 | 410.717 |
| MYSRVCLS | <= | 0,024 | 6.821 | 8.876 | | | 271.454 | 419.593 |
| MYSRVCLS | <= | 0,027 | 6.510 | 7.528 | | | 277.964 | 427.121 |
| MYSRVCLS | <= | 0,030 | 5.244 | 5.580 | | | 283.208 | 432.701 |
| MYSRVCLS | <= | 0,033 | 3.030 | 3.375 | | | 286.238 | 436.076 |
| MYSRVCLS | <= | 0,036 | 4.239 | 5.005 | | | 290.477 | 441.081 |
| MYSRVCLS | <= | 0,039 | 3.892 | 4.453 | | | 294.369 | 445.534 |
| MYSRVCLS | <= | 0,042 | 3.510 | 3.903 | | | 297.879 | **449.437** |
| MYSRVCLS | <= | 0,045 | 3.096 | 3.538 | | | 300.975 | 452.975 |
| MYSRVCLS | <= | 0,060 | 11.675 | 13.620 | | | **312.650** | 466.595 |
| MYSRVCLS | <= | 0,120 | 15.012 | 18.953 | | | 327.662 | 485.548 |
| MYSRVCLS | > | 0,120 | 7.972 | 9.916 | | | 335.634 | 495.464 |
| | | Total | 335.634 | 495.464 | | | | |
| | | | | | | | | |
| | | N | **302.071** | **445.918** | | PI | **2** | **1,4** |

# Performance Index

- Discretionary goal management
  - overachieving work will be capped by WLM in order to allow discretionary work to run
  - service classes eligible to capping if they have a response time goal of over 1 minute or a velocity goal less equal 30
- Capping begins when the PI drops below .71 and stops when it rises above .81 (fixed discretionary PI)
- Discretionary goal management uses internally Resource Groups; so it does not apply to service classes already belonging to a Resource Group



# Performance Index

- As we saw in the previous session, a new option is  provided to deactivate WLM's discretionary goal management
- The default for Deactivate Discretionary Goal Management is NO, which enables this kind of resource donation
- If  you specify YES, you deactivate this kind of resource donation and workload management will not cap processor resources in order to help discretionary work

# Performance Index

```
  Coefficients/Options  Notes  Options  Help
------------------------------------------------------------------------------
              Service Coefficient/Service Definition Options
Command ===> ____________________________________________________________

Enter or change the Service Coefficients:

CPU  . . . . . . . . . . . . . _______  (0.1-99.9)
IOC  . . . . . . . . . . . . . _______  (0.0-99.9)
MSO  . . . . . . . . . . . . . _______  (0.0000-99.9999)
SRB  . . . . . . . . . . . . . _______  (0.0-99.9)

Enter or change the service definition options:

I/O priority management  . . . . . . . . . NO   (Yes or No)
Enable I/O priority groups . . . . . . . . NO   (Yes or No)
Dynamic alias tuning management  . . . . . YES  (Yes or No)
Deactivate Discretionary Goal Management   NO   (Yes or No)
```

# Dispatching priority

# Dispatching priority

- SYSTEM and SYSSTC predefined service classes run at a fixed dispatching priority
- No goals can be set for them by the user
- You can assign workloads to SYSTEM and SYSSTC in the classification rules

| CPU dispatching priorities | | |
|---|---|---|
| 255 | FF | SYSTEM |
| 254 | FE | SYSSTC |
| 253-249 | FD-F9 | Reserved |
| 248 | F8 | Small consumers |
| 247-203 | F7-CB | Importance 1-5 |
| 202 | CA | Not used |
| 201-192 | C9-C0 | Discretionary |
| 191 | BF | Quiesce |

# Dispatching priority

- One special priority is used by WLM for service classes that consume very little CPU service

| CPU dispatching priorities | | |
|---|---|---|
| 255 | FF | SYSTEM |
| 254 | FE | SYSSTC |
| 253-249 | FD-F9 | Reserved |
| 248 | F8 | Small consumers |
| 247-203 | F7-CB | Importance 1-5 |
| 202 | CA | Not used |
| 201-192 | C9-C0 | Discretionary |
| 191 | BF | Quiesce |

# Dispatching priority

- For service class periods with a user goal (Importance 1-5) dispatching priority is dynamically managed to help them reach that goal

| CPU dispatching priorities | | |
|---|---|---|
| 255 | FF | SYSTEM |
| 254 | FE | SYSSTC |
| 253-249 | FD-F9 | Reserved |
| 248 | F8 | Small consumers |
| 247-203 | F7-CB | Importance 1-5 |
| 202 | CA | Not used |
| 201-192 | C9-C0 | Discretionary |
| 191 | BF | Quiesce |

# Dispatching priority

- Discretionary work is managed by a MTTW algorithm to improve throughput giving a higher dispatch priority to work doing more I/O and a lower dispatch priority to work using a lot of CPU (AVG CPU time between I/Os)

| CPU dispatching priorities | | |
|---|---|---|
| 255 | FF | SYSTEM |
| 254 | FE | SYSSTC |
| 253-249 | FD-F9 | Reserved |
| 248 | F8 | Small consumers |
| 247-203 | F7-CB | Importance 1-5 |
| 202 | CA | Not used |
| 201-192 | C9-C0 | Discretionary |
| 191 | BF | Quiesce |

# Dispatching priority

- The last used dispatching priority is 191 which is assigned to quiesced work

| CPU dispatching priorities | | |
|---|---|---|
| 255 | FF | SYSTEM |
| 254 | FE | SYSSTC |
| 253-249 | FD-F9 | Reserved |
| 248 | F8 | Small consumers |
| 247-203 | F7-CB | Importance 1-5 |
| 202 | CA | Not used |
| 201-192 | C9-C0 | Discretionary |
| 191 | BF | Quiesce |

# How many service classes?

# How many service classes?

- It’s strongly suggested to use not more than 25 to 35 active service class periods, with goals, at any point in time
- Having too many service class periods to manage has negative impacts on:
  - WLM responsiveness
  - WLM decisions quality

# How many service classes?

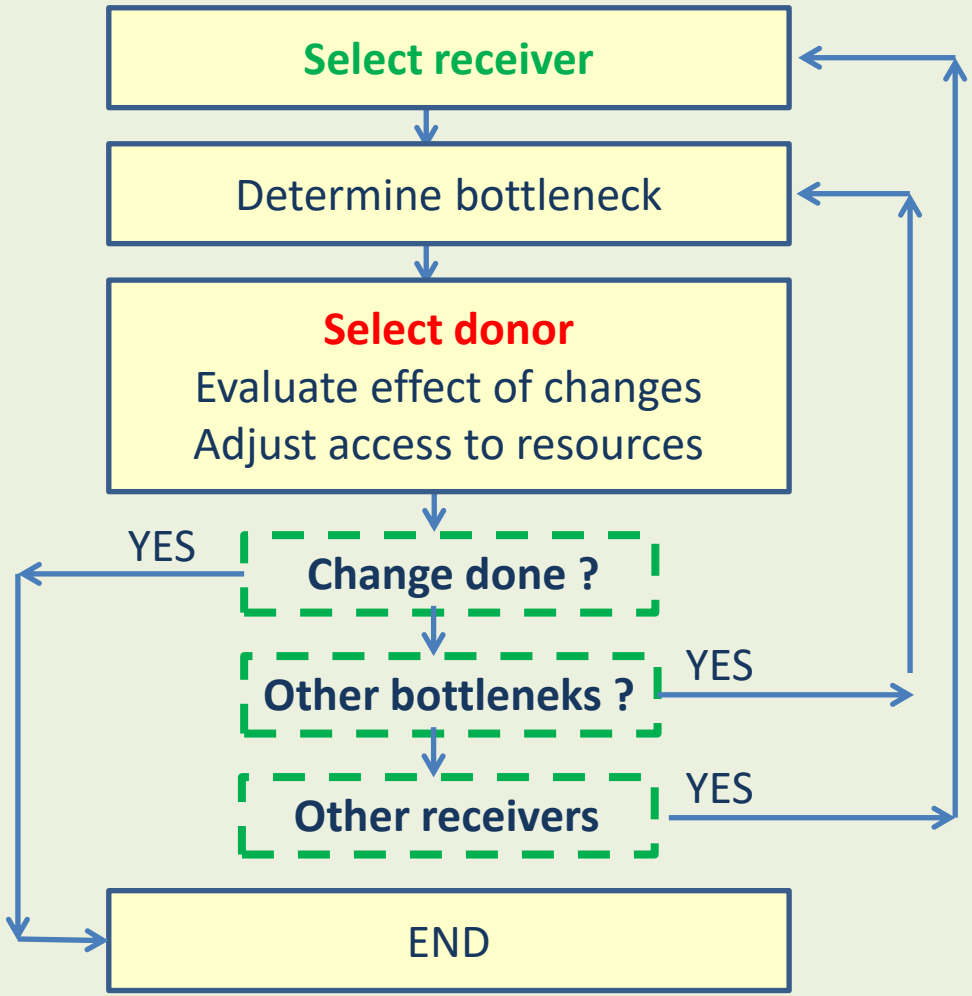


Most important SC period missing the goal first

Discretionary SC period first

Every 250 ms SRM
Collects measurements

Every 10 seconds WLM
Uses these measurements

# How many service classes?

- WLM responsiveness

Every 10 seconds, WLM selects the service class period in the most need of help and takes one policy action on its behalf

The more active service class periods you have that need assistance, the more 10-second cycles it will take WLM to take all the needed policy actions for these service classes

This could reduce WLM's responsiveness in managing to goals

# How many service classes?

- WLM decisions quality

As part of its decision-making, WLM makes projections for how work will behave when resource changes are made

The accuracy of these projections depends on the quality of the history data gathered for a service class

The more like work that can be put into a single service class, the more statistically valid will be its history data, and the better will be WLM's decisions

# Questions?