EPV Performance University 2021

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

24th February 2021

Resource groups

Agenda

- Introduction
- Exploiting resource groups
- Measuring resource groups

Introduction

Introduction

- Resource Groups are the only available technique able to limit the CPU (no zAAP, no zIIP) used by a workload inside an LPAR (they can also guarantee a minimum amount of CPU)
- Resource groups exploit the WLM capping function
- WLM service class granularity
- A service class can be assigned to only one resource group but more service classes can be assigned to the same resource group

Introduction

```
  Service-Class  Xref  Notes  Options  Help
------------------------------------------------------------------------------
                          Modify a Service Class                Row 1 to 3 of 3
Command ===> ____________________________________________________________________

Service Class Name . . . . . : BJOBTWS
Description  . . . . . . . . . JOBS scheduled by TWS_____________
Workload Name  . . . . . . . . BJOB____  (name or ?)
Base Resource Group  . . . . . ________  (name or ?)
Cpu Critical . . . . . . . . . NO      (YES or NO)
I/O Priority Group . . . . . . NORMAL  (NORMAL or HIGH)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

         -- Period --  ------------------ Goal ------------------
Action  #  Duration    Imp.  Description
  ___   _  _________    _    ____________________________________________
  ___   1  200000       4    Execution velocity of 30
  ___   2  _________    5    Execution velocity of 30
******************************** Bottom of data ********************************
```

Introduction

- Resource Groups main goal is to set a limit to the amount of CPU used by non important and CPU intensive workloads (e.g. batch) to:
  - protect business critical workloads
  - postpone an upgrade
  - reduce the software bill (when in WLC)
- Resource Groups can be used to stop not swappable address spaces when in loop
- They can be used also to fulfil the terms of an outsourcing contract
- Control enforced at WLM main cycle (10 seconds)

Introduction

- Resource groups limit the possibility for WLM to manage the workload
- You have to use them carefully and only if really needed
- You should avoid capping everything
- For most systems, you can let WLM decide how to manage the resources in the sysplex and not use resource groups

Introduction

- There are historically three types of Resource Groups:
  - Type 1: capacity limits are specified in un-weighted CPU (CPU+SRB) service units per second (SU/sec); it has a Sysplex wide scope
  - Type 2: capacity limits are specified as a percentage of the LPAR capacity; the sum of all Resource Groups minimum values should not exceed 99; it has a system scope
  - Type 3: capacity is specified as a percentage of a single general purpose CP; 100 represents the capacity of one CPU; the sum of all Resource Group minimum values should not exceed the number of processors x 100; it has a system scope

Introduction

  - Type 4; it has been introduced to support container pricing; capacity limits can be specified in accounted (captured) workload MSU; it has a sysplex-wide scope

- Every 10 sec, WLM converts type 2, 3, and 4 capacity limits into a SU/sec service rate based on current configuration

Introduction

- A new option allows to include Specialty Processor Consumption in capacity limits calculation
- NO means that WLM management of the resource group is based only on CPU consumptions (default)
- YES means that management of the resource group is based on CPU and zIIP total consumptions

```
 Resource-Group  Xref  Notes  Options  Help
------------------------------------------------------------
                          Modify a Resource Group
Command ===> ____________________________________________________
Resource Group Name  . . . . . : TSOLIMIT
Description  . . . . . . . . . . TSO capping

Define Capacity:
-    1.  In Service Units (Sysplex Scope)
     2.  As Percentage of the LPAR share (System Scope)
     3.  As a Number of CPs times 100 (System Scope)
     4.  In accounted workload MSU (Sysplex scope)
Minimum Capacity . . . . . . . 0
Maximum Capacity . . . . . . . 0

Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope):    GB
```

Introduction

- A memory limit, in GB, can be specified for any resource group
- It allows to limit the amount of memory capacity available to one or more service classes at system-level

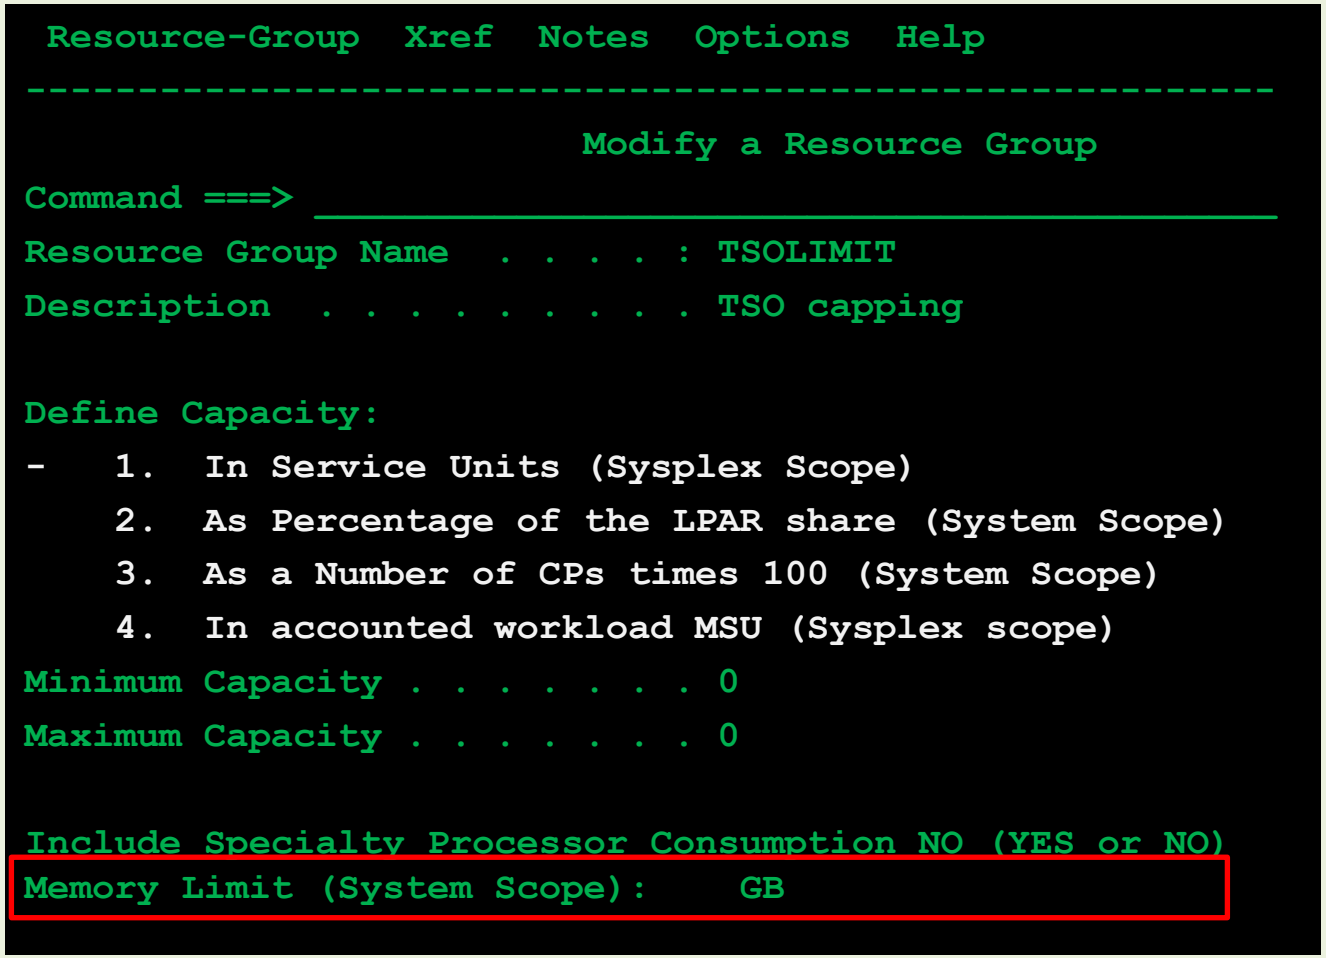

Introduction

- Even if quiesced, looping tasks running on non swappable AS may continue to use a lot of CPU
- They can be stopped by using a KILLIT resource group

Introduction

- A mono-period service class with a discretionary goal has to be defined
- It has to be associated to a resource group

```
 Service-Class  Xref  Notes  Options  Help
------------------------------------------------------------------------------
                          Modify a Service Class                    Row 1
Command ===>

Service Class Name . . . . . : KILLIT
Description  . . . . . . . . . Service class to block a task
Workload Name  . . . . . . . . USTC      (name or ?)
Base Resource Group  . . . . . KILLIT    (name or ?)
Cpu Critical . . . . . . . . . NO        (YES or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

        -- Period --  ------------------ Goal ------------------
Action  #  Duration    Imp.  Description
  __    _  _________   _     ____________________________________
  __    1  _________   _     Discretionary
******************************* Bottom of data ********************************
```

Introduction

- This is a type 1 resource group
- With a maximum capacity limit of 1 service unit per second

```
  Resource-Group  Xref  Notes  Options  Help
------------------------------------------------------------------------------
                          Modify a Resource Group
Command ===> ____________________________________________________________

Enter or change the following information:

Resource Group Name  . . . . : KILLIT
Description  . . . . . . . . . Resource group to block a task_

Define Capacity:
1   1.  In Service Units (Sysplex Scope)
    2.  As Percentage of the LPAR share (System Scope)
    3.  As a Number of CPs times 100 (System Scope)
Minimum Capacity . . . . . . . 1_______
Maximum Capacity . . . . . . . 1_______

Memory Limit (System Scope)    ________  GB
```

Introduction

- All the considerations about resource groups also apply to Tenant Resource Groups with the following exceptions:
  - you can't set a minimum capacity limit
  - you can't set a memory limit
- More details on Tenant Resource Groups in the last session of this training course

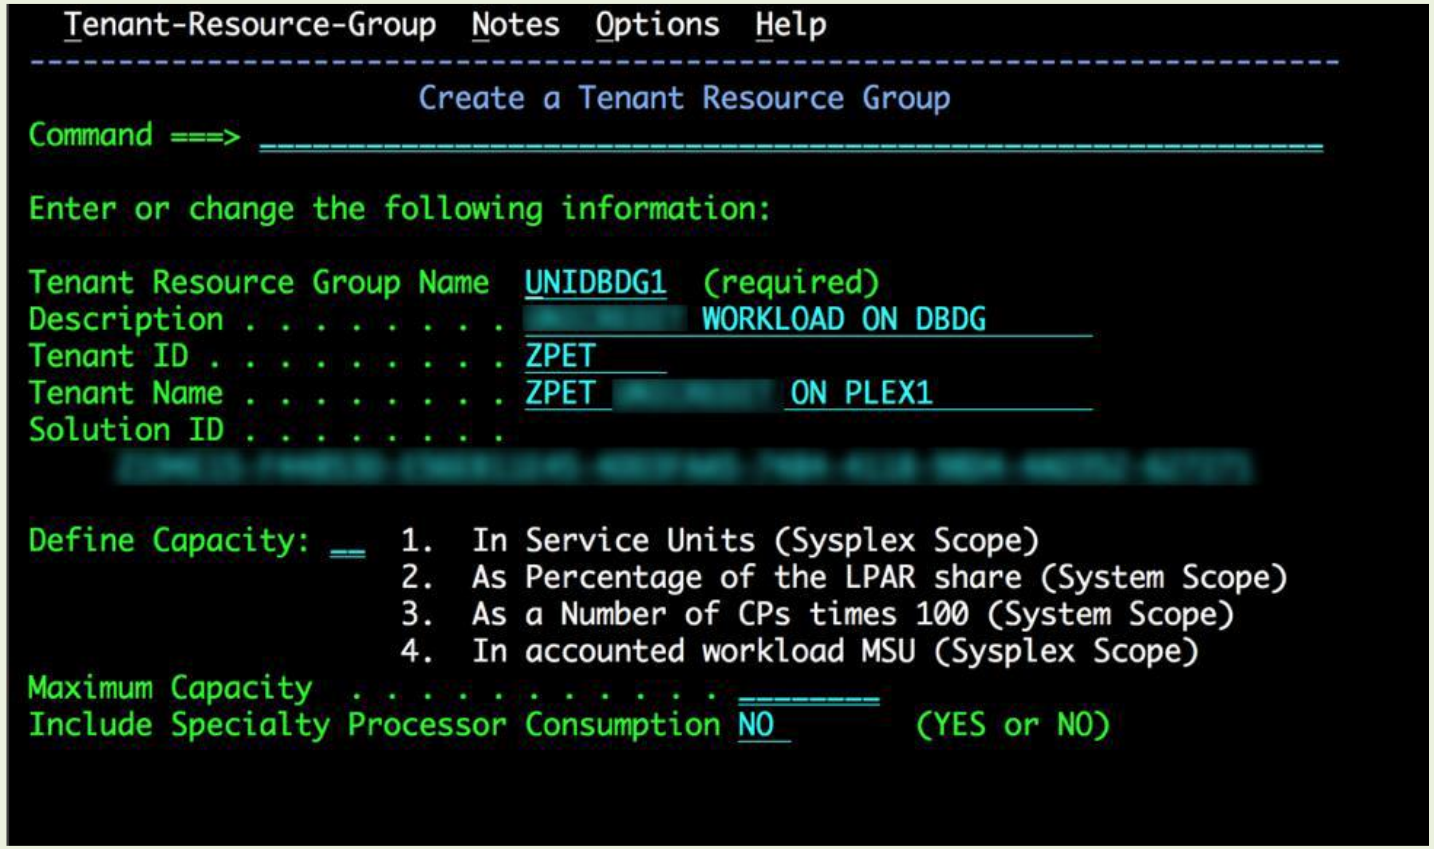

Exploiting resource groups

Exploiting resource groups

- In the next slides we show the effects on the different resource group types when migrating an LPAR from a 2964-751 (50.549 MIPS) to a 8561-739 (51.602 MIPS)
- Our goal is to assign the resource group the same maximum capacity as before

Exploiting resource groups –Type 1

- Type 1 resource groups have been available for many years but they are not often used
- Type 1 maximum capacity characteristics:
  - not easy to understand; the starting point is the number of service units per second (SU/sec) of one CPU
  - SU/sec depends on the number of CPUs assigned to the LPAR (not considering zAAP, zIIP and HiperDispatch);
  - it has a sysplex scope; it controls how much service the workloads receive in total, but not how much service is received on each system

Exploiting resource groups –Type 1

- The maximum capacity of a Type 1 resource group has been set to 59.480 CPU service units per second
- This is the system service units rate of 1 CPU when 19 logical CPUs are assigned to a LPAR hosted on an IBM 2964-751 machine
- The resulting limit in MIPS is about 991 MIPS

```
 Resource-Group  Xref  Notes  Options  Help
------------------------------------------------------------
                          Modify a Resource Group
Command ===> ____________________________________________________
Resource Group Name  . . . . . : TSOLIMIT
Description  . . . . . . . . . . TSO capping

Define Capacity:
1    1.  In Service Units (Sysplex Scope)
     2.  As Percentage of the LPAR share (System Scope)
     3.  As a Number of CPs times 100 (System Scope)
     4.  In accounted workload MSU (Sysplex scope)
Minimum Capacity . . . . . . . . 0
Maximum Capacity . . . . . . . . 59480

Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope):     GB
```

Exploiting resource groups –Type 1

- The following formula can be used to set the limit of a Type 1 resource group:
  - Maximum capacity = Max CPU * System SU rate
  - Max CPU = Maximum capacity / System SU rate
- If you want to set a limit of 1 CPUs
  - Maximum capacity = 1,0 CPU * 59.480 SU/sec = 59.480 SU/sec
- If you want to set a maximum capacity of 59.480 SU/sec
  - Max CPU = 59.480 SU/sec / 59.480 SU/sec/CPU = 1,0 CPU

Exploiting resource groups –Type 1

- The maximum capacity in MIPS, corresponding to Max CP, can be calculated by using the following formula:
  - Max MIPS = CEC MIPS / CEC CPUs * Max CP
  - Max MIPS = 50.549 MIPS / 51 CPU * 1 CPU = 991 MIPS

Exploiting resource groups –Type 1

- Exercise 1

In the new machine the production LPAR has now 15 logical CPUs assigned instead of 19; the system SU rate of the new machine is higher than the old one (79.208 vs 59.480) because of the more powerful processors used.

We want to keep the maximum capacity of the TSOLIMIT resource group fixed at 59.480

Calculate the maximum number of usable CP (Max CP) and MIPS (Max MIPS) in the new machine

Max CP =

Max MIPS =

Exploiting resource groups –Type 1

- Exercise 1

In the new machine the production LPAR has now 15 logical CPUs assigned instead of 19 while the system SU rate of the new machine is higher than the old one (79.208 vs 59.480) because of the more powerful processors used.

We want to keep the maximum capacity of the TSOLIMIT resource group fixed at 59.480

Calculate the maximum number of usable CP (Max CP) and MIPS (Max MIPS) in the new machine

Max CP = **59.480 SU/sec / 79.208 SU/sec/CPU = 0,75 CPU**

Max MIPS = **0,75 CPU * 51.602 MIPS / 39 CPU = 994 MIPS**

Exploiting resource groups –Type 2

- Type 2 maximum capacity is a percentage of the LPAR share (%Weight); easier to understand than Type 1
- Type 2 maximum capacity characteristics:
  - if for any reason the LPAR share changes (e.g. IRD, soft capping) the capacity allowed will change accordingly
  - It has a system scope; it controls how much service the workloads receive on each system; so if in a two systems sysplex one LPAR has twice the share of the other it will get double service
  - when configuration or hardware changes, it will provide the same percentage of the LPAR share but a different capacity

Exploiting resource groups –Type 2

- While the LPAR share concept is more straightforward than “unweighted service units”, it is not as simple as it may look
- The LPAR share can in fact change as a consequence of both manual actions, such as weight modifications and LPARs activation/deactivation, and automatic actions such as IRD weight management and soft capping, due to defined and group capacity limits
- In all these situations the capacity allowed to type 2 resource groups will change accordingly

Exploiting resource groups –Type 2

LPAR	Active	Weight	%Weight	Type 2 max	Type 2 capacity
LPAA	YES	600	60%	20%	12%
LPAB	YES	300	30%		
LPAC	YES	100	10%		
TOTAL		1.000	100%		

LPAR	Active	Weight	%Weight	Type 2 max	Type 2 capacity
LPAA	YES	600	86%	20%	17%
LPAB	NO				
LPAC	YES	100	14%		
TOTAL		700	100%		

Exploiting resource groups –Type 2

- Old machine capacity is 50.549 MIPS
- LPAR share is 32,7%
- The maximum capacity for the resource group has been set to 6%

```
 Resource-Group  Xref  Notes  Options  Help
------------------------------------------------------------
                          Modify a Resource Group
Command ===> ____________________________________________________
Resource Group Name  . . . . . : TSOLIMIT
Description  . . . . . . . . . . TSO capping

Define Capacity:
2    1.  In Service Units (Sysplex Scope)
     2.  As Percentage of the LPAR share (System Scope)
     3.  As a Number of CPs times 100 (System Scope)
     4.  In accounted workload MSU (Sysplex scope)
Minimum Capacity . . . . . . . . 0
Maximum Capacity . . . . . . . . 6

Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope):     GB
```

Exploiting resource groups –Type 2

- The maximum capacity in MIPS on the old machine can be calculated by using the following formula:
  - Max MIPS = CEC MIPS * LPAR share * Type 2 limit
  - Max MIPS = 50.549 MIPS * 0,327 * 0,06 = 992 MIPS
- If the LPAR share and the maximum capacity limit remain the same, the maximum capacity in MIPS on the new machine is:
  - Max MIPS = 51.602 MIPS * 0,327 * 0,06 = 1012 MIPS

Exploiting resource groups –Type 2

- This is a lucky case because the capacity of the old and new machine is very close
- Generally speaking even if the LPAR share remains the same, when the hardware changes you have to review maximum capacity value
- To set a maximum capacity value corresponding to the same amount of MIPS allowed in the old machine you should use the following formula
- New machine limit =

  old maximum limit in MIPS / new machine CEC MIPS / LPAR share

Exploiting resource groups –Type 2

- Exercise 2

In the new machine the production LPAR has now 40% share.

Calculate the maximum capacity for the TSOLIMIT resource group so that the maximum capacity in MIPS is still about 992 MIPS.

Maximum capacity =

Exploiting resource groups –Type 2

- Exercise 2

In the new machine the production LPAR has now 40% share.

Calculate the maximum capacity for the TSOLIMIT resource group so that the maximum capacity in MIPS is still about 992 MIPS.

Maximum capacity = **992 MIPS / 51.602 MIPS / 0,40 = 5%**

Exploiting resource groups –Type 3

- Type 3 maximum capacity is a percentage of one CP capacity multiplied by 100
- Type 3 maximum capacity characteristics:
  - easy to set;
  - It has a system scope; it controls how much service the workloads receive on each system; so if in a two systems sysplex one LPAR uses CPs twice powerful than the other it will get double service
  - when configuration or hardware changes, it will provide the same percentage of the CP but a different capacity

Exploiting resource groups –Type 3

- Old machine capacity is 50.549 MIPS
- The maximum capacity of a Type 3 resource group has been set to 100% of 1 CPU

```
 Resource-Group  Xref  Notes  Options  Help
------------------------------------------------------------
                          Modify a Resource Group
Command ===> ____________________________________________________
Resource Group Name  . . . . . : TSOLIMIT
Description  . . . . . . . . . . TSO capping

Define Capacity:
3   1.  In Service Units (Sysplex Scope)
    2.  As Percentage of the LPAR share (System Scope)
    3.  As a Number of CPs times 100 (System Scope)
    4.  In accounted workload MSU (Sysplex scope)
Minimum Capacity . . . . . . . . 0
Maximum Capacity . . . . . . . . 100

Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope):    GB
```

Exploiting resource groups –Type 3

- The maximum capacity in MIPS on the old machine can be calculated by using the following formula:
  - Max MIPS = **CEC MIPS / CEC CPUs * Type 3 limit**
  - Max MIPS = **50.549 MIPS / 51 * 1 = 991 MIPS**
- If the maximum capacity limit remains the same the maximum capacity in MIPS on the new machine will become:
  - Max MIPS = **51.602 MIPS / 39 * 1 = 1.323 MIPS**

Exploiting resource groups –Type 3

- Exercise 3

Calculate the maximum capacity for the TSOLIMIT resource group so that the maximum capacity in MIPS is still about 991 MIPS in the new machine.

Maximum capacity =

Exploiting resource groups –Type 3

- Exercise 3

Calculate the maximum capacity for the TSOLIMIT resource group so that the maximum capacity in MIPS is still about 991 MIPS in the new machine.

Maximum capacity = **991 MIPS / 51.602 MIPS * 39 = 75%**

Exploiting resource groups –Type 3

- You have to review type 3 limits when the hardware changes

Exploiting resource groups –Type 4

- Type 4 maximum capacity is the amount of MSUs allowed
- Type 4 maximum capacity characteristics:
  - easy to set;
  - it has a Sysplex wide scope;
  - when configuration or hardware changes, it will provide the same capacity

Exploiting resource groups –Type 4

- Old machine capacity is 50.549 MIPS and 5.927 MSUs
- The maximum capacity of a Type 4 resource group has been set to 116 MSUs
- It corresponds to 991 MIPS

```
 Resource-Group  Xref  Notes  Options  Help
------------------------------------------------------------
                          Modify a Resource Group
Command ===> ____________________________________________________
Resource Group Name  . . . . . : TSOLIMIT
Description  . . . . . . . . . . TSO capping

Define Capacity:
4    1.  In Service Units (Sysplex Scope)
     2.  As Percentage of the LPAR share (System Scope)
     3.  As a Number of CPs times 100 (System Scope)
     4.  In accounted workload MSU (Sysplex scope)
Minimum Capacity . . . . . . . . 0
Maximum Capacity . . . . . . . . 116

Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope):     GB
```

Exploiting resource groups –Type 4

- Exercise 4

Calculate the maximum capacity in MSUs corresponding to 991 MIPS

MSU maximum capacity =

MIPS maximum capacity / CEC MIPS * CEC MSUs =

991 MIPS / 50.549 MIPS * 5927 MSUs = 116 MSUs

Measuring resource groups

Measuring resource groups

- All resource group measurements are provided in SMF 72
- To collect them you need the WKLD parameter in RMF settings (default)

Measuring resource groups

- Two type 2 and two type 3 resource groups defined
- No limit to memory, zIIP consumptions not included

WLM RESOURCE GROUPS

WLM RESOURCE GROUPS - Wed, 30 Sep 2020 - PROD

SYSTEM	SYSPLEX	SERVICE DEFNAME	SERVICE POLICY	RESOURCE GROUP	TYPE	SCOPE	MIN	MAX	UNIT	MEMLIM (GB)	ZIIP_INC
PROD	PLX	WLME 00	WLMPOL00	JESGROUP	3	SYSTEM	0	140	% CP	0	N
PROD	PLX	WLME 00	WLMPOL00	SASGROUP	2	SYSTEM	0	10	% LPAR	0	N
PROD	PLX	WLME 00	WLMPOL00	TSOGROUP	2	SYSTEM	0	10	% LPAR	0	N
PROD	PLX	WLME 00	WLMPOL00	TWSGROUP	3	SYSTEM	0	250	% CP	0	N

Measuring resource groups

- Assignment of service classes to resource groups

- Wed, 30 Sep 2020 - PROD

SERVICE POLICY	SC	PERIOD	DUR	IMP	GOALTYPE	GOAL	PERC	WKL NAME	RESOURCE GROUP
WLMPOL00	ALLTSO	1	200,0	2	PCT	1,000	90,0	TSO	TSOGROUP
WLMPOL00	ALLTSO	2	300,0	3	PCT	5,000	90,0	TSO	TSOGROUP
WLMPOL00	ALLTSO	3	8.000,0	4	PCT	30,000	90,0	TSO	TSOGROUP
WLMPOL00	ALLTSO	4	0,0	5	VEL	20,0	0,0	TSO	TSOGROUP
WLMPOL00	BATCHIGH	1	0,0	3	VEL	30,0	0,0	BATCH	
WLMPOL00	BATCHLOW	1	100.000,0	4	VEL	30,0	0,0	BATCH	JESGROUP
WLMPOL00	BATCHLOW	2	0,0	5	VEL	20,0	0,0	BATCH	JESGROUP
WLMPOL00	BATCHOPC	1	100.000,0	4	VEL	50,0	0,0	BATCH	TWSGROUP
WLMPOL00	BATCHOPC	2	0,0	5	VEL	40,0	0,0	BATCH	TWSGROUP

Measuring resource groups

- Resource groups intensively used in the day hours
- Essentially to limit batch workload

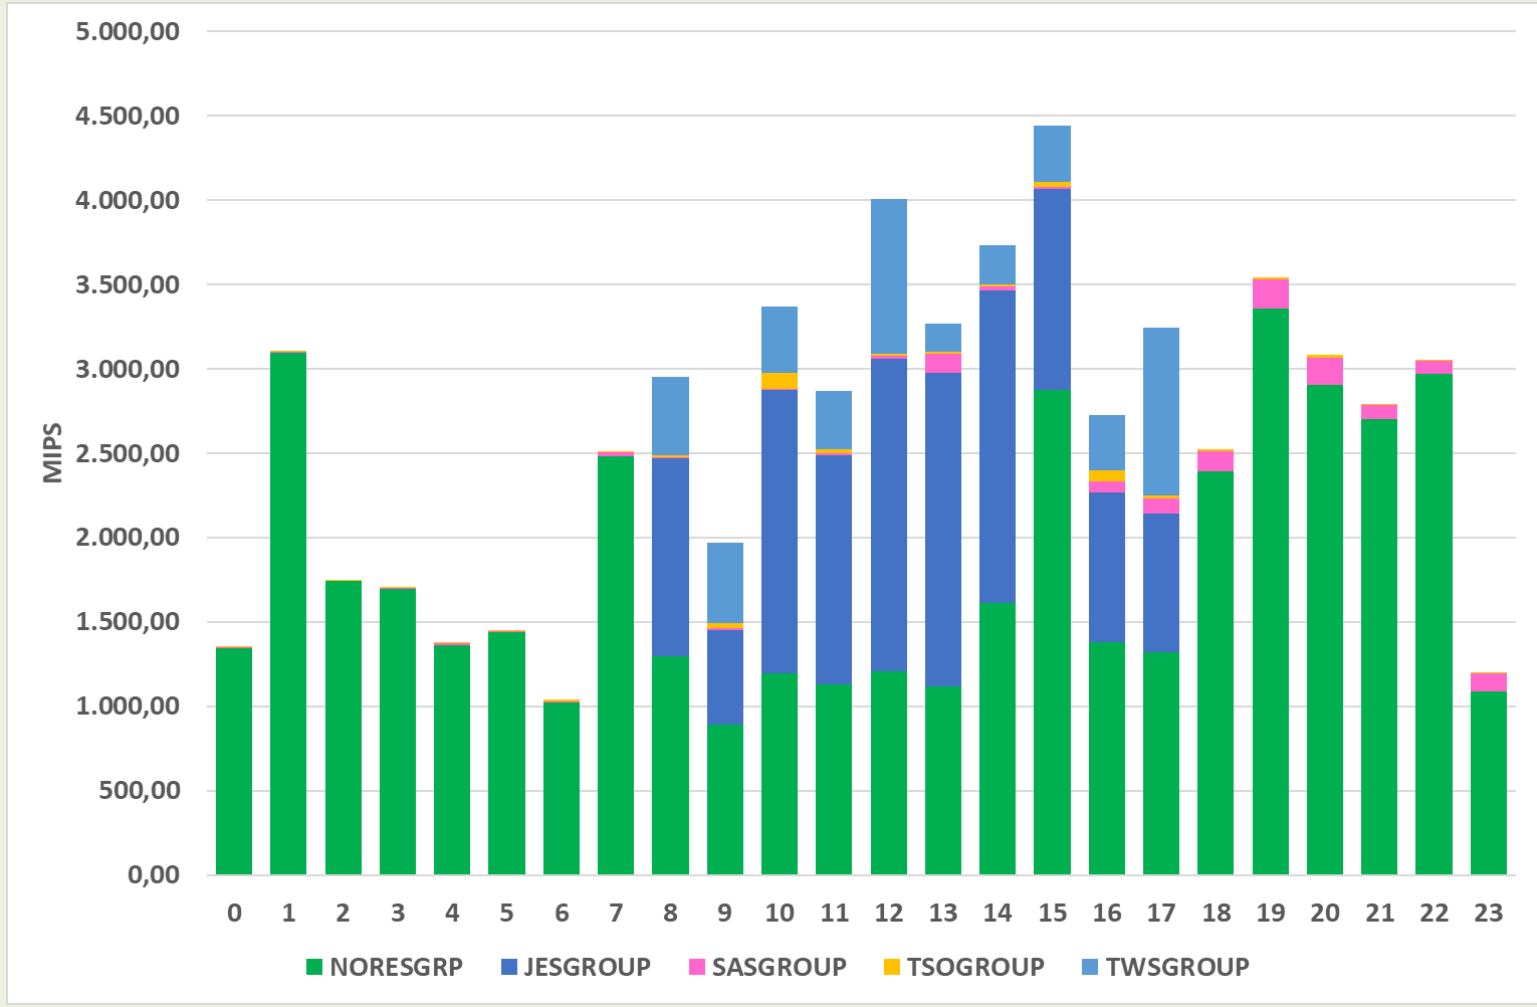

Measuring resource groups

- Where the graph flattens the maximum capacity limit has been reached
- Type 3 maximum capacity limit = 140
- Single CPU capacity is about 1320 MIPS

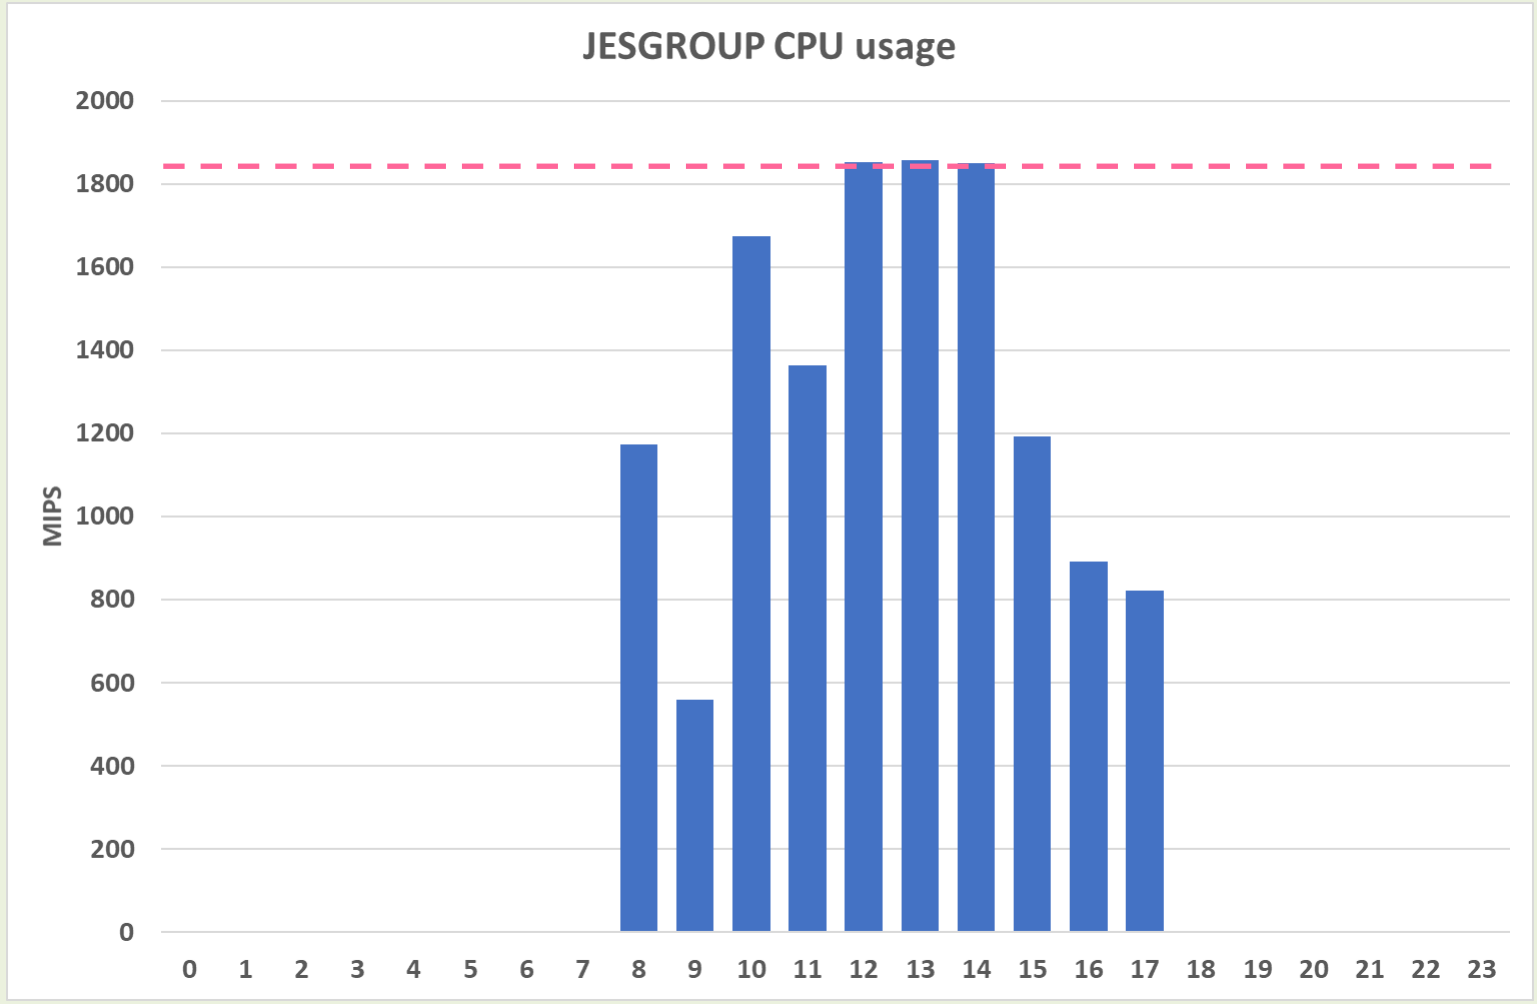

Measuring resource groups

- CPU capping delay samples can be used to evaluate the impact of resource group capping on the involved service classes
- CPU capping delays due to the JESGROUP max limit are measured and reported, for all the service classes associated to that resource group
- Only one service class period is presented in the next slide

Measuring resource groups

SC GENERAL DELAYS

SC GENERAL DELAYS - Wed, 30 Sep 2020 - PROD - BATCHLOW - PER=1 IMP=4 DUR=100000 GTYPE=VEL GOAL=30,000

SYSTEM	STATES	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
PROD	UNKNOWN DELAY	.	.	.	.	.	.	.	.	55,71	37,67	34,17	19,77	42,51	35,41	33,38	55,76	45,46	35,60	.	.	.	.	.	.
PROD	IDLE	.	.	.	.	.	.	.	.	23,91	0,47	4,22	0,47	3,83	13,02	2,49	4,60	4,29	1,59	.	.	.	.	.	.
PROD	TOTAL USING	.	.	.	.	.	.	.	.	8,90	32,48	25,36	12,13	33,63	28,38	52,04	29,70	32,09	36,46	.	.	.	.	.	.
PROD	TOTAL DELAY	.	.	.	.	.	.	.	.	11,48	29,39	36,25	67,63	20,03	23,19	12,09	9,93	18,17	26,35	.	.	.	.	.	.
PROD	VALID SAMPLES < 200																								
	PERCENT USING STATES																								
PROD	CPU USING	.	.	.	.	.	.	.	.	1,86	5,09	4,79	1,92	5,99	4,08	4,84	4,49	3,46	6,39	.	.	.	.	.	.
PROD	I/O USING	.	.	.	.	.	.	.	.	7,05	26,96	20,54	10,21	27,63	24,31	47,20	25,22	28,53	30,08	.	.	.	.	.	.
PROD	IIP USING	.	.	.	.	.	.	.	.	0,00	0,44	0,03	0,00	0,00	0,00	0,00	0,00	0,10	0,00	.	.	.	.	.	.
	PERCENT DELAY STATES																								
PROD	CPU DELAY	.	.	.	.	.	.	.	.	0,25	0,90	0,76	0,22	1,48	1,12	1,49	1,82	1,27	1,23	.	.	.	.	.	.
PROD	DISK DELAY	.	.	.	.	.	.	.	.	0,43	2,38	2,57	1,49	4,00	2,40	5,26	3,10	1,41	2,58	.	.	.	.	.	.
PROD	IIP DELAY	.	.	.	.	.	.	.	.	0,00	0,10	0,00	0,00	0,00	0,00	0,00	0,00	0,00	0,00	.	.	.	.	.	.
PROD	CPU CAPPING DELAY	.	.	.	.	.	.	.	.	0,12	0,00	0,92	0,28	0,34	0,08	3,27	1,60	0,00	0,55	.	.	.	.	.	.

Questions?