

Capacity Provisioning (CP) for z/OS V2R1 can manage Defined Capacity and Group Capacity limits. It can automatically adjust them based on a predefined policy as well as provides the end user with commands to make the adjustments from a z/OS console rather than the Support Element (SE).

In this blog entry we will share with you our experiences with this new enhancement. From here on we will refer to Defined Capacity as DC and Group Capacity as GC.

Specifically we will talk about:

- High level introduction to DC and GC concepts
- Steps we took to start exploiting DC and GC
- CP for z/OS policy updates we made to let it start managing our DC and GC limits
- CP for z/OS DC support in action
- Links to additional documentation

High level introduction to DC and GC concepts

In various other documentations Defined Capacity and Group Capacity are also referred to as Soft Capping or WLM Capping.

An installation can specify a Defined Capacity expressed in millions of service units per hour (MSU) for an LPAR. The Defined Capacity sets the capacity limit of an individual LPAR when soft capping is in effect.

WLM keeps a 4-hour rolling average of the CPU usage of the LPAR, and when the 4-hour average CPU consumption exceeds the defined capacity limit, WLM dynamically activates soft capping. When the rolling 4-hour average returns below the Defined Capacity, soft capping is removed.

The Group Capacity allows the definition of a group of LPARs on the same CPC and a limit for the combined capacity usage by those LPARs. This allows the system to manage the group in such a way that the limit for Group Capacity in MSU per hour will not be exceeded. PR/SM and WLM work together to enforce the capacity defined for the group and also enforce the capacity optionally defined for each individual LPAR.

An LPAR can be restricted by Defined Capacity and in addition belong to a capacity group.

For more details on Defined Capacity and Group Capacity concepts please refer to the following documentation:

z/OS Performance: WLM Soft Capping Support for Sub Capacity Pricing

- <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10099>

MVS: Planning: Workload Management

- <http://publibz.boulder.ibm.com/epubs/pdf/iea3w100.pdf>

Steps we took to start exploiting DC and GC

You need to have DC and GC limits in place before CP for z/OS can start managing them. In other words today CP for z/OS doesn't support initial assignment of these limits but just managing them once they are assigned.

Please note that the DC and GC limits... etc used in this blog entry are by no means recommendations for you. These are just values from one of our various tests. Limits you assign will depend on your configuration and requires capacity & performance planning.

For our test, we decided to assign DC and GC to one of our zEC12 CPCs, referred to as P91 in the examples below. These changes were done using the SE.

First we assigned DC to a set of our LPARs using the *Change Logical Partition Controls* task:

Logical Partition	Active	Defined Capacity	WLM	Current Weight	Initial Weight	Min Weight	Max Weight	Current Capping	Initial Capping	Absolute Capping	Number of Dedicated Processors	Number of Not dedicated Processors
JB0	Yes	1035	<input checked="" type="checkbox"/>	320	300			No	<input type="checkbox"/>	None	0	61
JD0	Yes	690	<input checked="" type="checkbox"/>	180	200			No	<input type="checkbox"/>	None	0	61
JF0	Yes	518	<input checked="" type="checkbox"/>	150	150			No	<input type="checkbox"/>	None	0	61
JI0	Yes	518	<input checked="" type="checkbox"/>	150	150			No	<input type="checkbox"/>	None	0	61
JJ0	Yes	518	<input checked="" type="checkbox"/>	150	150			No	<input type="checkbox"/>	None	0	61
JK0	No	0	<input checked="" type="checkbox"/>	0	150			No	<input type="checkbox"/>	None	0	60
J90	Yes	1035	<input checked="" type="checkbox"/>	300	300			No	<input type="checkbox"/>	None	0	61
P91AWARE	Yes	0	<input type="checkbox"/>	100	100			No	<input checked="" type="checkbox"/>	None	0	2
P91AWARE2	No	0	<input type="checkbox"/>	0	100			No	<input checked="" type="checkbox"/>	None	0	2
Z2	Yes	0	<input checked="" type="checkbox"/>	142	150			No	<input type="checkbox"/>	None	0	61
Z4	Yes	0	<input checked="" type="checkbox"/>	158	150			No	<input type="checkbox"/>	None	0	61
PETLVS	Yes	0	<input type="checkbox"/>	10	10			No	<input checked="" type="checkbox"/>	None	0	4
ZG5	Yes	0	<input type="checkbox"/>	75	75			No	<input checked="" type="checkbox"/>	None	0	8

Next we created an LPAR Group, called PLEX1, using the *Change LPAR Group Controls* task. We added the above mentioned LPARs to it and gave it a GC of 4313.

Change LPAR Group Controls - P91

Edit ▾

Input/Output configuration data set (IOCDS): a0

Group Name	Member Partitions	Group Capacity Value
DEFAULT		0
TESTPLEX		600
PLEX1	JB0 JD0 JF0 JI0 JJ0 J90	4313
PLEX2		1306

Save to Profiles Change Running System Save and Change Reset Cancel Help

Note that these changes can be made dynamically to the running systems and/or saved to the LPAR profiles to take effect in the future. We made the changes dynamically using the *Save and Change* button, shown above, and were able to see the results in RMF right away. Take a peek at the *Image Capacity, Group* and *Limit* fields for system JB0 from the *RMF Monitor III CPC Capacity* panel below.

Samples: 118		System: JB0		Date: 12/04/13		Time: 15.55.00		Range: 120	
Partition: JB0		2827 Model 758							
CPC Capacity: 5986		weight % of Max: ****		4h Avg: 856		Group: PLEX1			
Image Capacity: 1035		WLM Capping %: 0.0		4h Max: 1010		Limit: 4313			
Partition	MSU	Def	Act	Cap	Proc	Logical	Util %	Physical	Util %
					Num	Effect	Total	LPAR	Effect
*CP					490			0.7	37.6
JB0	1035		802	NO	58.0	13.3	13.4	0.1	13.3
JD0	690		279	NO	60.0	4.5	4.5	0.0	4.6
JF0	518		365	NO	60.0	5.8	5.9	0.1	6.0
JI0	518		228	NO	57.0	3.8	3.9	0.0	3.8
JJ0	518		311	NO	59.0	5.1	5.1	0.0	5.2
J90	1035		0	NO	60.0	0.0	0.0	0.0	0.0

Also even though we hadn't made any CP for z/OS changes yet we noticed a few messages in it's job log indicating that it was aware of the above changes:

```
05.45.16 S0022149 CPO3987I Group capacity observed. Current capacity is 4313 MSU for capacity
group PLEX1 of CPC P91
05.45.16 S0022149 CPO3961I Capacity base is 4313 MSU for capacity group PLEX1 of CPC
P91
05.49.17 S0022149 CPO3984I Defined capacity change detected. New defined capacity is 1035
MSU for LPAR JB0 of CPC P91 with system JB0 in sysplex
UTCPLXJ8
05.49.17 S0022149 CPO3960I Defined capacity base is 1035 MSU for LPAR JB0 of CPC P91 with
system JB0 in sysplex UTCPLXJ8
...
```

Since CP for z/OS was already aware of the DC and GC updates we decided to take a look at some of the reports.

DEFINEDCAPACITY report

```
F CPOSERV,APPL=R DC PLEX=UTCPLXJ8 SYS=JB0
```

OR

```
F CPOSERV,APPL=R DC CPC=P91 LPAR=JB0
```

CP01095I Defined capacity report generated at 08/05/2013 10:51:26

Defined capacity for system JB0 in sysplex UTCPLXJ8

```
CPC.LPAR:                P91.JB0
Sysplex.System:          UTCPLXJ8.JB0
Management state:        not increased
Policy limit:             0 additional MSU
Management base:         1035 MSU
Managed capacity:        0 additional MSU
Current capacity:         1035 MSU
Remaining time until capping: 240 minutes
4 hour rolling average consumption: 856 MSU
End of report
```

GROUPCAPACITY report

```
F CPOSERV,APPL=R GC CPC=P91 GROUP=PLEX1
```

OR

```
F CPOSERV,APPL=R GC PLEX=UTCPLXJ8 SYS=JB0
```

CP01096I Group capacity report generated at 08/05/2013 10:53:24

Group capacity for group PLEX1 on CPC P91

```
Group name:              PLEX1
CPC name:                P91
Management state:        not increased
Policy limit:             0 additional MSU
Management base:         4313 MSU
Managed capacity:        0 additional MSU
Current capacity:         4313 MSU
End of report
```

CP for z/OS policy updates we made to let it start managing our DC and GC limits

Before we describe the updates we made let us familiarize you with our CP for z/OS environment. Our domain consists of a z10, z196 and a zEC12 with 16 z/OS LPARs distributed across them and managed

by CP for z/OS. We have three rules within our policy.

Actions

	Name Filter	Description Filter	Default Status Filter
<div><div></div><div></div></div>	EveningRule	6 PM to 6 AM	<div><div></div>Enabled</div>
<div><div></div><div></div></div>	PrimeRule	6 AM to 6 PM	<div><div></div>Enabled</div>
<div><div></div><div></div></div>	WeekendRule	Weekend 24x7	<div><div></div>Enabled</div>

Each rule has a set of recurring time conditions as well as workload conditions and take effect only when both set of conditions are met. For instance the recurring time condition of our Evening rule is specified such that it is only active on week days between 6:00 PM and 6:00 AM.

Nonrecurring Time Conditions			Recurring Time Conditions							Workload Conditions			
Actions ▾													
	Name <i>Filter</i>	Start Date <i>Filter</i>	End Date <i>Filter</i>	Mon <i>Filter</i>	Tue <i>Filter</i>	Wed <i>Filter</i>	Thu <i>Filter</i>	Fri <i>Filter</i>	Sat <i>Filter</i>	Sun <i>Filter</i>	Start Time <i>Filter</i>	Deadline <i>Filter</i>	End Time <i>Filter</i>
<input type="checkbox"/>	Evening	Jul 25, 2011	Jul 30, 2029	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	6:00 PM	5:30 AM	6:00 AM

The workload condition portion of our Evening rule is setup such that if the system Performance Index (PI) for any of our workloads (regardless of priority) exceeds 1.5 for 5 minutes additional resources that are available in our On/Of CoD record will be activated. Similarly once the system PI is 1.1 or less for 5 minutes unnecessary resources will be deactivated.

Importance Filters		Included Service Classes		Excluded Service Classes		
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Actions ▾				
	Importance Filter	Provisioning PI Filter	Provisioning Duration (Minutes) Filter	Deprovisioning PI Filter	Deprovisioning Duration (Minutes) Filter	PI Scope Filter
<input type="checkbox"/>	1	1.5	5	1.1	5	System
<input type="checkbox"/>	2	1.5	5	1.1	5	System
<input type="checkbox"/>	3	1.5	5	1.1	5	System
<input type="checkbox"/>	4	1.5	5	1.1	5	System
<input type="checkbox"/>	5	1.5	5	1.1	5	System

Now let's go back to our discussion on DC and GC support. First, per system, we added a *Maximum Defined Capacity Scope* to our policy. This value represents the total amount of DC MSU that can be added to that system by all the rules within the policy.

Maximum Processor Scope		Logical Processor Scope		Maximum Defined Capacity Scope		Maximum Group Capacity Scope		Rules	
Actions									
	System Filter	Sysplex Filter	Max. Increase (MSU) Filter		Primary Increment (MSU) Filter		Secondary Increments (MSU) Filter		
<input type="checkbox"/>	J90	UTCPLXJ8	5198		500		100		
<input type="checkbox"/>	J80	UTCPLXJ8	5198		500		100		
<input type="checkbox"/>	J00	UTCPLXJ8	5543		500		100		
<input type="checkbox"/>	JF0	UTCPLXJ8	5715		500		100		
<input type="checkbox"/>	J10	UTCPLXJ8	5715		500		100		
<input type="checkbox"/>	JJ0	UTCPLXJ8	5715		500		100		

Next we assigned a *Defined Capacity Scope (Max. Increase MSU)*, per system, in all our rules within the policy. This value limits the DC MSU that can be added by the specific rule.

Processor Scope		Defined Capacity Scope		Group Capacity Scope	
<div><div><input checked="" type="checkbox"/></div><div><input type="checkbox"/></div><div>Actions ▾</div></div>					
	System Filter	Sysplex Filter		Max. Increase (MSU) Filter	
<input type="checkbox"/>	J90	UTCPLXJ8		5198	
<input type="checkbox"/>	JB0	UTCPLXJ8		5198	
<input checked="" type="checkbox"/>	JD0	UTCPLXJ8		5543	
<input type="checkbox"/>	JF0	UTCPLXJ8		5715	
<input type="checkbox"/>	JJ0	UTCPLXJ8		5715	
<input type="checkbox"/>	JJ0	UTCPLXJ8		5715	

We repeated the same actions for GC. First we added a *Maximum Group Capacity Scope* to our policy.

Maximum Processor Scope		Logical Processor Scope		Maximum Defined Capacity Scope		Maximum Group Capacity Scope		Rules	
<input checked="" type="checkbox"/>		<input type="checkbox"/>		Actions ▾					
Group Filter		CPC Filter		Max. Increase (MSU) Filter		Primary Increment (MSU) Filter		Secondary Increments (MSU) Filter	
<input type="checkbox"/> PLEX1		P91		1920		500		100	

Then we updated all our rules with a *Group Capacity Scope (Max. Increase MSU)*.

Processor Scope		Defined Capacity Scope	Group Capacity Scope
Actions			
	Group Filter	CPC Filter	Max. Increase (MSU) Filter
<input checked="" type="checkbox"/>	PLEX1	P91	1920

Finally we installed and activated our policy!

CP for z/OS DC support in action

Here are some sample messages related to CP for z/OS DC and GC support.

When CP for z/OS starts you will notice messages, per system and per LPAR group, on whether or not if it is managing DC and GC for those systems.

```

11.50.02 S0031098 CPO3986I Defined capacity observed. Current
capacity is 1518 MSU for LPAR JF0 of CPC P91 with system JF0 in
sysplex UTCPLXJ8
11.50.02 S0031098 CPO3960I Defined capacity base is 1018 MSU for
LPAR JF0 of CPC P91 with system JF0 in sysplex UTCPLXJ8

11.50.02 S0031098 CPO3987I Group capacity observed. Current
capacity is 4313 MSU for capacity group PLEX1 of CPC P91
11.50.02 S0031098 CPO3961I Capacity base is 4313 MSU for capacity
group PLEX1 of CPC P91

11.50.03 S0031098 CPO3966I Defined capacity is turned off for
LPAR TPN of CPC H91 with
    793                system TPN in sysplex UTCPLXJ8. Defined
capacity management for this
    793                system stops

```

As CP for z/OS adjusts DC and GC you will see messages similar to these:

```

12.05.02 S0031098 CPO3964I Defined capacity decrease initiated to
1018 MSU for LPAR JF0 of CPC P91 with system JF0 in sysplex
UTCPLXJ8
12.06.13 S0031098 CPO3984I Defined capacity change detected. New
defined capacity is 443
    443                1018 MSU for LPAR JF0 of CPC P91 with system
JF0 in sysplex UTCPLXJ8

```

Links to additional documentation

For more details on CP for z/OS please refer to: MVS Capacity Provisioning User's Guide

- <http://publibz.boulder.ibm.com/epubs/pdf/iea3u100.pdf>

Make sure to take a look at their web site as well. I find the *What's New* and *FAQs* sections to be very handy. Also if you are new to z/OS capacity provisioning take a look at the *Further Info* section for some introductory presentations.

- <http://www-03.ibm.com/systems/z/os/zos/features/cpm/>