



VWLC

Part 3: How does GCL work

- Technical terms and abbreviations used in this document.
 - MSU (Million Service Units) : reference unit for Software billing.
 - IMSU : instantaneous consumption of MSU.
 - R4H : Rolling 4 Hours.
 - DC : Defined Capacity.
 - GCL: Group Capacity Limit.
 - CPC Capacity : full capacity of the System z server.
 - Billing space : MSU Value of consumption used for the bill (SCRT).
 - White Space : CPC Capacity – Billing Space.
 - ShareW%: This is the %SHARE calculated from the WEIGHT of the LPAR.
 - ShareW-GCL%: This is the %SHARE calculated from the WEIGHT in the GCL of the LPAR.
 - ShareD%: This is the %SHARE calculated from the DEFINED CAPACITY of the LPAR.
 - Target MSU from Weight: ShareW% translated in MSU.
 - Target MSU from GCL-Weight: ShareW-GCL% translated in MSU.

ASC / GCL : Group Capacity Limit

- Enhancements to the IBM System z9 with the LPAR Group Capacity Limit feature:
- **GROUP CAPACITY LIMIT with z9 + μ code & zOS1.8**
- **Solution:** the LPAR group capacity limit feature is designed to allow you to specify a capacity limit for each LPAR group, making management easier.
- **Capacity Group**
 - Consists in multiple LPARs on the same CEC
 - LPARs must run on z/OS 1.8
 - If they do not, the group limit may not be implemented correctly
 - It is possible to define more than one group on a CEC
 - A partition can only belong to one group
 - A capacity group is independent of a sysplex and a LPAR cluster
 - WLM will only manage partitions with shared CPs
 - Dedicated partitions and partitions with weight completion

ASC / GCL : Group Capacity Limit

LPAR Group Capacity Limit

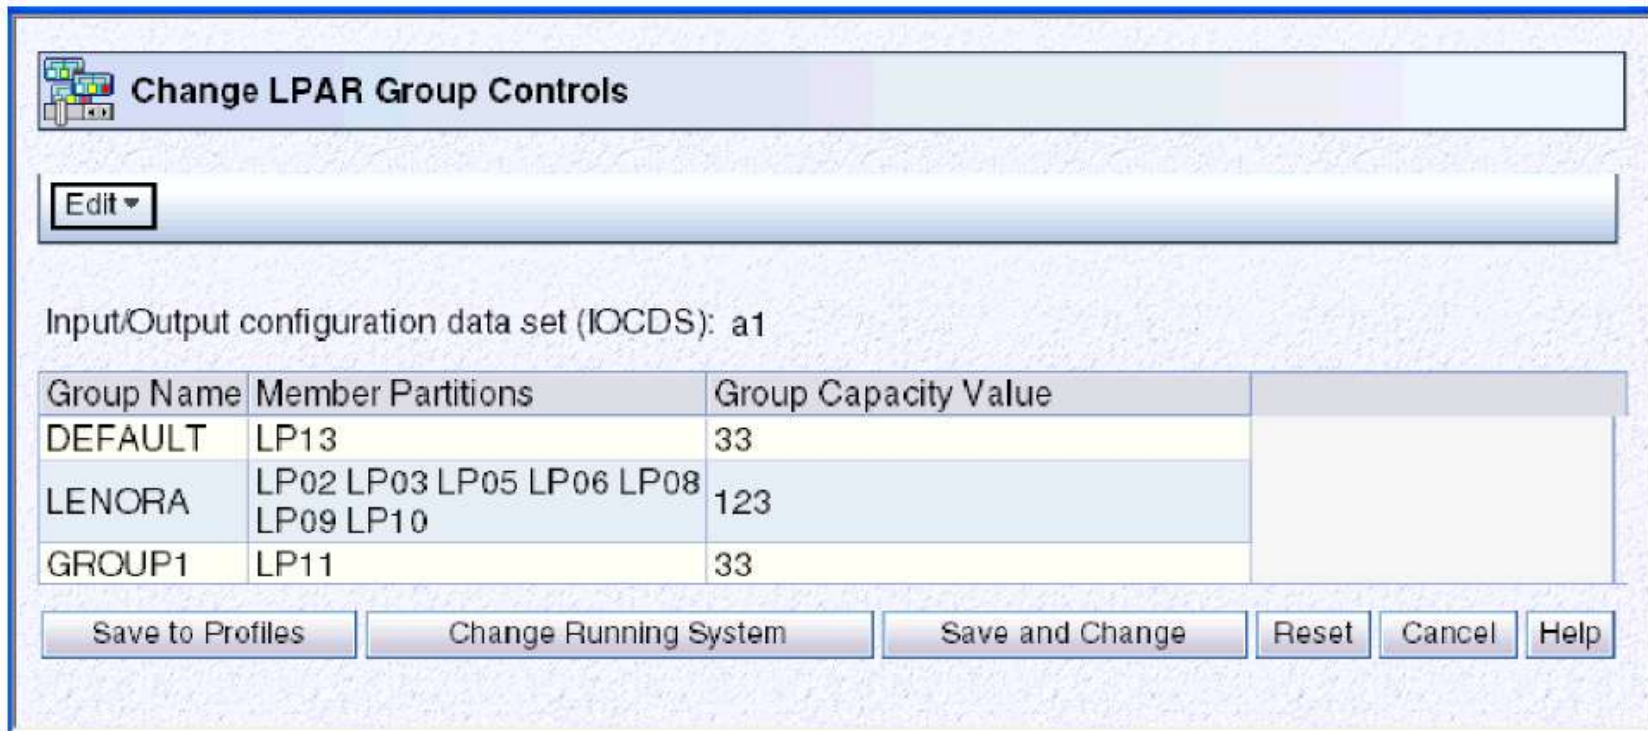
With the logical partition Group Capacity Limit on z9 EC and z9 BC, you can now specify LPAR group capacity limits allowing you to define each LPAR with its own capacity and one or more groups of LPARs on a server.

This is designed to allow z/OS to manage the groups in such a way that the sum of the LPARs' CPU utilization within a group will not exceed the group's defined capacity. Each LPAR in a group can still optionally continue to define an individual LPAR capacity limit.

LPAR group capacity limit requires that all LPARs managed in the group are running at z/OS or z/OS.e V1.8 or later. LPAR group capacity limits may help provision a portion of a System z9 server to a group of LPARs allowing the CPU resources to float more readily between those LPARs, resulting in more productive use of "white space" and higher server utilization.

How to define a GCL ?

- Support GROUP CAPACITY LIMIT – HMC)



Change LPAR Group Controls

Edit ▾

Input/Output configuration data set (IOCDS): a1

Group Name	Member Partitions	Group Capacity Value
DEFAULT	LP13	33
LENORA	LP02 LP03 LP05 LP06 LP08 LP09 LP10	123
GROUP1	LP11	33

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

GCL : Technical function → single sample


LPAR Name GCL	Weight	SHAREW% global	GCL Group	SHAREW% GCL	Target MSU based on Weight GCL	MSU limit from Weight	MSU Limit from #LP	Number of LP
PROD	600	60,00%	GCL1	60,00%	262	319	532	8
PREPROD	200	20,00%	GCL1	20,00%	87	106	466	7
DEV	180	18,00%	GCL1	18,00%	79	96	399	6
SYST	20	2,00%	CGL1	2,00%	9	11	200	3
Total	1000				437	532		

Capacity 532


GCLimit 437

Total Weight of GCL 1000

ASC / GCL : Technical function → single sample



When the sum of the LPARs' R4Hs exceeds the GCLimit, then all the LPARs are capped (WLM Capping% in pink) on their weight or target MSU (below – static): the sum of the LPARs' IMSUs is brought back to the sum of the LPARs' DCs.



VWLC with IBM GCL: Pros & Cons

- VWLC with IBM GCL +:
 - IBM GCL is easy to setup on the machine.
 - It allows you to control the bill at a chosen level (with fixing GCLimit).
 - IBM GCL is a technical enhancement to facilitate Soft-Capping management.

- VWLC with IBM GCL - :
 - The weight is statically defined with a linear value which is rarely the right value at the right time when WLM asks PR/SM to cap the GROUP.
 - Initially, the weight guarantees a quantity of power when the CPC runs at 100%.
 - The user often sets a higher GCLimit than needed in order to avoid potential performance problems (capping).
 - A loop in one LPAR can disturb all the LPARs by generating WLM Capping.