

# Container Pricing for IBM Z

## The WLM View

06 Dec 2018

—

Horst Sinram  
STSM z/OS Workload Management and Capacity Management  
[sinram@de.ibm.com](mailto:sinram@de.ibm.com)

# IBM Z



Central Europe  
Computer Measurement Group

# Trademarks

## The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

## The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* All other products may be trademarks or registered trademarks of their respective companies.

### Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

# Agenda

- **Container Pricing Rationale**
- **WLM Tenant Resource Group functions**
  - **Definitions**
  - **Usage**
- **Memory management enhancements**
  - **Material by Tobias Orth**

# Container Pricing for IBM Z

## Announcement excerpt:

IBM is introducing **Container Pricing for IBM Z** for qualified solutions running on IBM z13 and z14 servers. Container Pricing will provide simplified software pricing for qualified solutions, combining flexible deployment options with competitive economics that are directly relevant to those solutions.

Container Pricing can scale from collocated solutions within existing LPARs, through to separate LPARs, up to multiple-LPAR solutions, without directly impacting the cost of unrelated workloads. Additionally, Container Pricing will simplify pricing and billing on the IBM Z platform, by superseding a number of existing price offerings and by fully automating the billing process.

IBM initially announces three solutions that will be enabled with Container Pricing:

- The **New Application Solution** will provide a highly competitive stand-alone priced offering for new z/OS applications, such as CICS® TS or WebSphere applications. The New Application Solution is the strategic replacement for the current zCAP and IWP priced offerings.
- The **Application Development and Test Solution** will provide highly competitive stand-alone pricing for z/OS based development and test workloads. Modern DevOps tooling can be optionally added at uniquely discounted prices.
- The **Payments Solution** will provide a "per payment" pricing option for IBM Financial Transaction Manager for z/OS deployments. This new offering directly ties operational cost to business value by basing the price on the number of payments processed, rather than capacity used to process them.

**For more information, see Whitepaper**


**<https://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102719>**

**or**

**<https://www-03.ibm.com/systems/z/resources/swprice/container.html>, and [New Application Solution introduces consumption-based pricing for new IBM z/OS-based applications](#)**

# Infrastructure for Container Pricing – the Basic Idea

- Identify a subset of the work running in a system in the WLM service definition
- WLM automatically gathers consumption data (4HRA, service units) for that subset
- SCRT then subtracts this 4HRA from the system 4HRA (which determines your software charges), and prices it separately and/or differently (discount, free of charge, other metric like number of transactions, etc.)



# Why two new WLM service definition objects?

- Ubiquitous cloud workload paradigm asks for new ways of metering workloads in multi-tenant environments
- IBM Z business asks for an infrastructure to support novel pricing options

- A Tenant Report Class (TRC) is similar to a WLM Report Class. TRCs are assigned through WLM classification and are always associated with a Tenant Resource Group.

- A Tenant Resource Group (TRG) is somewhat similar to a WLM Resource Group and can be associated with tenants or solutions.
- TRGs aggregate consumption data and can optionally be used to apply consumption limits.

WLM TRG technology is intended for collocated Container Pricing solutions, and tenant implementations.

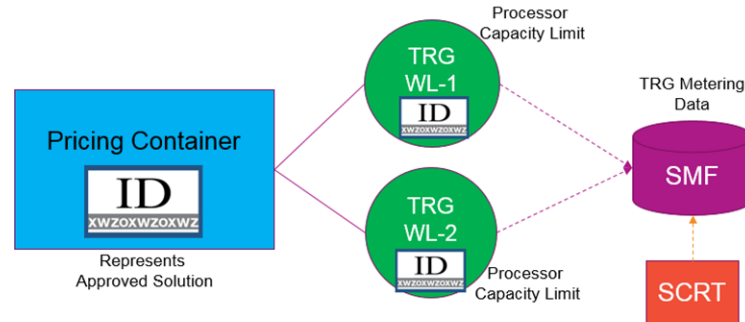
They are not used for non-collocated Container Pricing (LPAR level) solutions.

# Agenda

- **Container Pricing Rationale**
- **WLM Tenant Resource Group functions**
  - **Definitions**
  - **Usage**
- **Memory management enhancements**
  - **Material by Tobias Orth**

# Container Pricing workflow overview

- For an eligible solution, IBM provides you with a “Solution ID” (key)
- In your WLM service definition you
  - Define one or more TRGs and paste the Solution ID into the definition
    - Dummy solution IDs for test/education are documented at [www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSL03543USEN](http://www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSL03543USEN)
  - Define one or more TRCs associated with TRG
  - Change or add classification rules to classify eligible work (only!) and assign service class and TRCs
  - Install and activate WLM service definition
- Monitors query WLM and write new SMF70 data sections for TRGs
  - In addition, the TRC and TRG data is reported via the existing report class and resource group mechanism in the WLMGL report (SMF72.3)
- SCRT consumes SMF70 and SMF89 data for billing
  - Verifies solution ID, applies pricing rules





# Overview of WLM/SRM Enhancements for Container Pricing for IBM Z

- New WLM service definition panels (z/OSMF and ISPF)
  - For defining and modifying Tenant Resource Groups
  - For defining and modifying Tenant Report Classes
- Resource Group (RG) and Tenant Resource Group (TRG) enhancements
  - New Sysplex-wide Type 4 limit expressed at a scale of “MSU”
  - Optionally, speciality processor consumption can be counted towards the limit
- Enhanced WLM/SRM programming services
  - A new IWM4QTNT service allows monitoring products to retrieve TRG-level consumption data
  - Various APIs provide TRC and TRG indications
- A new service definition option that allows to disable “Discretionary Goal Management” globally
- Summary of WLM changes in <http://publibz.boulder.ibm.com/zoslib/pdf/OA52312.pdf>
- WLM TRG level data is reported by monitoring products, such as RMF
  - Including reporting in SMF type 70 record
- SMF70 TRG data can be consumed in SCRT
  - Depending on business rules (solution ID), SCRT can reduce the system 4HRA by TRG level 4HRA partially or entirely. Optionally, other charge models may apply.
  - Product usage information (SMF89) may be evaluated
- Various other components are updated (SMF, SDSF, z/OSMF, ...)
  - z/OSMF [Workflows](#) simplify solution implementation
  - Comprehensive Knowledge Center [Content Collection](#)

# Enablement of Container Pricing for IBM Z

<b>z/OS release</b> <b>Function</b>	z/OS V2.3	z/OS V2.2	z/OS V2.1
WLM	<b>OA52312</b> <b>OA54590</b>	<b>OA52312</b> <b>OA54590</b>	Coexistence only: <b>OA52312</b>
RMF	<b>OA52694</b>	<b>OA52694</b>	
z/OSMF WLM	<b>PI82551</b>	<b>PI82551</b>	
z/OSMF RMF	<b>PI82546</b>	<b>PI82546</b>	
SMF	<b>OA53033</b>	<b>OA53033</b>	
SCRT and Billing System Support	<b>OA53047</b>	<b>OA53047</b>	

Container Pricing  
 FIXCAT category:  
 Keyword:

[Content collection \(Knowledge center\)](#)  
 IBM.Function.PricingInfrastructure  
 PRICINGINFR/K

# Tenant Resource Group (TRG) and Tenant Report Class (TRC) Definition

- The WLM Administrative Application Level increases to 32.
- Tenant Resource Groups and Tenant Report Classes can be defined via new menu items.
- Specification of these new objects will increase the functionality level of the service definition to 32.

```
File Utilities Notes Options Help
-----
Functionality LEVEL032          Definition Menu          WLM Appl LEVEL032
Command ===> -----

Definition data set . . . : 'WLM.DEMO.SERVDEF.XML'

Definition name . . . . . PROD01      (Required)
Description . . . . . Production service definition

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 Mgmt Provider
```

# Tenant Resource Group (TRG) Definition

- The TRG name is mandatory (8 char)
- Description, Tenant ID, Tenant Name are optional (expected to be used in a z/OS cloud context)
- For qualified offerings, a 64 char Solution ID needs to be provided.
  - Exactly enter (paste) the IBM provided Solution ID string; WLM performs sanity check only. Solution IDs failing that check are rejected.
  - Attributes encoded into the ID may change how the system processes the work.
  - Solution ID is acted upon during SCRT processing.
  - Multiple TRGs may specify same Solution ID
- Optionally, a consumption limit can be specified.
  - TRG capacity limits should not be specified unless there is a need to limit processor consumption.
  - The “Include Specialty Processor Consumption” switch indicates whether the combined CP and specialty processor consumption determines the cap limit.
- Unlike standard resource groups there is no minimum consumption limit

```

Tenant-Resource-Group  Notes  Options  Help
-----
                                Create a Tenant Resource Group
Command ==> -----

Enter or change the following information:

Tenant Resource Group Name TRGDEM01 (required)
Description . . . . . Sample TRG
Tenant ID . . . . . -----
Tenant Name . . . . . Solution Newapp
Solution ID . . . . .
                                Z194E15-F1078F4-CEBBF9F075-099853BF-60EF-4A05-A0D1-EF925B-992C90

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)

Maximum Capacity . . . . . -----
Include Specialty Processor Consumption NO (YES or NO)

Memory Limit (System Scope) . . . . . 12 GB
    
```

- The new type of capacity limit, and the switch to include speciality processor consumption are also available for standard resource groups

# Tenant Report Class (TRC) Definition

- The TRC name is mandatory (8 char)
  - Name must be unique (also across report classes)
  - Up to 2047 Report Classes and Tenant Report Classes can be defined
- The TRG name is required, i.e. any TRC must be associated with a TRG
- Monitoring interfaces and monitors report on TRC as on standard report classes

```
Tenant-Report-Class  Notes  Options  Help
-----
                                Create a Tenant Report Class
Command ==> _____

Enter or change the following information:

Tenant Report Class Name . . . T_CDC      (Required)
Description . . . . . TRC including CDC region

Tenant Resource Group Name . . TRGDEM01  (Required; name or ?)
```

# Using Tenant Report Classes (TRCs) in classification rules

- Work units are assigned a service class for performance management and optionally a Report Class or Tenant Report Class for reporting and optional capping
- The Report Class or Tenant Report Class specified must have already been created. Entering “?” allows to view and select from the list of defined Report Classes and Tenant Report Classes.
- The use of Report Classes and Tenant Report Classes in a classification rule is mutually exclusive

```

Subsystem-Type  Xref  Notes  Options  Help
-----
                                Modify Rules for the Subsystem Type          Row 1 to 6 of 6
Command ==>  _____ Scroll ==> CSR

Subsystem Type . : STC          Fold qualifier names?  Y (Y or N)
Description . . . Use Modify to enter YOUR rules

Action codes:  A=After      C=Copy      M=Move      I=Insert rule
               B=Before    D=Delete row R=Repeat  IS=Insert Sub-rule
                                   More ==>

          -----Qualifier-----          -----Class-----
Action  Type      Name      Start          Service      Report
-----  ---
          1 TN      TOR      ___          VEL80      VELR80
          1 TN      AOR1     ___          VEL50      VELR50
          1 TN      AOR2     ___          VEL50      VELR50
          1 TN      TOR1     ___          VEL60      VELR60
          1 TN      STC*     ___          SYSSTC     _____
          1 TN      CDC*     ___          VEL80      T_CDC

***** BOTTOM OF DATA *****

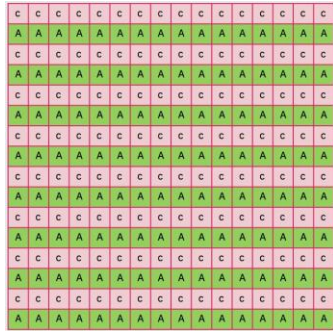
```

# Using Tenant Report Classes in classification rules (cont.)

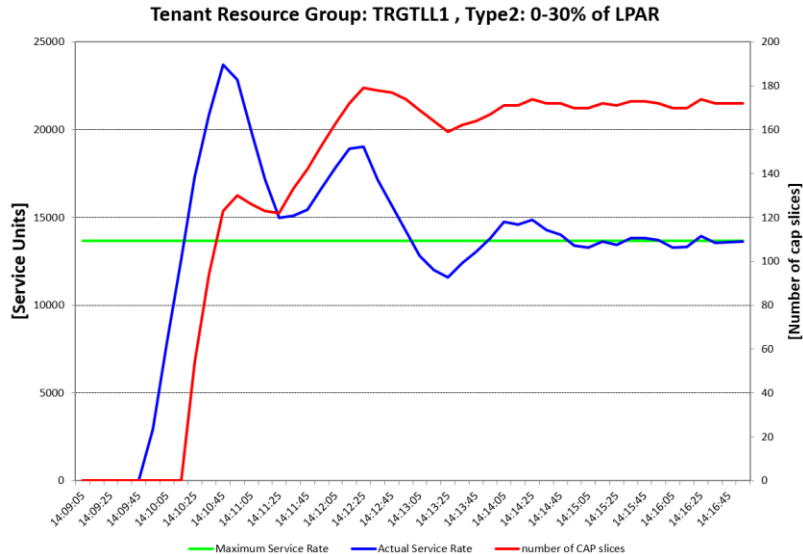
## Restrictions:

- Non-default values of the Reporting Attribute (MOBILE, CATEGORYA or CATEGORYB) must not be used with a Tenant Report Classes in the same classification rule
- **A Tenant Report Class must not be used with a service class that is associated with a Resource Group in the same classification rule**
- If a Tenant Report Class is used in classification rules that assign different service classes, the Tenant Report Class might become heterogeneous
  - This means that work may run in different service classes but is reported altogether in this one Tenant Report Class
  - For heterogeneous Tenant Report Classes, reporting products might provide less meaningful data than for homogeneous Tenant Report Classes
  - WLM tolerates that, but issues a warning panels and messages, e.g., IWMAM916W Tenant Report Class T\_CDC might become heterogeneous by combining work running in service classes VEL80 VEL50
  - **Strong recommendation is to use only homogeneous TRCs: Create different Tenant Report Classes for each service class, and connect them all to the same Tenant Resource Group, instead**

# TRG Capping



Tenant Resource Group Overview



TRG capping is based on Resource Group capping and inherits its characteristics:

- Time is divided into 256 “slices”. In any slice the whole (T)RG can be set
  - dispatchable (called awake slice)
  - non-dispatchable (cap slice).
- The cap pattern is adjusted every 10 sec based on the average of the last minute
- (T)RG consumption will vary based on demand, mix of dispatch priorities, number of dispatchable units and number of processors.
  - The system will attempt to over-cap the work, i.e., the consumption will be throttled to remain below the limit. Depending on the characteristics this may not always be possible.
  - Usually consumption stabilizes within minutes
- Very latency sensitive work is not a good candidate for capping. Multiple TRGs with same Solution ID may be used when needed.



# Comparison of TRG and RG Capping Types

Type 1	Type 2	Type 3	Type 4
Raw CPU+SRB service units (“Raw” meaning that Service Definition Coefficients are not applied)	Percent of CP LPAR share (even if specialty processor consumption included). May exceed 100%.	Percent of one CP processor (even if specialty processor consumption included)	Processor consumption expressed in “accounted workload MSU”
Limit applies to Sysplex	Limit applies to each System	Limit applies to each system	Limit applies to Sysplex
For all (T)RG types only captured TCB and SRB times are counted towards the limit. The limit is enforced based on a one minute average (i.e., no 4HRA). Up to 32 RGs plus 32 TRGs may be defined.			

# RG and TRG “MSU” limits

- Background: Technical and pricing related performance/capacity numbers are based on different views. This remains unchanged.
  - The adjustment factor for service units (technical view) is based on the logical configuration (number of CPs online to the LPAR) on the respective CPC. Refer to [Processor version codes and SRM constants](#)
  - The adjustment factor for pricing purposes (MSU) is based on the physical configuration (CPC model capacity rating). Refer to [Large Systems Performance Reference for IBM Z](#)
- Every 10 sec, WLM converts type 2, 3, and 4 limits into a SU/sec service rate based on current configuration
- Therefore, a type 4 (MSU) limit is converted using the CEC and LPAR adjustment factors
- A type 4 (MSU) limit is intended to simplify the specification of a limit expressed in MSU, **but...**
- ... it cannot be expected that RG and TRG MSU limits will closely match the resulting LPAR MSU consumption or 4HRA TRG consumption:
  - The (T)RG limit applies only to the accounted (captured) TCB and SRB times. System management time (uncaptured time) is not included
  - The limit is not a 4 hour rolling average, it is managed to a short interval (~ 1 min)

# Use of Include Specialty Processor Consumption

- In the past, resource group limits were only based on general purpose processor consumption
- Now, if “Include Specialty Processor Consumption=YES” is specified, also the specialty processor consumption counts towards the limit
  - Available in RG and TRG definitions
- Primary use case:
  - When used with a service class that specifies “Honor Priority = No” it is possible to limit such work to a specific consumption level.
  - Fine grain control for SPARK or Java batch workloads

# When to use resource groups or tenant resource groups

## **Only or preferentially use standard resource groups...**

- When a resource group minimum is required
- When it is required to use WLM managed initiators as part for the RG

## **Only or preferentially use tenant resource groups...**

- For authorized pricing container solutions
- **For all functions as long as the TRGs is only used for aggregation (i.e. no limits)**

# WLM Topology and Management of Limits

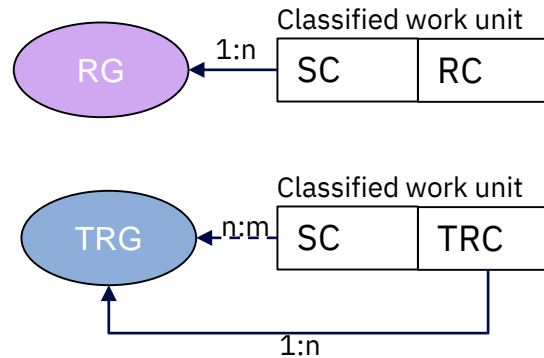
## WLM management

- Granularity is and remains service class period
- With TRGs a service class period may become “heterogenous”  
n:m associated to TRGs  
(irrelevant when the TRG does not specify limits)

## Processor and memory limits

As always


- Processor management is based on the transaction, and therefore transaction SC and TRC
- Memory management is based on the address space, and therefore address space SC and TRC



# RMF Monitoring

- TRG related overview conditions for SMF70.1 data
  - See [CPU Activity - SMF record type 70-1](#)
  - Includes 4HRA condition TRGLACS
- As an example, these overview statements will report on several consumption related metrics by 60 min intervals

```
OVW(4HRA(LACS))  
OVW(TRG4HRA(TRGLACS(UNITRG)))  
OVW(TRGCPSU(TRGCP(UNITRG)))  
OVW(INTVMSU(LACTMSU(LP01)))  
OVW(TOTAPPL(APPLPER(POLICY)))  
OVW(UNIAPPL(APPLPER(R.UNITRC.1)))  
DINTV(C060)
```



4HRA	TRG4HRA	TRGCPSU	INTVMSU	TOTAPPL	UNIAPPL
635	351	578134	1717	1444.7	857.5
1064	613	588761	1764	1486.9	872.0
1487	869	576245	1742	1466.9	853.9
1660	1003	508639	1461	1217.5	753.2
1652	808	135730	1666	1366.9	871.9
1646	546	0	1744	1468.7	
1641	482	0	1670	1401.5	

# RMF Monitoring (cont.)

- TRC level data on WLMGL report as for standard report classes
- TRG level data as for standard resource groups on the WLMGL policy page:

```
RESOURCE GROUPS
--NAME--  -----DESCRIPTION-----  -SYSTEM-  ---CPU CONSUMPTION---  -----CPU CAPACITY-----  ----MEMORY----
          #CPS  MSU  SU/SEC  MIN  MAX  DEFINED AS  USAGE  LIMIT
UNITRG   UNI WORKLOAD ON LP01
          8.53  980  576K
          LP01  8.53  980  576K  354M
          -----REPORT CLASSES  UNITRC  8.53  980  576K
```

# New service definition option to deactivate Discretionary Goal Management

- Discretionary Goal Management (DGM) can improve throughput of discretionary work by throttling certain vastly over-achieving non-discretionary work.
- It “throttles” through dynamically created resource groups.
- The effect on non-discretionary workload can be unexpected and many installation try to avoid DGM by defining dummy resource groups, or particular goals.
- With this support, DGM can optionally be disabled.

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

Enter or change the Service Coefficients:

CPU . . . . . 1.0      (0.1-99.9)
IOC . . . . . 0.0      (0.0-99.9)
MSO . . . . . 0.0000   (0.0000-99.9999)
SRB . . . . . 1.0      (0.0-99.9)

Enter or change the service definition options:

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



# Agenda

- **Container Pricing Rationale**
- **WLM Tenant Resource Group functions**
  - **Definitions**
  - **Usage**
- **Memory management enhancements**
  - **Material by Tobias Orth**

# How to limit memory consumption on z/OS

## **Pre-existing (hard limits for virtual storage):**

- Region size on the job or exec statement in JCL
- System-wide memory limit for **all address spaces** via MEMLIMIT parameter in parmlib SMFPRMxx or IEFUSI exit
- System-wide memory limit for **different classes of work** (for example, JOB, TSO, STC) via MEMLIMIT parameter in parmlib SMFLIMxx or IEALIMIT exit


## **New (limit for real storage):**

- Limit real storage for address spaces associated with a WLM resource group with a memory limit

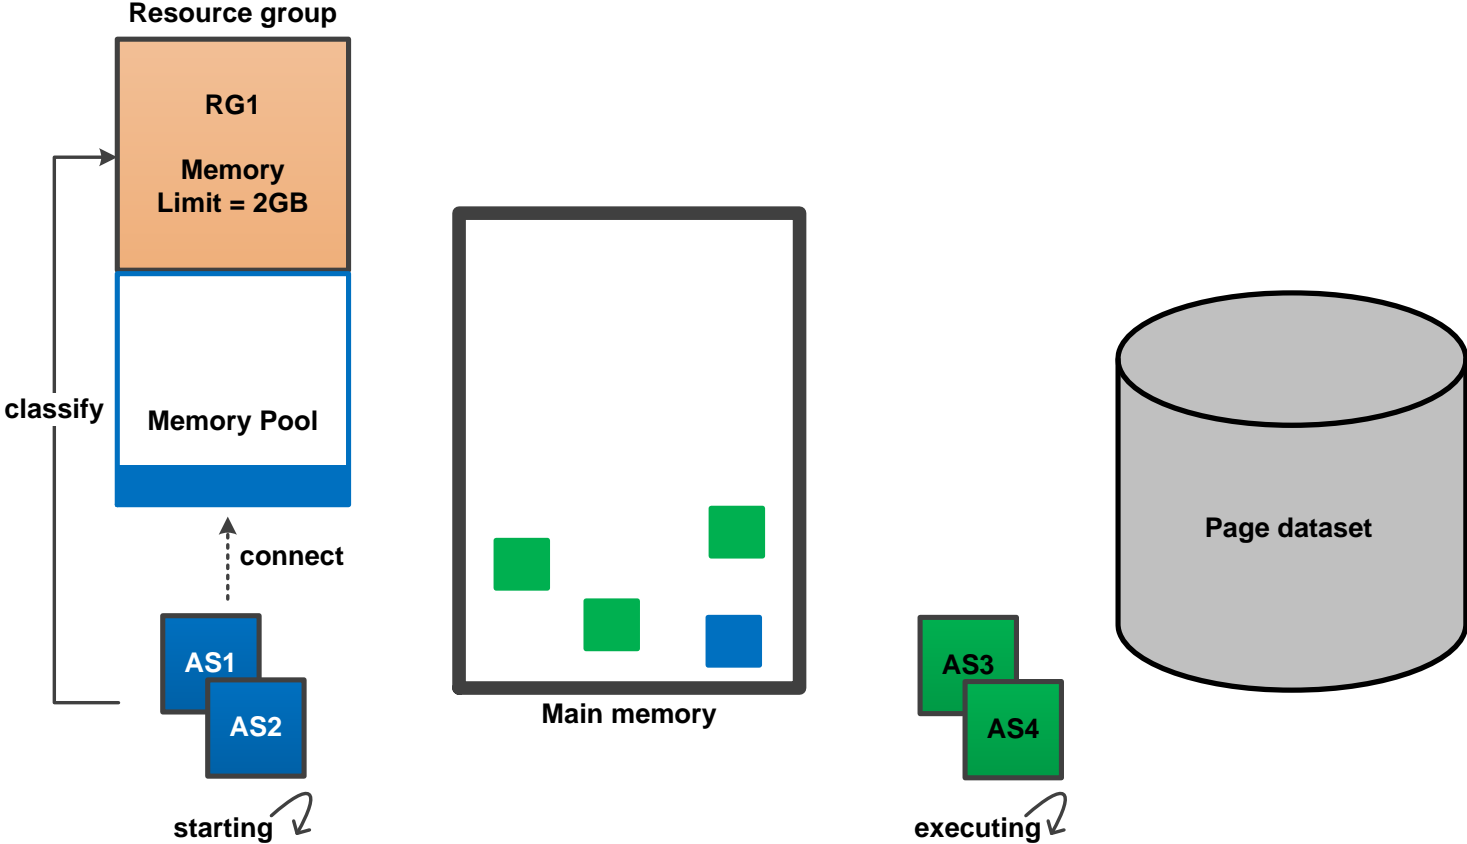
# z/OS memory pools

- z/OS memory pool – **restrict real storage usage by applications:**
  - Defined by the (WLM) admin and instantiated during WLM policy activation
  - Address spaces of an application connect to the pool at initialization time
  - Traditional storage interface are used to obtain and free storage
  - When pool storage is exhausted, page stealing within the pool is started
  - System enforces the limit so jobs in the pool may compete for storage
  
- **Memory pools aim to protect important workload outside a pool by keeping workloads with abnormal or unexpected storage demands within their pool limit**

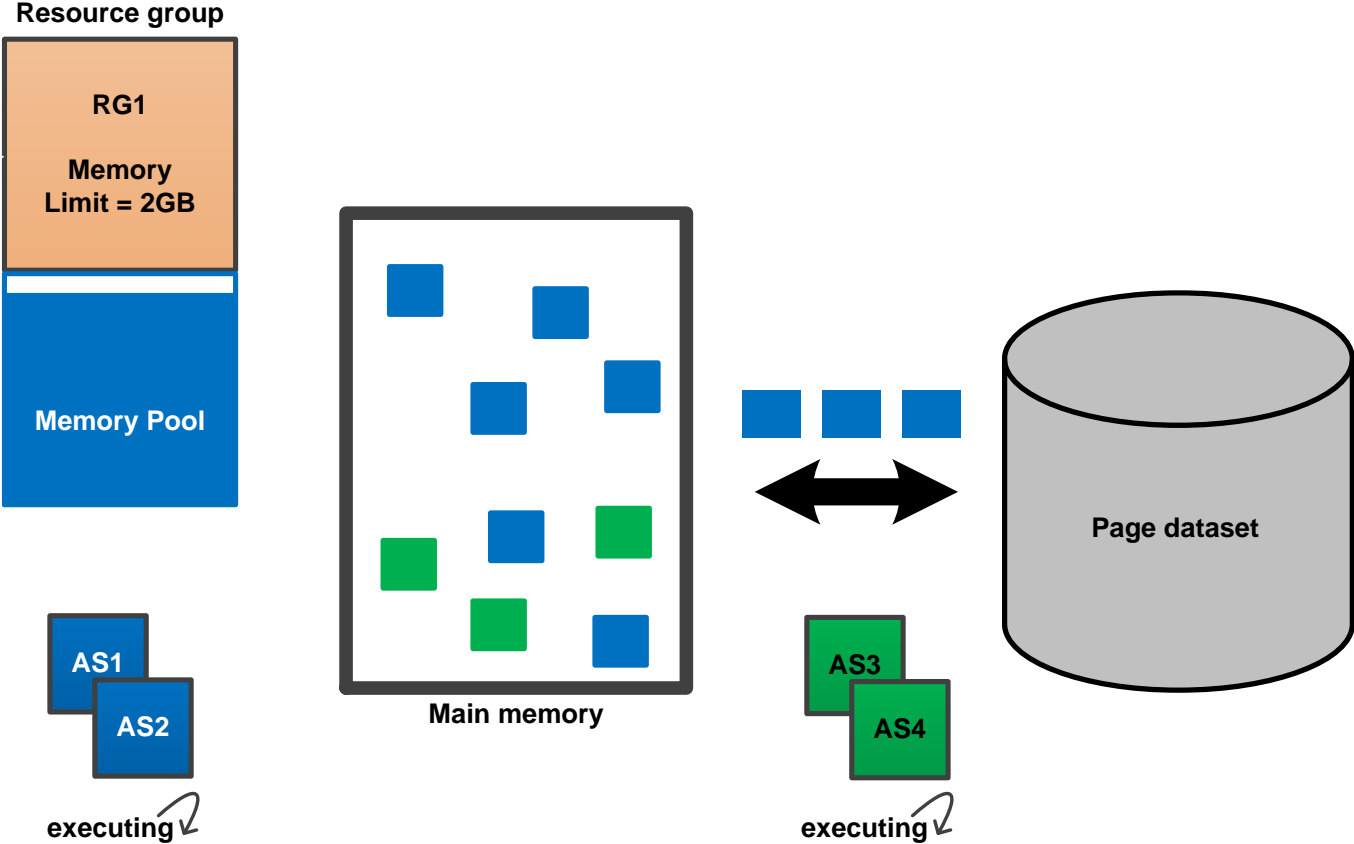
# How do memory pools work?



# How do memory pools work?




# How do memory pools work?



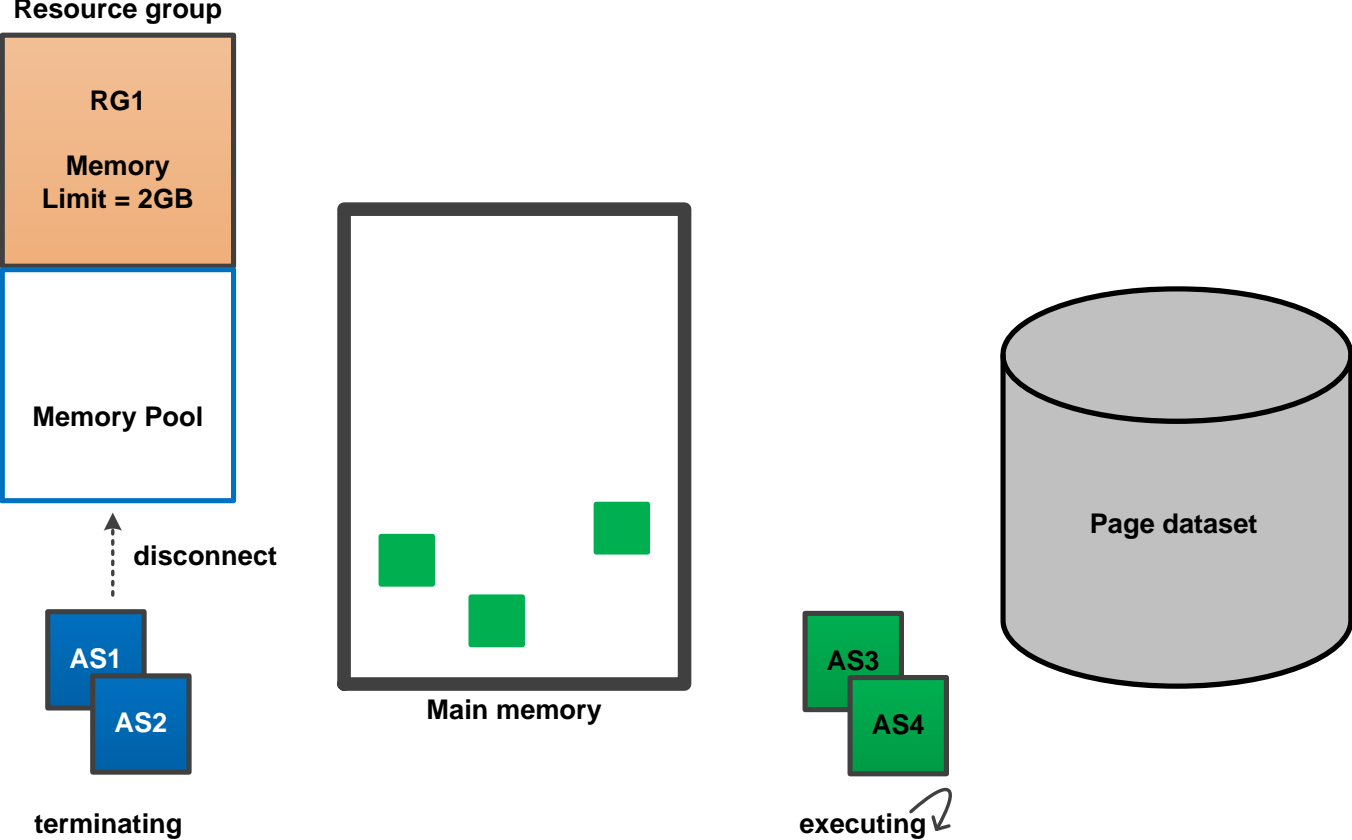
# Is work impacted when the memory pool is full?

- When the memory pool reaches its limit, self-stealing is initiated to keep the number of frames within the limit
  - Address spaces may be suspended when using storage services (getmain, IARST64, etc.)
  - New address spaces connecting to the pool may be temporarily deferred
  - If the shortage can not be resolved, jobs will get canceled with Abend x'E22'
- **Performance of pool members might be impacted severely by paging**

**Note:** Memory pool paging uses system resources that *might* result in a performance impact on work running **outside** the memory pool



# How do memory pools work?





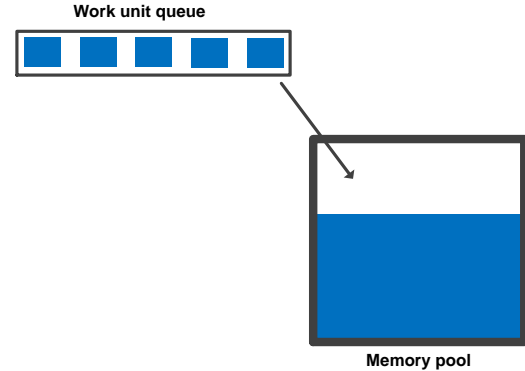
# How to work with memory pools?

```
Modify a Resource Group
Command ==> _____
Enter or change the following information:
Resource Group Name . . . . . : MPLIMIT
Description . . . . . : _____
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 . . . . . : _____
Maximum Capacity . . . . . : _____
Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope) . . . . . 12 GB
```

- Create / delete / size change of a memory pool depends on the memory limit configured for a resource group
- Address spaces can be changed to another memory pool or back to global pool by changing their service class via:
  - RESET command (overwrite service class)
  - SDSF (overwrite service class)
  - WLM Policy activation (WLM classification rules)
- Applications can query the amount of free storage of their connect pool by using the Sysevent STGTEST service

# How WLM can help to prevent paging and job cancellation in the memory pool?

- WLM management available for:
  - Application environments server
  - WLM managed initiators
  - Buffer pools
- WLM anticipates the average storage consumption of your (homogenous) workload running in a service class and starts additional work units based on available storage in the memory pool
- WLM does not start any further work units when the memory pool is in a shortage
  - Available frame shortage
  - Pageable frame shortage



**Note:** A memory limit overrules the storage critical attribute of classification rules and also any protective storage target managed through SRM

# Continuous delivery – Memory pool stages

Function	Stage 1		Stage 2		Stage 3	
	RG	TRG	RG	TRG	RG	TRG
Create/delete a memory pool		✓		✓	✓	✓
Increase the size of a memory pool		✓		✓	✓	✓
Reduce the size of a memory pool		✗		✓	✓	✓
Connect/disconnect a job to/from a memory pool		✓		✓	✓	✓
Switch job between memory pools		✗		✓	✓	✓
WLM managed memory pools for: <ul style="list-style-type: none"> <li>• Application environment server</li> <li>• WLM managed initiators</li> <li>• Buffer pools</li> </ul>				✓	✓	✗
Sysevent STGTEST considers memory pool storage				✓	✓	✓
<b>APARS</b>	WLM OA50845 RSM OA51171 zOSMF PI71084		WLM OA52611 RSM OA53197		WLM OA54452 zOSMF PI96460	

# Summary

- Memory intensive applications like SPARK or sort-jobs can be run in a memory pool **to protect your important business application**
- A memory limit for a group of address spaces can be set in a WLM resource group
- WLM manages the amount of jobs running for a service class to prevent the memory pool from paging
- Running address spaces in memory pools increases the likelihood of paging in the system
- Remember to monitor memory pools for paging activity
  
- **Resources:**
  - [MVS Planning: Workload Management](#)
  - [MVS Initialization and Tuning Guide](#)
  - [Open Data Analytics for z/OS: Installation and Customization Guide](#) (Section Configuring z/OS workload management for Apache Spark)

# Heads up!

```
Coefficients/Options  Notes  Options  Help
-----
                Service Coefficient/Service Definition Options
Command ==> -----
Enter or change the Service Coefficients:

CPU . . . . . 1.0      (0.1-99.9)
IOC . . . . . 0.0      (0.0-99.9)
MSO . . . . . 0.0000   (0.0000-99.9999)
SRB . . . . . 1.0      (0.0-99.9)

Enter or change the service definition options:

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

**It is highly recommended to converge to this set of service definition coefficients!!!**

# z/OS Workload Management

## – More Information –

IBM Industries & solutions Services Products Support & downloads My IBM Search

IBM Systems > Mainframe servers > Operating systems > z/OS >

## IBM Workload Manager for z/OS

### The Power of Orchestration

Overview What's New FAQs Further Info

One of the strengths of the zSeries platform and the z/OS operating system is the ability to run multiple workloads at the same time within one z/OS image or across multiple images. Such workloads have different often competing completion and resource requirements. These

Contact IBM  
Email z/OS WLM

- z/OS WLM Homepage: <http://www.ibm.com/systems/z/os/zos/features/wlm/>
- z/OS MVS documentation
  - z/OS MVS Planning: Workload Management: <http://publibz.boulder.ibm.com/epubs/pdf/iea3w101.pdf>
  - z/OS MVS Programming: Workload Management Services: <http://publibz.boulder.ibm.com/epubs/pdf/iea3w201.pdf>
- IBM Redbooks publications:
  - System Programmer's Guide to: Workload Manager: <http://www.redbooks.ibm.com/abstracts/sg246472.html?Open>
  - ABCs of z/OS System Programming Volume 12 <http://www.redbooks.ibm.com/abstracts/sg247621.html?Open>

Thank you

