

## What's New in z/OS V1R3 WLM

Jürgen Holtz  
IBM Germany  
Schönaicher Str. 220  
D-71032 Böblingen  
✉ [holtz@de.ibm.com](mailto:holtz@de.ibm.com)

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

### Abstract

Find out what's new in z/OS WLM. Come to see which functions WLM offers in addition to its large contribution for Intelligent Resource Director and License Manager. You will hear about WLM reporting enhancements, support for cryptographic devices and other interesting WLM features. This session was presented at SHARE in Minneapolis but has been modified to also reflect new functions available in the next release of z/OS.

## Agenda

---

- ▶ Removal of WLM compatibility mode
- ▶ Enclave Enhancements
- ▶ Sub-Capacity Pricing & WLM Softcapping
- ▶ Performance Improvements

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

---

## Removal of Compatibility Mode

- ▶ How Does it Affect You?
- ▶ Install Definition Utility

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

## How Does it Affect You?

▶ z/OS V1R3 runs only in goal mode

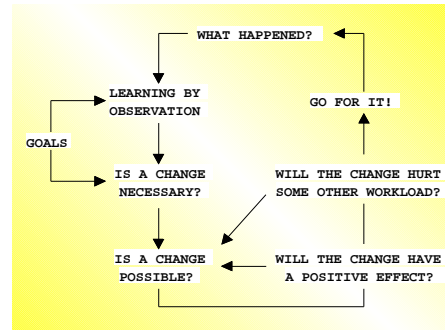
- IPL is done in goal mode no matter what
- IPS=xx and ICS=xx parameters are ignored in IEASYS-member

▶ Any compat-mode function is rejected

- F WLM MODE=COMPAT
- D DMN
- E jobname, PERFORM=xxx

▶ Service definition installation / policy activation is a lot easier

- Provision of Install Definition Utility
- Run it as a batch job or as a started task



IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

Beginning in z/OS V1R3, systems will IPL and run in **WLM goal mode** only. If a non-blank IPS or ICS is specified in the IEASYSxx parmlib member, IRA903I (see below) is shown, and the system IPLs in goal mode.

When the system IPLs successfully, the installation runs with any one of the following service definitions:

The existing service definition created for that installation

A sample service definition provided in SYS1.SAMPLIB(IWMSSDEF)

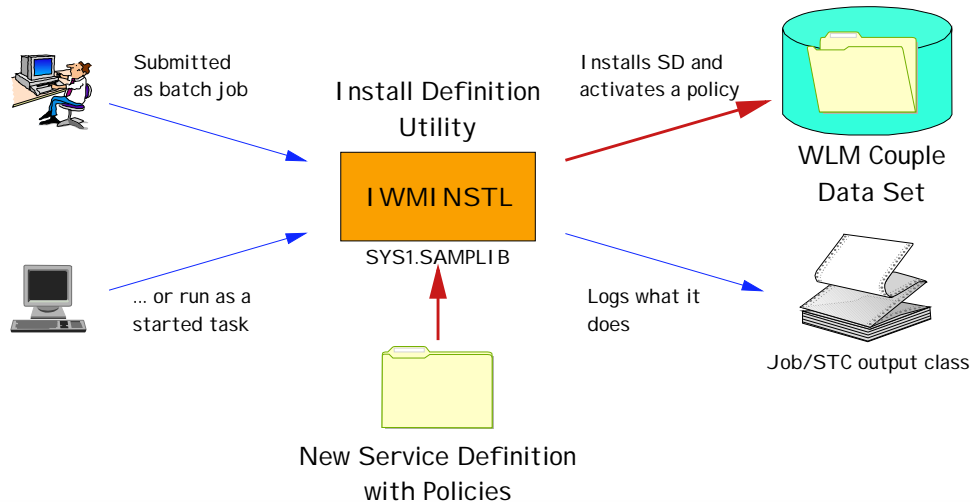
The WLM default "internal" service definition

It is highly recommended to **migrate to goal mode before moving on to z/OS V1R3**. This way you IPL and run the system with a service definition that is already tailored for your specific installation's needs.

In an attempt to switch to compat mode, message *IRA903I WLM COMPATIBILITY MODE IS NOT SUPPORTED* is issued by the system.

The **Install Definition Utility** makes the process of installing a new/changed service definition and activating a service policy a lot easier. Administrators do not have to use the ISPF application for that purpose but can submit a batch job or issue a command to run a started task, instead (see next page).

# Install Definition Utility



IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

The Install Definition Utility allows you to install a service definition and activate a policy without having to use the ISPF WLM Application. You can simply run the utility as a batch job (for example using the sample JCL) or start it from the operator console as a started task.

Any valid service definition can be installed and any policy within that service definition can be activated using the utility.

Usage (refer to the prolog of the sample JCL for details)

Batch: Edit the sample JCL and submit the job

Console: S IWMINSTL,<operator-specified parameters>

Examples:

- (a) Install new service definition without activating a policy  
s iwminstl,force=y,svdefpds=your.definition.dataset
- (b) Activate new policy within installed service definition  
s iwminstl,polname=weekend
- (c) Do steps (a) and (b) at once  
s iwminstl,force=y,svdefpds=your.definition.dataset,polname=weekend

WLM provides a sample service definition in SYS1.SAMPLIB(IWMSSEDEF) for installations that are in the process of migrating their systems to goal mode.

It is a broad definition on functionality level 6 (OS/390 V2R6 and higher)

However, you have to revisit it in order to match your installation's specific needs.

---

## Enclave Enhancements

- ▶ Definition and Examples
- ▶ Reporting
- ▶ Enclave Service Class Reset

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002


# Enclaves: Definition

## ► Characteristics

- Represents business transaction or business unit of work
- Often comprised of multiple dispatchable units (tasks, enclave SRBs)
- Executes in several address spaces even across subsystem bounds
- As a transaction of its own, it can be classified and managed independently from where it is executing
- Single system or sysplex scope

## ► Types

- Independent
- Dependent



IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

## ► What is an Enclave?

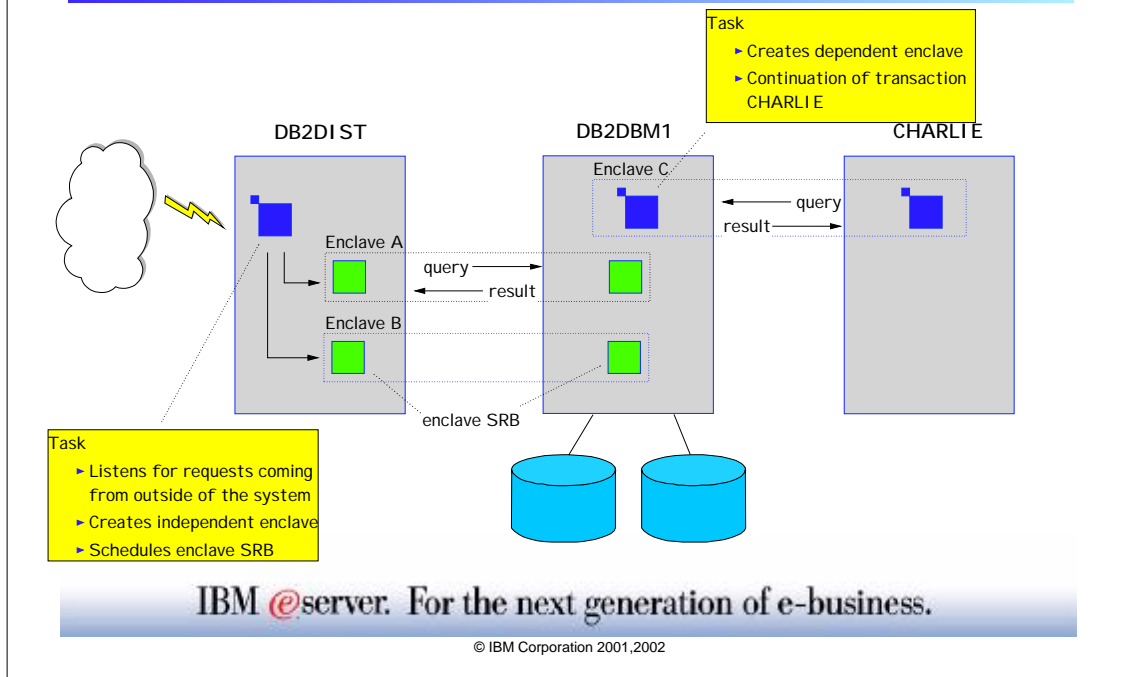
An enclave is a construct introduced in MVS/ESA V5R2 and extended in OS/390 V1R3 that represents a business unit of work. An enclave is an anchor for accumulating the resources consumed by a transaction, regardless of where it is executing. The intent is to have the enclave map closely to the end user's view of a transaction. The transaction represented by an enclave consists of pieces that can span many server address spaces and simultaneously share these address spaces with other transactions.

One of the most novel aspects of enclave transactions is that OS/390 can manage them independently, even though several may be executing concurrently in the same address space. Enclaves are classified by the subsystem creating them and can be associated with distinct service classes and report classes.

Since OS/390 V2.9 an enclave transaction can have a sysplex scope, i.e. all service consumed is attributed to the sysplex transaction rather than to the individual piece of such a transaction executing on a single system.

Basically two types of enclaves are distinguished: (a) independent enclaves and (b) dependent enclaves (since OS/390 V1.3). An independent enclave represents a new transaction, while a dependent enclave is created as a continuation of an existing transaction. In the first case, the enclave is classified using classification attributes associated with a request, and in the second case, the enclave inherits the service class and all service attributes accumulated so far from the existing transaction.

## Usage Example: DB2



- ▶ The following scenario illustrates the use of enclaves in a DB2 environment. DB2-requests from outside the current system are received by the Distributed Data Facility (DDF). A network task takes a request, creates an independent enclave and schedules an enclave SRB into the DB2 master address space to process the request. DDF creates an independent enclave because there is no transaction on the current system related to the query from the network. Such transactions must be classified under subsystem DDF in the WLM classification. Failure to do so causes all enclaves created by DDF to run in service class SYSOTHER.
- ▶ There could be also a TSO address space CHARLIE that issues a DB2-query. When the database call PCs to DB2DBM1, a dependent enclave is created under which the SQL query is finally executed. DB2 creates a dependent enclave because the caller is known and because it already resides on the local system. So the SQL query really represents a continuation of the current transaction processed under address space CHARLIE. A dependent enclave inherits all service attributes of the address space it depends on (service class period, service consumed,...)

# Enclave: Reporting

## ► Online

- RMF Monitor III
  - CPU usage
  - Using / delays
  - Only long running enclaves
- SDSF
  - CPU usage
  - All enclaves

```

RMF 2.7.0 Enclave Report
Command ==>
Samples: 10 System: RMP5 Date: 08/26/98 Time: 17.51.00 Range: 10 Sec
Current options: Subsystem Type: ALL -- CPU Util --
                  Enclave Owner: Appl% EAppl%
                  Class/Group: 16.3 74.1
Enclave Attribute CLS/GRP P Goal % D EAppl% TCPU USG DLY IDL
*SUMMARY
ENC00002 DDF DDFWORK 1 70 12.60 1.260 50 50 0.0
          BHIM
ENC00001 IWEB WEBMISC 1 20.00 9.181 0.918 13 88 0.0
          BHOL
  
```

## ► Postprocessing

- SMF 72: transaction counts and resource usage by service class / report class
- SMF 30: total enclave CPU consumption by owner address space
- SMF 89: total enclave CPU consumption by home address space

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

Enclaves are reported online by RMF Monitor III Enclave report introduced with OS/390 V2.7 RMF. RMF lists all enclaves classified by a certain subsystem, owned by a certain job or running in a specific service class or performance group. The focus of the Monitor III enclave report is on enclaves that run for a longer period in time and that consume a good portion of CPU service during the reported range. Detailed information (using/delay states and classification attributes) is available on a popup panel.

A snapshot of all enclaves within the system/sysplex is also shown on SDSF's ENC report introduced with z/OS V1R2. RMF collects resource consumption data and transaction response time data by service class and report class in the SMF 72 subtype 3 records. The CPU service consumed for enclaves as well as transaction response time data of enclaves is reported for the service class/report class the enclave was running in.

There is no SMF 30 record or equivalent for enclaves themselves. The SMF 30 record of the owning address space (creator of the enclave) reports the total CPU service of all the enclaves it owns.

CPU time consumed by enclaves is recorded in the SMF 89 records of each home address space in which an enclave SRB or enclave task runs (the participating address space(s)) for usage-based pricing but is not included in the SMF 30 or 72 records for those address space(s).

## Enclave Service Class Reset

---

- Allows installation to change performance characteristics of an original independent enclave transaction
  - ▶ Assign enclave to different service class
  - ▶ Quiesce all work in an enclave
  - ▶ Resume, i.e. reclassify, the enclave
  
- Objective
  - ▶ Give an installation more control of independent enclave transactions
  - ▶ Allow sysplex-wide reset for multisystem enclave reset
  - ▶ Similar usage and characteristics as reset function for address spaces
  
- Restrictions
  - ▶ Dependent enclaves are reset through their owning address space
  - ▶ Reset into service classes SYSSTC and SYSTEM is not allowed

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

Enclave transactions can consist of many dispatchable work units (tasks and SRBs) which in turn run in one or more address spaces. Common to all these work units is the fact that they are managed independently from the address space once they are associated with the enclave.

However, once an enclave transaction is created and classified to the system, the enclave remains associated with the service class determined by classification over its whole lifetime. Neither was it possible to accelerate an enclave to push it out of the system, nor could it be tamed or quiesced once it lost control. The only means so far to accomplish these kind of tasks was to change the classification rules, re-install the service definition, and activate a policy. This process, however, is circumstantial and might also conflict with competencies in an IT-shop's organization.

What is really needed is a reset-capability similar to the address space reset function. With z/OS V1R3, operators or system programmers can reset an enclave to another service class or quiesce the enclave.

There is no operator command available to reset enclaves from an operator console, but instead, SDSF for z/OS V1R2 provides an SPE (PQ50025) that includes reset options on its ENC display.

# Enclave Reset: SDSF's Exploitation

- Using SDSF SPE PQ50025 for z/OS V1R2

Overtyping service class here

```

Display Filter View Print Options Help
-----
SDSF ENCLAVE DISPLAY WLMJ ALL LINE 1-3 (3)
COMMAND INPUT ==> SCROLL ==> CSR
PREFIX=* DEST=(ALL) OWNER=* SYSNAME=
ACTION=//-Block,=-Repeat,+-Extend,I-Info,M-Match,R-Reset
NP TOKEN SSType Status Quiesce SrvClass Per PGN OwnerJob CPU-T
-- 2000000002 STC ACTIVE YES RESETSC 1 SPACE0 0.27
-- 2400000003 STC ACTIVE NO STCDEF 1 SPACE0 0.01
-- 2800000004 STC ACTIVE NO TSODEF 1 SPACE0 0.00
  
```

Resume enclave with 'R'

Quiesce enclave with 'RQ'

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

The System Display and Search Facility (SDSF) is the first exploiter of WLM's reset API IWMERES. It enhances the enclave display ENC to allow overtyping of the SrvClass column to reset to another service class. The action commands 'R' and 'RQ' are defined to resume or to quiesce an enclave, respectively.

SDSF will allow the operator to reset a foreign enclave by issuing the reset request on the originating system. SDSF will also allow you to reset a dependent enclave by issuing a reset request for the dependent enclave's owner on the system that the owner is running on.

Access to the reset function is protected through a new RACF-resource. Enclave reset can be granularly allowed by subsystem type and by subsystem name. The resource name is ISFENC.*subsystem-type.subsystem-name*.

---

## Sub-Capacity Pricing and WLM Softcapping

- ▶ Improved Reporting Capabilities

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

# Sub-Capacity Pricing and WLM Softcapping

## ▶ Sub-capacity workload license charge

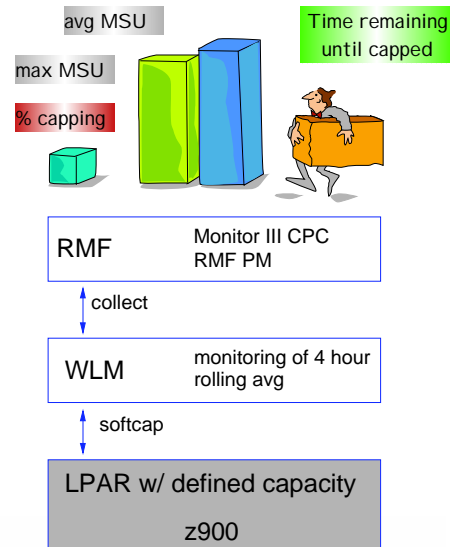
- New pricing model applicable with z900 servers running z/OS in LPAR-mode
- LPAR has defined capacity set in terms of MSUs

## ▶ WLM support

- Monitors the actual capacity in MSU
- Calculates 4h rolling average capacity
- Turns softcapping on when  
4h rolling average > defined capacity

## ▶ RMF support

- Monitor III CPC Capacity report
- Additional metrics in RMF PM



IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

For z900 servers running z/OS V1R1+, LPARs have a new external called "defined capacity" where the customer can (optionally) specify the average LPAR size in MSUs. The rolling 4-hour average utilization of an LPAR is compared, real-time, to the defined capacity of the LPAR by z/OS Workload Manager (WLM). The actual utilization of the LPAR may exceed the defined capacity or dip below defined capacity. However, if rolling 4-hour average exceeds the defined capacity, WLM will constrain the LPAR capacity at the customer-specified defined capacity. This constraint is called softcapping and the softcapping only lasts until the rolling 4-hour average dips back below the defined capacity. The softcap is implemented using existing S/390 technology: when WLM senses that the rolling 4-hour average exceeds the defined capacity it sends a message to PR/SM who, in turn, applies the capping mechanism. The softcap is truly a temporary PR/SM hardcap. For more information, see the PR/SM Planning Guide, Page 3-47. (source: [http://www-1.ibm.com/servers/eserver/zseries/wlc\\_lm/image\\_capacity.html](http://www-1.ibm.com/servers/eserver/zseries/wlc_lm/image_capacity.html))

The Resource Measurement Facility (RMF) provides the new Monitor III CPC Capacity Report that lists the LPARs defined within the CEC, their characteristics, and their current utilization. In addition, the CPC report also shows WLM softcapping information for the partition, the report was requested from. This information includes

The partition's defined capacity, the current average and the maximum. All units are in Million Service Units (MSU).  
Percentage of soft-capping in the reporting interval.

Besides the Monitor III CPC report, RMF PM, the graphical Java client, has been extended by additional resources and metrics. Especially with respect to sub-capacity workload license charge the new metric "remaining time until capping in seconds" is offered to allow analysts to monitor the availability of an LPAR based on software license management. For additional details, refer to the RMF Newsletter 14 "Support for IBM License Manager (ILM)" that is linked off the RMF homepage (refer to appendix for URL).

---

## Performance Improvements

- ▶ Paging I/O Subsystem
- ▶ Application Environments

IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

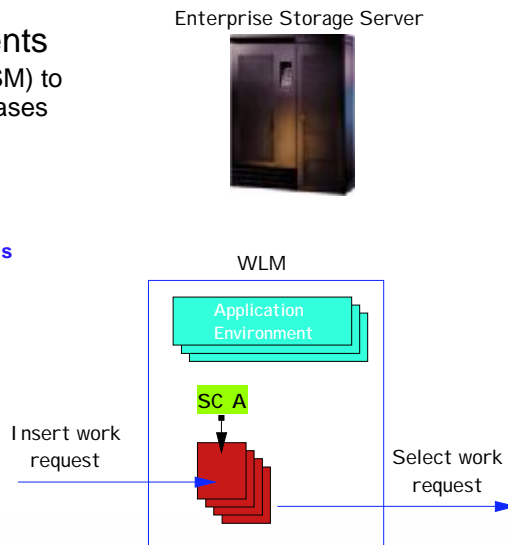
# Performance Improvements

## ▶ Paging availability enhancements

- Allows Auxiliary Storage Manager (ASM) to state the minimum number of PAV aliases needed for a PAV paging device
- Benefits
  - Reduced I/O contention
  - Higher paging throughput
  - Reduced lockout times, e.g. when a dump is written

## ▶ Application environment performance improvements

- Reduced pathlength for work queue requests
- Beneficial to all exploiters



IBM @server. For the next generation of e-business.

© IBM Corporation 2001,2002

On z/OS V1R3, the paging subsystem is enhanced to allow installations the use of Parallel Access Volumes (PAV) for PAV devices containing page data sets. The Auxiliary Storage Manager (ASM) informs SRM about the minimum number of PAV alias devices needed to

- Reduce the overall paging I/O contention
- Provide higher throughput for page I/O
- Reduce the periods of lockouts, for example when a dump is written

The pathlength within the application environment logic has been shortened to improve the overall performance of work managers that insert work requests in a WLM work queue, or that select work requests from such a queue. This improvement is transparent to its exploiters and therefore beneficial to all subsystems using WLM application environments to let the system manage the number of server address spaces depending on the actual needs.

## References


---

### ► Documentation

- z/OS V1.3 MVS Planning: Workload Management (SA22-7602-03)
- z/OS V1.3 MVS Programming: Workload Manager Services (SA22-7619-02)
- Redbook OS/390 Workload Manager Implementation and Exploitation (SG24-5326)

### ► Visit our home page and other interesting sites at

- WLM <http://www.ibm.com/s390/wlm>
- SRM <http://www.ibm.com/s390/srm>
- RMF <http://www.ibm.com/s390/rmf>
- WSC <http://www.ibm.com/support/techdocs>



**IBM @server. For the next generation of e-business.**

© IBM Corporation 2001,2002