

## What's New for WLM in z/OS V1R4

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 2002

# Copyright And Trademarks

---

(c) Copyright IBM Corporation 2002

The following names are trademarks of the IBM Corp. in USA and/or other countries and may be used throughout this presentation:

CICS, DB2, eLiza, IBM, IMS, MVS/ESA, MQSeries, NetView,  
RMF, RACF, S/390, Tivoli, VTAM, VSE/ESA, VM/ESA,  
WebSphere, z/OS, z/VM, zSeries

Other company, product and service names may be trademarks or service marks of others.

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

© IBM Corporation 2002

# Agenda

---

- ▶ WLM Managed Batch Initiator Enhancements
- ▶ Application State Reporting

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

© IBM Corporation 2002

---

## WLM-Managed Initiators

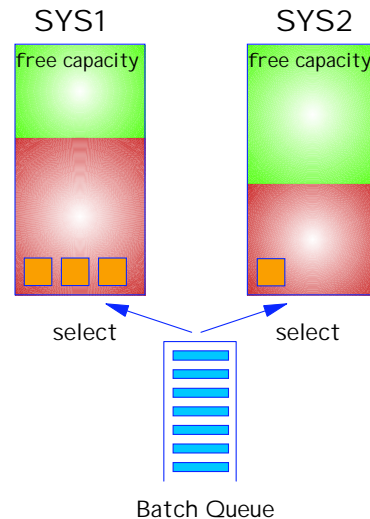
- ▶ Review
- ▶ Limitations
- ▶ Enhanced Batch Initiator Balancing with z/OS V1R4

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

© IBM Corporation 2002

# WLM-Managed Initiators - Review

- Individual batch queues can be WLM managed
- Initiators are dynamically started by WLM to meet service class period goals by reduction of batch queue delays and to optimize resource usage
- WLM selects the system based on
  - ▶ Available system resources
  - ▶ Availability of waiting batch jobs
  - ▶ Importance of the service class
- Initiators are dynamically stopped by WLM
  - ▶ When significantly more initiators exist than needed ( $> 1.5$  times avg. queue length)
  - ▶ In case of CPU or memory shortage
  - ▶ After one hour of inactivity




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

© IBM Corporation 2002

- ▶ WLM batch initiator management was introduced with OS/390 R4 and OS/390 JES2 R4. Later, in OS/390 R8, JES3 started to provide similar functionality for JES3-environments.
- ▶ Beginning in OS/390 R4, WLM has the capability of controlling the rate at which queued jobs are initiated. Moreover, the WLM will dynamically change the number of WLM initiators and/or their work selection criteria in an attempt to meet installation defined goals.
- ▶ The installation can choose between JES managed batch initiators and WLM managed batch initiators by job class. Both types of initiators can coexist. With WLM managed batch initiators, it is WLM who controls the number and placement of the initiator address spaces.
- ▶ New initiators are started when service class goals are missed, when a system is underutilized and there are jobs waiting to be selected, or when jobs have an affinity to a system on which no initiator is available yet. Decision factors are the available CPU and memory resources on a system, the service class' importance, and the projected net value on overall goal achievement.
- ▶ WLM stops initiators, when the number of started initiators is 1.5 times of the long term average queue length, when a system runs short of CPU or memory, or when the last initiator was inactive for 1 hour.

# Limitations Before z/OS V1R4

- Balancing only when new initiators are started
- When initiators are already available, they will be used
- New jobs can be started on any system where idle initiators are, regardless of system load
- Often tends to batch jobs selected on heavily loaded systems, while other systems in the sysplex have relaxed capacity available
- Hurts discretionary workload on loaded systems



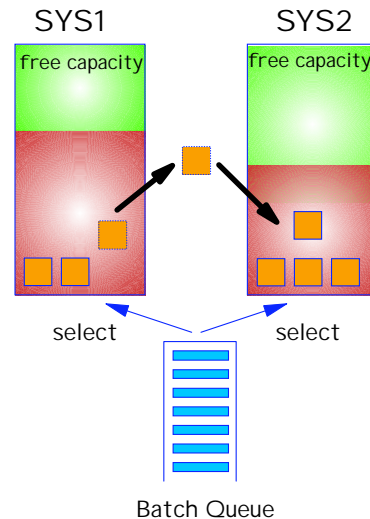
IBM @server. For the next generation of e-business.

© IBM Corporation 2002

- ▶ The existing WLM batch management support has some limitations. One limitation is that often jobs are started on a highly utilized system, while there is enough capacity on other systems in the sysplex. This limitation is addressed by the z/OS V1R4 enhanced batch initiator balancing function explained on the next foil. What are the problems?
- ▶ A balancing decision is only made when new initiators have to be started. In practice, initiators are most likely started on the system where job conversion is done. As long as the system is not short of CPU or memory resources, there is nothing wrong with this decision. Once initiators are available, they can be used to select jobs that wait for execution.
- ▶ Over the time, this leads to configurations where one system provides enough initiators to fulfill the service class goals but these initiators are lumped on one system that is then highly utilized, while other systems have available capacity but lack of initiators to process batch work.
- ▶ This behavior does not only look odd, it also hurts the discretionary work on the highly utilized system in that it prevents it from access to the CPU.

# Enhanced Batch Initiator Balancing

- Balance distribution of initiators over sysplex members
- More aggressive reduction of initiators on heavily loaded systems (< 5% CPU capacity) when other systems have enough idle capacity
- Start of new initiators on systems with relaxed capacity (up to 5 initiators at once)
- Balancing evaluation done every 10 seconds
- WLM keeps enough initiators available for jobs with system affinity to constrained systems



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

© IBM Corporation 2002

- ▶ With z/OS V1R4, WLM improves the balancing of WLM managed batch initiators between the systems of a sysplex. On highly utilized systems the number of initiators can be reduced while more are started on low utilized systems. This enhancement improves the performance of the sysplex with better use of the processing capacity of each system.
- ▶ WLM attempts to distribute the initiators across all members in the sysplex to balance the utilization of the systems while taking care that jobs with affinities to specific systems are not hurt by WLM decisions.
- ▶ When the available CPU capacity of a system decreases to less than 5%, WLM stops an initiator address space when the current system is observed as the system with the highest CPU demand and when there is another remote system, that has enough available resources to start a new initiator. This evaluation is done every 10 seconds. The order of decrease is to stop initiators serving lower importance service classes first.
- ▶ WLM increases the number of initiators on lower utilized systems. To speed up job selection for a high volume of waiting jobs, up to 5 initiators can be started at once compared to the increment of 1 that was used before z/OS V1R4.

---

# Reporting Enhancements

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


© IBM Corporation 2002

# Application State Reporting

---

## ► Problem

- Subsystems like CICS and IMS can present subsystem states for reporting and analysis purposes
- Applications using enclaves for transaction performance management lack of such detailed reporting information



### Before:

- RT, endet trx
- Service
- Using & Delays

### Now:

- + Subsystem work manager delays

## ► New

- Enhanced subsystem work manager delay reporting
- Report-only performance blocks (PBs)
- Supporting all goal types and multi-period service classes
- OW51848 for z/OS V1R2 and higher
- RMF Support: OW52227

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

© IBM Corporation 2002