

WLM Migration Guide and Checklist

**IBM
Advanced Technical Support
Washington Systems Center
Gaithersburg, MD 20879**

May 28, 2003
Version 3.1

About This Document

Disclaimer

The information contained in this document was created by IBM performance specialists. The document is aimed at providing information to help customers successfully migrate their OS/390 systems to WLM goal mode. This information is based on the experiences of IBM performance specialists as they helped numerous customers migrate to WLM goal mode.

The information in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis without any warranty expressed or implied. Use of this information or the implementation of any of these techniques is a user responsibility and depends on the user's ability to evaluate and integrate them into the user's operational environment. While each item may have been reviewed for accuracy in a specific situation there is no guarantee the same or similar results may be achieved elsewhere. Users attempting to adapt these techniques to their environments do so at their own risk.

Permission to Copy

IBM grants permission to the installation to copy this guide, in part or in entirety, but does not waive IBM's statutory right to the copyrighted material contained in this document.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States and other countries and may be used throughout this document:

ACF/VTAM, CICS, CICS/ESA, CICS/MVS, CICSplex, COBOL/2, COBOL/370 DB2, DFSMS, DFSMS/MVS, DFSORT, Distributed Relational Database Architecture, ECKD, Enterprise System/3090, Enterprise System/9000, Enterprise Systems Architecture/370, Enterprise Systems Architecture/390, ES/3090, ES/9000, ESA/370 Enterprise Systems Connection Architecture, ESA/390, ESCON, Hiperbatch, Hipersorting, Hiperspace, IBM, IMS, IMS/ESA, MVS, MVS/XA, Language Environment, MVS/DFP, MVS/ESA, MVS/SP, MQ, NetView, OpenEdition, OS/390, Parallel Sysplex, PR/SM, RACF, Processor Resource/System Manger, RAMAC, RMF, S/370, S/390, System/390, 3090, Large Systems Performance Reference, LSPR, z/OS, z/Architecture, zSeries.

IBM is a registered trademark of the IBM Corporation

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/OPEN Company Limited.

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

OVERVIEW

The IBM Washington Systems Center WLM Migration Guide and Checklist document is intended as a planning aid to assist customer migrations to WLM goal mode. The document is intended to help customers plan and organize the migration effort, listing issues commonly raised by installations during their migration. The document may not answer all of your questions but it will point out topics you should consider.

The document is organized as a checklist with information organized around two major themes:

1. Systems Management Issues or Concerns
2. Setting Goals

The tasks listed may need to be performed by technical staff other than performance or capacity planning specialists. Input from operations, and systems programming staffs will often be required.

The checklist will point the reader to a more detailed section in the document. The detailed sections will give an overview of the item, give pointers to the topic in other published information, and give information on issues and alternatives regarding the topics. Not all actions on the checklist are automatic. Installation policy and procedures may dictate favoring one technique over another. Recommendations regarding the item will be given when appropriate. Not all items will apply to all installations.

Migration Scope

The scope of the WLM Migration Guide and Checklist document is to support an installations successful migration to running their production systems in WLM goal mode. The scope of this effort **does not** include the following:

- Definition of transaction goals for CICS and IMS workloads
- Definition of WLM managed JES initiators
- Definition of Scheduling Environments
- Definition of scalable WEB Servers
- Introduction and Use of Dynamic Parallel Access Volumes (PAVs)
- Explicit R10 enhancement definitions. References are made to R10 enhancements where appropriate with further information available at the WSC and WLM home pages.

The exploitation of the above listed items should be viewed as a follow-on effort once the installation has successfully migrated to running goal mode in production status.

Additional Information:

If you need additional information or want some clarity on some of the checklist items please e-mail the authors at walshk@us.ibm.com. A WSC staff member will respond to your email within 1 business day if possible.

List of Materials

In order to review and complete the checklist in the document additional technical materials will be used. The migration team should have access to these materials. In the checklist the abbreviation will be used when referencing these documents. Different dash levels of publications may make it at times difficult for the migration team to always find the exact reference listed. We understand this issue exists, and though we may be sympathetic, a total cross-referencing of all dash manuals or versions of all documents used in this guide is not practical.

Another fact to keep in mind is the documented reference may often refer to an entire chapter or major section of a publication. It is expected the reader will decide if more information is required, and it will be the reader's decision on how much additional documentation should be reviewed. One of the major goals of this checklist was to organize and direct readers to available material as a means of improving productivity. Any user of this document is invited to bring additional documentation sources to our attention for possible inclusion in the migration guide. The list of materials will include:

| Abbreviation | Document |
|---------------------|--|
| WLMPLAN | MVS Planning: Workload Management |
| WLMSERV | Workload Management Services |
| REDBOOK | OS/390 WLM Implementation and Exploitation |
| RMFREPORT | Resource Measurement Facility Report Analysis |
| RMFGUIDE | Resource Measurement Facility Users Guide |
| MVSETUP | MVS Setting Up a Sysplex |
| MVSCMDS | Operator Commands Reference |
| MVSMMSG | Messages |
| INITREF | Initialization and Tuning Reference |
| INITGUIDE | Initialization and Tuning Guide |
| SMF | Systems Management Facility |
| JCLREF | JCL Reference |
| WSCGOALS | WSC, Presentation, Setting Goals, WSC Home Page |
| WSCMIG | WSC, Presentation, Migration Considerations, WSC Home Page |
| WSCOPS | WSC, Presentation, Operational Considerations, WSC Home Page |
| WSCRMF | WSC, Presentation, Understanding RMF Reports, WSC Home Page |
| WLMR10 | WLM Home Page, see URL below |
| WSCIRD | WSC Presentation, Intelligent Resource Director, WSC Home Page |
| WLMHOMEPAGE | WLM Home Page, URL: http://www-1.ibm.com/servers/eserver/zseries/zos/wlm/ |
| WSCHOT | WSC Presentation, HOT Topics, WSC Home Page |
| WSCISPFEN | WSC Home Page, URL: www.ibm.com/support/techdocs |
| WSCSHORT | WSC Presentation, Short Subjects, WSC Home Page |
| RMFHOMEPAGE | RMF Home Page, URL: http://www-1.ibm.com/servers/eserver/zseries/zos/rmf/rmfhp.html |
| WSCHOMEPAGE | WSC Home Page, URL: www.ibm.com/support/techdocs |
| CHECKLIST | This document |

Estimation of Work Effort

Most migrations to WLM will take an installation about 4-8 weeks, not including the time it takes to enroll and take needed education courses. There are a number of factors which will influence how long it will take. Some of the major factors are:

- Number of systems, or the size of sysplex, which needs to be migrated
- Extent to which the installation has invested in reporting, (Performance Management, Capacity Planning, and Accounting)
- How plentiful or scarce system resources are, especially CPU capacity. Let's face it, a WLM policy doesn't have to be "as perfect" when resources are not in contention, as when there is either CPU or storage constraints.

Generally, multiple skills will be required, not just performance and capacity planning. Training will also be required for operations, production control and perhaps to end users.

Education Plan

This migration effort assumes the technical team lead has some basic knowledge of workload manager, and of basic performance concepts. This basic knowledge can be acquired from many sources including IBM Learning Services courses and from technical publications.

IBM Learning Services

The IBM Learning Services course catalog can be viewed on the Internet at URL <http://www-3.ibm.com/services/learning/>. Several Workload Manager courses are available:

- ES540 Workload-Based Performance Management - 4.5 Days
- H4113 Workload Manager Goal Mode Migration - 2.5 Days

Technical Publications (Reading List)

Another way to get training on WLM is to read the technical documentation which has already been developed. In fact, many of the checklist items used in this document will point you to currently developed documentation.

IBM Services

Other IBM service offerings are available to assist an installation's migration to WLM goal mode. Below is a brief review of available offerings.

1. WSC Pre/Post-Migration Policy Review

This service offering provides an installation a chance to review their WLM project migration plan and their proposed policy before migrating to goal mode or after they have migrated. WSC performance experts will review the customers proposed policy and make suggestions for improvement as well as be a forum for answering questions, and addressing concerns. For more information please send an email to walshk@us.ibm.com.

2. IBM Performance Management Services for S/390 Workload Manager - Goal Mode

IBM Global Services, and the Performance Management Services Offerings, will provide consulting services to migrate a customers IPS/ICS to a WLM policy. More information on this offering can be obtained by calling 1-800-426-4682.

| | |
|--|---------|
| Chapter 1 - Systems Management | Page 9 |
| Evaluating Your System | Page 10 |
| ___ Review IEAIPSxx: | Page 10 |
| ___ Review IEAICSxx: | Page 12 |
| ___ Review IEAOPTxx: | Page 14 |
| ___ CPU Busy > 100% (RCCCPUT): | Page 16 |
| ___ Setting User Coefficients | Page 17 |
| ___ Storage Isolation: | Page 18 |
| ___ Paging DS Configuration: | Page 19 |
| ___ Diverse Systems in a Sysplex: | Page 20 |
| ___ Multiple Jes-plex in Sysplex (JES2): | Page 22 |
| ___ Multiple Jes-plex in Sysplex (JES3): | Page 24 |
| Build Goal Mode Environment | Page 25 |
| ___ Fulfill Sysplex Requirements: | Page 25 |
| ___ Install and Activate WLM Policy: | Page 26 |
| ___ Secure The WLM Policy: | Page 27 |
| Operational Considerations | Page 28 |
| ___ Train Operators on New Commands and Procedures: | Page 28 |
| ___ Check Automation for any Changes: | Page 30 |
| ___ Impact of Reset Command on Policy Activation: | Page 31 |
| ___ Managing Runaway Tasks: | Page 32 |
| ___ Version Management for Service Definitions: | Page 33 |
| ___ Printing the Policy: | Page 34 |
| ___ I/O Priority Management and Tape Work: | Page 36 |
| ___ Location of WLM Application Messages: | Page 37 |
| ___ I Can't Find All of the Subsystem Classification Rules I Need: | Page 38 |
| ___ SMF 99 Management | Page 39 |
| Availability | Page 40 |
| ___ Develop a Fall Back Plan: | Page 40 |
| ___ Develop a Testing Plan: | Page 42 |
| ___ Check on WLM Maintenance: | Page 44 |
| Performance Management and Reporting | Page 45 |
| ___ Develop Naming Conventions | Page 45 |
| ___ PM Tools Update: | Page 47 |
| ___ RMF Workload Activity Goal Mode Reports: | Page 49 |
| ___ Validate Resource Allocation: | Page 51 |
| ___ What is a PI and How do I Use it to Tune: | Page 53 |
| Capacity Planning | Page 56 |
| ___ Migrate Report PGNs to Report Service Classes: | Page 56 |
| ___ Update and Validate Capacity Planning Tools and Reports: | Page 58 |
| Chapter 2 - Setting Goals | Page 60 |
| Building a Policy | Page 61 |

| | | |
|-----|--|---------|
| ___ | Default Service Class For Unknown Work: | Page 61 |
| ___ | Manage the Number of Service Classes: | Page 62 |
| ___ | Requirement For Spare Service Classes: | Page 63 |
| | Workload Considerations | Page 64 |
| ___ | How Do I Manage STC's: | Page 64 |
| ___ | Handling Different Populations of TSO: | Page 66 |
| ___ | TSO Emergency ID: | Page 69 |
| ___ | Discretionary Work and Capping: | Page 70 |
| ___ | The Use of Velocity Goals: | Page 72 |
| ___ | Multiple Period Batch: | Page 74 |
| ___ | PVLDP - There Just Different: | Page 76 |
| ___ | Enclave Information: | Page 78 |
| ___ | When to use I/O Priority Management | Page 79 |
| ___ | Using DB2 Classification Rules: | Page 80 |
| ___ | DDF Service Class: | Page 81 |
| ___ | SPAS Service Class: | Page 83 |
| ___ | UNIX Service Class: | Page 84 |
| ___ | IWEB Subsystem Rules: | Page 86 |
| ___ | CB(WAS) Service Class: | Page 87 |
| ___ | NETVIEW Subsystem Rules: | Page 89 |
| ___ | ORACLE Subsystem Rules: | Page 90 |
| ___ | R10 Enhancements: | Page 91 |

Chapter 1 - Systems Management

The first chapter of the checklist will be used to review different systems management issues which an installation may need to address in the migration to WLM goal mode. The topics are organized in the following manner:

1. Evaluating Your System

This section will discuss issues regarding your current configuration. Some changes can be made to the current system running in compat mode to allow an easier migration.

2. Build Goal Mode Environment

This section will describe the steps needed to create the infrastructure required for a system to run in goal mode.

3. Operation Considerations

Operational considerations are also reviewed. It is important to discuss and educate operations and production control personnel with the changes which will occur during this migration.

4. Availability

This section will discuss both the requirements for testing and for developing a fall-back plan

5. Performance Management and Reporting

The biggest impact of the migration may be on the changes needed for performance management and reporting. WLM offers new capabilities as well as new challenges.

6. Capacity Planning

Capacity Planning systems will also need to be changed to support WLM goal mode, this section reviews the most common items. Don't forget subjects like accounting and charge-backs and how they will be impacted by goal mode.

Evaluating Your System

___ Review IEAIPSxx:

Overview:

The IEAIPSxx member is not used in WLM goal mode. In fact, it is the absence of the IPS= statement in IEASYSxx which causes an OS/390 system to be IPL'ed in goal mode. Certain parameters which were previously defined in the IEAIPSxx such as user coefficient for service units are still used but they are defined in the WLM service definition rather than the parameter library.

Documentation References:

| Document | Section or Page Reference |
|-----------|--------------------------------------|
| CHECKLIST | Storage Isolation |
| WSCMIG | IEAIPSxx Parameters Foil |
| REDBOOK | 2.2.2.1 Dispatching Priority Control |
| INITREF | Chapter 40. IEAIPSxx |

Issues and Differences:

Many of the parameters specified in IEAIPSxx will be dynamically set by WLM in goal mode. Some parameters have been externalized into the service definition, while others are no longer available.

Sysplex members running in compat mode can be managed with different IEAIPSxx members, though many installations share a common IEAIPSxx across all systems in the sysplex and even at times the enterprise. If your systems use different IEAIPSxx members, be aware, once these members migrate to goal mode they will all be managed with a common service policy. This requires the consolidation of each IEAIPSxx member into one service policy.

Note: Review the section on R10 enhancements which allows for classification using SYSPLEX, SYSTEM and JES Subsystem Collection name qualifiers.

| IPS Parameter | Considerations |
|--|---|
| CPGRT CWSS PPGRT PPGRTR PWSS | These parameters are used for storage isolation. Goal mode can set these values dynamically if a service class is missing an objective due to storage contention. |
| TUNIT TSPTRN TSDP TSGRP | These parameters are used for time slicing. Goal mode does not use this facility. Instead, it will dynamically adjust dispatch priority. |
| IOQ | Goal mode will use IOQ=PRTY |
| IOP | Replaced by I/O Priority Management in Service Definition. I/O priority can be dynamically managed. |
| SRVC | Goal mode will use SRVC=COUNT. If a compat mode system uses SRVC=TIME, the I/O service unit calculation will change. This will affect anything which is sensitive to service unit consumption, and may impact billing routines. |
| RTO | The TSO Response Time Option is not available in goal mode. If RTO is specified on the compat mode system, you should consider reducing the response time objective in goal mode by the value specified by RTO. |
| UNT=R | This option is not available in goal mode. Any performance group period using this value should consider using a velocity objective in goal mode. |
| PVLDP | See the IEAOPTxx memberand the section on PVLDP in Chapter 2 |

Recommendations:

Review the WLM Migration Considerations Presentation for a list of IEAIPSxx parameters to review. Determine if your system is using any parameters which have moved to the service definition or are no longer available.

If storage isolation has been specified in the IEAIPSxx, (PWSS,CWSS), determine if it is needed in the current environment or is left over from a previous configuration. If it is no longer needed, consider removing the parameters while still running in compat mode to create a baseline for the Goal Mode migration.

Exceptions: Do not remove the storage isolation parameters if you require this facility.

Release 10 of OS/390 provides for storage critical specification for service classes. See section on R10 enhancements in Chapter 2.

___ Review IEAICSxx:

Overview:

The IEAICSxx member is used to classify work entering a compat mode system to control and reporting performance groups. This function is replaced by the classification rules in the service definition.

Several subsystems, (e.g. DDF), make use of the WLM SRVCLASS parameter to provide performance management reporting and control. This technique allows WLM to associate work entering the system to a service class while the system continues to operate in compat mode. The service class is then mapped to a control performance group. This provides more granular control over the workload, and enhanced reporting capabilities.

Documentation References:

| Document | Section or Page Reference |
|-----------|--|
| INITGUIDE | 3.3.3.3 Assignment of Control Performance Groups |
| | 3.3.3.5 Special Considerations for Subsystems Using Enclaves |
| WLMPLAN | 9.1.1 Defining Work Qualifiers |
| | 9.3.5.1 Using a Sysplex with Systems at OS/390 R4 or Later |

Issues and Differences:

The IEAICSxx member documents a search order hierarchy of potentially present keywords. For example, ACCTINFO is searched first, if no match is found TRXNAME is searched. Additional information on the search order imposed by the IEAICSxx rules is contained in the references above. The classification rules used in the WLM service definition do not impose a hierarchy. Entries are searched in the sequence entered in the definition, with the first match being used. This change should not present a problem in the migration as the goal mode classification rules provide more flexibility than the IEAICSxx hierarchy.

The length of the ACCTINFO parameter in compat mode is 1-8 characters. Because JCL supports 143 characters in accounting information, and the WLM ISPF application allows only eight characters per rule, the application allows "nesting" for accounting information. Again, goal mode provides more flexibility for classifying work by accounting information.

If you use SUBSYS=OAM in the IEAICSxx, no exploitation of workload manager is done by the OAM subsystem. Instead, OAM now provides a SMF record, (SMF record id 85), which can be used to get performance and capacity planning information on the OAM transactions.

Recommendations:

Consider using the SPM qualifier type to classify the system address spaces into SYSTEM and SYSSTC. See the checklist item "What goes into SYSSTC" for more information on classification.

If the SRVCLASS parameter is used in compat mode, consider making the name the same in both goal mode and compat mode. This will allow the same service definition to be used in both modes.

___ Review IEAOPTxx:

Overview:

WLM goal mode eliminates the need to specify many of the IEAOPTxx parameters. The specification of an IEAOPTxx member is still required in goal mode. Parameters specified which are no longer valid in goal mode are ignored when running in goal mode.

Documentation References:

| Document | Section or Page Reference |
|-----------|---|
| INITREF | 42.2 Statements/Parameters for IEAOPTxx |
| REDBOOK | 2.2.1.2 Swap Control |
| | 2.2.1.3 Swap Protect Time |
| | 3.1.1 IEAOPTxx Usage |
| WSCMIG | WLM & DUMPSRV Foil |
| | WLM & IEAOPTxx Parameters Foil |
| CHECKLIST | CPU Busy >100% (RCCCPUT) |

Issues and Differences:

Systems running in compatibility mode with non default parameters may see a difference in performance characteristics once WLM starts using the default values. It is important to review the current IEAOPTxx settings for your current system to see if you are using settings which are different than the WLM goal mode settings.

A common IEAOPTxx parameter which installations often code controls the processor storage management for dumping services. Some installations have coded expanded storage criteria tables, (ESCTxxxx), to force DUMPSRV pages directly to auxiliary storage. Specifications for the criteria tables are no longer possible in goal mode. However, DUMPSRV pages are forced to auxiliary storage automatically in goal mode unless sufficient ESTOR is available.

DUMPSRV address space is given special treatment when it uses ESTOR for paging, WLM determines if there is a need for ESTOR by another workload, the DUMPSRV pages will be the first to be migrated out to AUX even though it is classified to SYSSTC.

Another common parameter often coded in the IEAOPTxx is RCCCPUT, which controls the multiprogramming levels based on CPU utilization. When the RCCCPUT default value is used in compat mode or when the system is running in goal mode, reported CPU utilization will not exceed 100% because the number of ready address spaces not receiving service is not included. Refer to checklist item RCCCPUT for more information.

In z/OS a new keyword is available in IEAOPTxx via APAR Owxxxxxx.

The name of the new keyword member is INITIMP and it is used to set the priority of initiators. If this parameter is NOT defined or if it is defined as INITIMP=0, all JES, APPC and OMVS initiators will run at DP='FE'x as they normally do in goal mode by default. This is done to ensure that the APAR will NOT change the existing system behavior of systems currently in goal mode.

If the parameter is defined as INITIMP=x, with x={1,2,3} the dispatching priority for the initiators will be set below the smallest dispatching priority (at that respective level) for any service classes with the cpu critical attribute specified.

Example: INITIMP=2 will set the DP of initiators below the lowest DP derived from any service class with importance 1 or 2 that has the cpu critical attribute specified. **Note, if any service class with importance 1 or 2 and the cpu critical attribute is defined in the policy, then this keyword applies, even if there are no tasks actually executing in this service class.** If only a service class with importance 1 is defined with the cpu critical attribute, then it is used to derive the DP of the initiators. Finally the customer can define INITIMP=E. In that case the initiators are set to the enqueue promotion dispatching priority which is dynamically calculated.

The use of this keyword provides a somewhat similar way to manage initiators in the same manner it works in compatibility mode where the initiators are set to the value you specify for PVLDP in IEAIPsxx.

Recommendations:

Review the parameters specified in IEAOPTxx. If your system is using a value which will be ignored, consider setting the value to the default while still in WLM compat mode. This will allow a baseline to be established with the new values prior to migrating to goal mode. If you require a parameter be set to a non default value in compat mode, it is likely WLM will manage the requirement in a different fashion after the migration.

Exceptions: Do not remove the criteria table definitions for DUMPSRV if you use this facility

___ CPU Busy > 100% (RCCCPUT):

Overview:

A system running in compat mode can report CPU busy greater than 100% if the RCCCPUT value in IEAOPTxx is coded with values less than 128. The maximum CPU busy percentage which is reported on a goal mode system is 100%.

Documentation References:

| Document | Section or Page Reference |
|----------|---|
| INITREF | 42.2 Statements/Parameters for IEAOPTxx |
| RMFGUIDE | Chapter 2. RMF Sysplex Data Services |
| SMF | 13.76 Record Type 70 (46) -- RMF CPU Activity |

Issues and Differences:

CPU busy greater than 100% can be used by operations as well as automation as an indicator of latent demand on a compat mode system. They can use this value to influence operational changes such as shutting down initiators. Since this capability is eliminated, another metric will be needed, to determine if actions should be taken for an over-committed system when migrating to goal mode.

Recommendations:

The In and Ready value reported by RMF can be used to estimate latent demand on the system. This value is contained in the RMF CPU Post Processor Report and is derived from data in the SMF type 70 record. Automation code could use the RMF sysplex data services API to access the data contained in the type 70 record and determine the In and Ready value for a system real time. If this capability is needed it could be implemented prior to the migration to goal mode to reduce the amount of change during the actual migration.

Setting User Coefficients

Overview:

Goal mode introduces a new construct called a resource group. This construct can be used to provide a minimum or maximum amount of service to a service class. The amount of service provided is defined in un-weighted, (the user specified coefficient is not used), CPU service units.

Documentation References:

| Document | Section or Page Reference |
|----------|---|
| WLMPLAN | 7.0 Chapter 7. Defining Resource Groups |

Issues and Differences:

If resource groups are defined they will specify un-weighted CPU service units, while a service class or performance group period's duration will use weighted CPU service units. RMF workload activity reports will also report service units with the user coefficients applied. The difference in how service units are reported can lead to confusion. By changing the CPU and SRB user coefficients to 1 you can eliminate any confusion in the value of the service units reported. So our recommendations would be : CPU=1, SRB=1, MSO=0.0000 and the IOC=.5. If the coefficients are changed you will need to adjust the duration for all multi-period performance groups and service classes. The change will also affect any other facilities sensitive to weighted service units.

Recommendations:

If resource groups will be used, it may be advantageous to change the CPU and SRB user coefficients to 1. If they will not be used, you should probably continue using the current user coefficients. MSO and IOC should be adjusted as stated above.

Migrate to the final set of coefficients prior to migrating to goal mode. This will provide a baseline in compat mode to allow valid comparisons after the migration to goal mode.

___ Storage Isolation:

Overview:

Storage isolation provides the ability to modify the default algorithms used to manage storage on an OS/390 system. The PPGRT, PPGRTR, and PWSS parameters in the IEAIPSxx are used to activate storage isolation. It is often used to guarantee a minimum number of frames of CSTOR and ESTOR are available or to keep the demand paging rate below a specified value for an address space. Goal mode continues to do storage management to manage working sets or paging, but the parameters are no longer externalized.

Documentation References:

| Document | Section or Page Reference |
|----------|---|
| INITREF | 42.2 Statements/Parameters for IEAOPTxx |
| RMFGUIDE | Chapter 2. RMF Sysplex Data Services |
| SMF | 13.76 Record Type 70 (46) -- RMF CPU Activity |

Issues and Differences:

Goal mode should successfully handle storage isolation requirements in a real time environment. However, some installations need to use storage isolation to protect the pages of their online systems during times of low utilization when other workloads make heavy use of storage. An online system which requires significant amounts of storage during the prime shift may not use many frames during third shift while the batch processing stresses the memory subsystem.

As the work increases in the CICS address space the next morning, the working set of CICS may need to be rebuilt via page fault processing. This can mean unacceptable response times for CICS while the working set is rebuilt. Storage isolation in compat mode addresses this requirement. WLM's historical perspective is not long enough to protect across extended idle periods and so workloads which need this expanded protection need to exercise care when implementing goal mode.

Recommendations:

Review the IEAIPSxx to determine if storage isolation is used. If not, there should be no issues. If it has been defined, an evaluation should be made to determine if it is still needed or is left over from a previous configuration. If it is no longer needed, it should be removed while still running compat mode. If it is needed, evaluate if it is needed for real time memory management or needed to protect frames during low usage periods. If it is needed for real time memory management, the goal mode environment should not be an issue. If it is needed for long term, low usage memory management, the migration to goal mode should be delayed until the installation of OS/390 R10.

WLM provides a new facility called Storage Protection in OS/390 R10. See the section on R10 Enhancements in Chapter 2.

___ Paging DS Configuration:

Overview:

The multiprogramming level can be controlled on a compat mode system with the DMN statements. The multiprogramming level controls are not externalized on a goal mode system. Some installations are reporting a higher multiprogramming level on their systems after moving to goal mode. The paging data set configuration should be reviewed to make sure it can support a higher multiprogramming level.

Documentation References:

| Document | Section or Page Reference |
|-----------|--|
| INITGUIDE | 2.0 Chapter 2. Auxiliary Storage Management Initialization |

Issues and Differences:

The local page data sets normally contain a combination of demand pages and swap pages. The algorithms used by Auxiliary Storage Manager to allocate page slots make use of contiguous slots whenever possible to improve performance. The contiguous slot algorithm is used most effectively when the number of allocated slots is below 25% on the local page data sets. This is one of the reasons installations should keep the number of allocated slots on the locals below 25%. If the multiprogramming level of the system increases after migrating to goal mode the number of slots allocated can also increase.

Recommendations:

Review the percentage of slots allocated for the local page data sets. If the value is approaching 25% consider increasing the size or number of locals to reduce the percent allocated. After the installation migrates to goal mode the percent allocated should again be reviewed. If the percent allocated has approached 25%, consider increasing the number of slots to reduce the allocation percentage to ensure the contiguous slot algorithm continues to perform optimally.

If you are planning to go to 64 bit mode review the WSC white paper on z/OS Performance: Managing Processor Storage in a 64-bit Environment. Document WP100269 in techdocs at URL: <http://w3-1.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100269>

___ Diverse Systems in a Sysplex:

Overview:

Certain environments require different copies of the same started task to be managed to different performance objectives. These include MVS systems with different workloads running in the same sysplex as well as different sysplexes using the same centrally managed and created operating system code loads.

Compatibility mode systems address this requirement with different IPS/ICS members. All members of a sysplex in goal mode will be managed to a common set of goals. Sharing a service definition across multiple sysplexes increases the possibility for different objectives for different copies of the same started task. If the installation has these diverse workload requirements, additional work will need to be done during the migration to goal mode.

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| MVSCMDS | 4.46 START Command |
| JCLREF | 7.0 Chapter 7. Started Tasks |
| WLMPLAN | 9.1 Defining Classification Rules for Each Subsystem |
| WLMR10 | R10 Enhancements presentation on WLM home page |
| REDBOOK | Merging Systems Into A Sysplex Sg24-6818 |

Issues and Differences:

Started tasks are normally classified in the IEAICSxx or the WLM classification rules using the name of the task. Compat mode systems normally use multiple IEAICSxx members to allow the same started task name to have different performance objectives depending on the system it is executing on. A common set of classification rules on a goal mode sysplex invalidates this technique.

The following techniques can be considered for the goal mode migration. The JOBNAME parameter on the start command can be used for STCs using procs or jobcards and will define the name of the task regardless of the PDS member name. The JOBACCT = parameter on the start command can be used for STCs using jobcards and will override the accounting information on the jobcard. Using these parameters provide the ability to assign different service classes to different copies of the same STC.

Recommendations:

If the installation will be running diverse workloads in a sysplex or on multiple sysplexes, an evaluation of these two techniques should be made. If JOBNAME is used, any software sensitive to jobname will need to be reviewed. This may include automation, Automatic Restart Manager, (ARM), etc. If JOBACCT is used, the JCL for the task requires a jobcard as the first statement. Whichever implementation is chosen, it should be implemented while the systems are still running in compat mode.

The examples below use the following classification rules:

| | | | | | | |
|-------|---|----|--------|-------|----------|-------|
| _____ | 1 | TN | GTF | _____ | OPS_HIGH | _____ |
| _____ | 2 | AI | ABCD* | 1 | SYSSTC | _____ |
| _____ | 1 | TN | GTFSYB | _____ | OPS_LO | _____ |

S GTF

The GTF PROC/Started Job will be classified to OPS_HIGH and is named GTF.

S GTF, JOBNAME=GTFSYB

The GTF PROC/Started Job will be classified to OPS_LO and is named GTFSYB.

S GTF, JOBACCT=ABCD

The GTF Started Job will be classified to SYSSTC and is named GTF.

WLM provides additional classification qualifiers in OS/390 V2R10. These include sysplex name, system name, collection name, and scheduling environment name. These qualifiers may provide additional techniques for classifying diverse workloads across the sysplex(es). See the section on R10 Enhancement in Chapter 2.

Multiple Jes-plex in Sysplex (JES2):

Overview:

Customers may run more than one JES2 Multi-Access Spool complex in a single sysplex. The WLM span of control is the entire sysplex not just a given JES2 MAS. When two or more JES2 MAS environments have different performance goals for work in the same jobclass, or with the same jobname then additional classification work will need to be done. JES2 and WLM do not support a method to define JES rules based on JES MAS until JES2 is at a R10 level. As such some other classification attribute must be used, such as accounting codes, to enable classifications of work with the same name or job class to different service classes by MAS.

Documentation References:

| Document | Section or Page Reference |
|----------|-------------------------------|
| WLMR10 | R10 enhancements presentation |

Issues and Differences:

In compat mode the scope of control was the individual system, and each system had an IEAICSxx member to control how work was assigned. Now the scope of control is the sysplex and there is a single set of classification rules for all JES work in the sysplex.

Recommendations:

If accounting codes are not able to be used then some modifications can be made to JES2 to allow the use of the SI classification attribute. The SI attribute refers to the JES subsystem name, typically specified as JES2. It is possible to change the primary subsystem to use a name other than JES2, such as JESA.

For example, if there is an 8-way sysplex, which is composed of 2 4-way MASes, then the 4 JESes in the first MAS could be called JES2, and the other 4 JESes in the second MAS could be called JESA. The subsystem names of the two JES MAS environments would then be different, paving the way for the use of the Subsystem Instance qualifier to be used in the JES definition rules. Then WLM classification rules for jobclass A would be done based on which MAS would run the job. It is recommended the JES service classes be unique across the different MAS configurations. It is possible to have a service class called BATHOT, for instance, in use in both of the MAS complexes, but unexpected results may occur.

R10 enhancements:

Clarification on JES classification rules using Subsystem Collection.

In OS/390 R10 new classification support allows work to be classified by system name. System name, however, cannot be used for JES. This is because the WLM classification call will be made on the system where the job ends conversion, which may or may not be the system where the job ends up being executed. Another classification qualifier, Subsystem Collection Name (SSC), is provided which allows JES work to be managed by MAS. For JES shared spool complexes, the XCF group name, (in WLM it's called Subsystem Collection Name), should be used to provide classification integrity across the JES-PLEX members. The use of the system name qualifier is prohibited for JES rules.

For JES2 you can display the appropriate name to use with the \$DMASDEF command. The "XCFGRPNM" field will be displayed with the name of the qualifier you need to specify with the SSC type parameter under the JES classification rules. Using the SSC technique for JES rules will allow customers with multiple JES MAS's to classify similar work differently among those JES MAS's. For example, a batch job called "compile" can run in a jobclass of C and be assigned service class "BATFAST" in one JES MAS and service class "BATSLOW" in a different JES MAS as shown below. Each JES MAS of course may have multiple members (images). It is also possible to have multiple systems in a sysplex and each JES on each image is its own shared spool environment. In this case each JES system on the image would be considered a MAS with 1 member. Even in these situations, where conversion and execution must take place on the single system, (as it's the only system), the use of system name is not allowed and the installation must use Subsystem collection as the classification attribute.

| | | | | | |
|-----|---|-----|---------|-----|----------------|
| ___ | 1 | SSC | WSCSYS1 | ___ | |
| ___ | 2 | TC | C | ___ | <u>BATFAST</u> |
| ___ | 1 | SSC | WSCSYS2 | ___ | |
| ___ | 2 | TC | C | ___ | <u>BATSLOW</u> |

___ Multiple Jes-plex in Sysplex (JES3):

Overview:

JES3 complexes should match the sysplex. If they do not then classification issues will arise since WLM JES classification rules are sysplex wide. Unless job classes are unique across the different JES3 complexes additional classification work will be required which differentiates the work based on another attribute such as accounting code information.

Documentation References:

| | |
|------|--|
| None | |
|------|--|

Issues and Differences:

None.

In OS/390 R10 new classification support is provided which allows JES work to be segregated by JES3 Complex. See the R10 Enhancement section in Chapter 2.

Build Goal Mode Environment

____ Fulfill Sysplex Requirements:

Overview:

The implementation of WLM requires the system to be IPL'd into a basic sysplex. This can be a multi-system sysplex or a single system sysplex, (monoplex). A basic sysplex requires a sysplex couple data set. WLM also requires a formatted couple data set which is used to store the WLM service definition. The couple data sets need to be accessible from all systems in the sysplex. The use of multi-system enclaves requires a Parallel Sysplex, that is, a coupling facility and structure for WLM's use of this function.

Documentation References:

| Document | Section or Page Reference |
|-------------|-----------------------------------|
| WLMPLAN | Chapter 15 and 16 |
| REDBOOK | Chapter 3.2.3 |
| MVSSETUP | Chapters 3, 15, 16 and appendix A |
| CHECKLIST | Operator training |
| WLMHOMEPAGE | WLM Quickstart Policy Article |

Issues and Differences:

Mixing various levels of the operating systems within a sysplex requires additional awareness with regards to toleration maintenance, functionality levels of the WLM ISPF application, and the format of the couple data set. New couple data set formats are supported by different levels of the operating system. A service definition may have a different functionality level depending on the release of the operating system. Maintenance PTF's may also be required for different levels of WLM to coexist. If you plan on using the new functionality of multi-system enclaves available in OS/390 R9, (retrofitted back to R6 via APARs), then the definition of a structure in a coupling facility is required. The coupling facility must be running CFCC code level 9 or higher.

Recommendations:

Review the above documentation for specifics on the actions you need to take on these issues. See the section on security in this document for authorization required to allocate the couple data sets and create a structure. Allocating an alternate, (secondary), couple data set is highly recommended as well as allocating a spare couple data set. Review with operations the SETXCF command to switch couple data sets.

The WLM couple data set is **only** accessed at image IPL time and whenever a service definition is installed or a policy is activated, therefore the placement of the couple data sets is not as critical as other couple data sets such as the sysplex (XCF) couple data set.

It is also recommended the installation save the JCL used to allocate the couple data set. If it is ever necessary to expand the size of the couple data set, this task will be easier if the original allocation is retained. This step can be avoided if the installation allocates couple data sets from the utilities pull-down in the WLM application. This technique allows the retrieval of the current definitions, which can then be modified to make the data set larger.

___ Install and Activate WLM Policy:

Overview:

Your service policies and classification rules make up your service definition and can be stored in the couple data set and/or a partitioned data set. You can use the WLM ISPF application to install the service definition and activate the policy. You can also use an operator command to activate a policy but not install a service definition.

Documentation References:

| Document | Section or Page Reference |
|-----------|--|
| WLMPLAN | Chapter 17 |
| REDBOOK | Section 3.3.14 |
| CHECKLIST | Version Management for Service Definitions |
| | Secure the WLM Policy |

Issues and Differences:

Installing the service definition overwrites any service definition previously installed on the WLM couple data set. When you extract the service definition, a copy remains on the WLM couple data set until the next install. If you just install the service definition but do not activate the policy then the currently active policy will remain in effect, even after an IPL of the image. This is possible because the active policy and the service definition are stored in different areas of the couple data set. Install and activation of a policy is sysplex wide in scope, there is no need to issue this command on each system running in goal mode.

Recommendations:

Review the checklist items in this document on Version Management for Service Definitions and Securing the WLM Policy.

Before installing and activating a policy, use the VALIDATE option in the WLM application UTILITIES pull down menu to insure you have not committed any coding violations.

___ Secure The WLM Policy:

Overview:

The WLM ISPF Application is protected by a SAF call. Access to the WLM ISPF application is controlled through the definition of a facility class profile in RACF or a security product. Previously the installation would have controlled access to the performance definitions by securing the SYS1.PARMLIB which contained the IEAIPSxx / IEAICSxx.

Documentation References:

| Document | Section or Page Reference |
|-----------|--|
| WLMPLAN | Restricting Access to the Administration App |
| | Authorizing AE Servers, Chapter 12 |
| CHECKLIST | Printing the WLM Policy |

Issues and Differences:

READ or UPDATE access to the **entire** WLM service definition is the only option available via the RACF security rules. READ access allows users access to all functions except installing and activating a service definition/policy. Some installations may want other personnel to have access only to certain functions of the WLM application such as application environments and / or scheduling environments yet not be able to modify any other definitions such as services classes and goals. This is not possible with the WLM ISPF application.

Recommendations:

For general viewing of the WLM policy, make the policy available on the Web (intranet) to the user community. Refer to the section on Printing the WLM Policy.

Specify a UACC of NONE for the facility class resource and limit the number of users authorized to read and update the WLM application to those maintaining the WLM policy and/or performance personnel. Review the requirement for operations to have access to install a service definition versus activating an existing policy. Installing a new service definition can only be done from the WLM application, while activating a new policy can be done by operator command.

Operational Considerations

___ Train Operators on New Commands and Procedures:

Overview:

The migration to WLM introduces a number of new commands to manage the workload environment. Operational procedures need to be updated and the new operator commands need to be discussed with the operations staff so they are aware of when and how to use the commands. If using SDSF service class information is displayed and changeable via the SDSF panels.

Documentation References:

| Document | Section or Page Reference |
|-----------|---------------------------------------|
| REDBOOK | Section 3.2.4 |
| MVSCMDS | Various sections |
| CHECKLIST | Use of RCCCPUT |
| | Requirement for spare service classes |

Issues and Differences:

Some of the new commands are sysplex in scope and others are not. As a *general* rule the VARY commands are sysplex wide and the MODIFY commands are single image in scope.

If current procedures allow the operations staff to alter the priority or performance group of jobs under special circumstances, then new procedures should be documented for how to do this in goal mode.

Activation of a policy, (V WLM,POLICY=policyname), is sysplex wide and therefore applies to all members of the sysplex regardless of whether they are running in goal mode or compatibility mode.

Switching in and out of goal mode is single image in scope. Using the command, F WLM,MODE=COMPAT, the sysplex member switches to compat mode. Special operator training should be done to ensure operations knows the correct procedures for switching modes. And it is always good to remember just because you are not in goal mode doesn't mean the active WLM policy isn't being used for other functions such as scheduling environments, and/or application environments.

If a switch is made to MODE=COMPAT, then if the system image used compatibility mode any time during the current IPL, then the IEAIPSxx member in parmlib used the last time the system ran in compat mode is reused. If the system was never in compat mode during the current IPL, then the system uses a default IEAIPSxx, which places all work into performance groups 0,1, or 2. A WTO, "IRA850E SKELETON IPS IN EFFECT" is issued to identify this situation. You do **NOT** want to run this way. Issue the required SET IPS command or have automation pick up the WTO and reset the IEAIPSxx to the appropriate member.

Reclassification of work into other service classes

Installations today process different types of work with different completion and resource requirements. Every installation wants to make the best use of its resources and maintain the highest possible throughput and achieve the best possible system responsiveness. Workload management (WLM) makes this possible.

With workload management, you define performance goals and assign a business importance to each goal. You define the goals for work in business terms, and the system decides how much resource, such as CPU and storage, should be given to it to meet the goal. Workload Manager will constantly monitor the system and adapt processing based on current system conditions to meet the goals.

Customers spend considerable time in developing their workload manager definitions to match their business goals. When operations staff members or other personnel manually change the service class of jobs which have been assigned by WLM they can affect several areas.

1) WLM's decision making process:

When classifying jobs to specific service classes one objective is to attempt to place similar types of work in the same service class. By doing this WLM's ability to better manage the workload is improved. Through sampling and the use of historical data for the service class WLM can more accurately allocate the resources necessary to achieve the goals of that service class. If jobs are manually reclassified to another service class then you risk changing the service class's workload profile with the potential to skew both the sampling and historical data for both service classes (the old one and the new one). This could result in WLM making resource decisions that may no longer be appropriate or in line with the business objectives initially established.

2) Capacity planning:

As part of the capacity planning process, service class statistics are often measured and plotted over time to establish trends and predict growth patterns. In order for capacity plans to have validity, the stability of the work in the service class needs to be assured. By manually reclassifying work to other service classes, the capacity planning methodology is compromised. Naturally, there are always circumstances when work may need to be moved to a different service class, but these cases should be the exception not the rule. A separate unique service class should be used for these reclassification exceptions so they can be tracked and accounted for accurately. Wholesale reclassification is not an appropriate course of action. If work is not meeting certain objectives as established by the WLM directives, then this needs to be communicated to the performance/capacity planning group for proper resolution through the WLM policy constructs.

Recommendations:

The system programming staff should set up an emergency service class or empty service classes with higher goals specified which operations can use to reset work. These service classes can be used when work needs to be reclassified in an *emergency or in special cases*.

One option open to the installation is to actually remove the DP (dispatching priority) column from the SDSF display. Since the direct control of dispatching priority is removed in goal mode some installations may favor removing it from SDSF-like screens so users do not see and focus on this performance metric.

Review the SETXCF switch couple data set command with operations.

___ Check Automation for any Changes:

Overview:

Some automation packages may have been enhanced by your installation to key on the performance group number (PGN) of jobs, or to initiate a SET IPS/ICS/OPT command based on the time of day or some other event, or you may have set up responses or action items based on specific messages issued.

Documentation References:

| Document | Section or Page Reference |
|-----------|---------------------------|
| MVSPLAN | Chapter 17 |
| CHECKLIST | Training operations |
| MVSMMSG | MVS System Messages Vol 5 |

Issues and Differences:

Performance group numbers are no longer used in WLM goal mode and are replaced with service class definitions. Any automation rules using PGNs will need to be changed. The IPS/ICS no longer can be set in goal mode, but new WLM policies can be activated by command. Any automation rules which switches parameter libraries will need to be changed.

The WLM application issues a set of new messages (IWMAMnn) which are not documented in OS/390 MVS System Messages, Vol 5 (IGD-IZP) but are documented in the Workload Management manual chapter 17. IWM00nn messages are still found in Vol 5 of MVS System Messages.

Recommendations:

Review and update your automation procedures as appropriate. If you are using a WLM structure (used in R9 and up for multi-system enclaves which currently has only one exploiter, Data Miner/POE) set up automation to intercept message IXC585E which indicates a structure full condition. Discuss how to handle the WTO (IRA850E) which indicates you are using the default IPS if this condition arises. Refer to checklist item on training operators on new commands and procedures.

Research the possibility of using automation to prevent operations from manually reclassifying jobs to different service classes at their discretion.

___ Impact of Reset Command on Policy Activation:

Overview:

You may find yourself in the situation where you need to use the reset operator command (E JOBNAME,SRVCLASS=nnnnnn) or use the overtime fields in SDSF to change the service class of a job or an address space. When this occurs be aware a policy activate will not change or reset the service class of a job which has been reset by operator command or through SDSF.

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| MVSPPLAN | System Provided Service Classes, chapter 9 |
| REDBOOK | Section 4.4.3.1 System Service classes |

Issues and Differences:

Prior to OS/390 R4, a job cannot be reset into service classes SYSSTC or SYSTEM. Once it was reset out of these service classes or if it was classified into a different service class, the job would have to be restarted in order to be classified back into SYSSTC, (system address space may have required an IPL). Changes were made to OS/390 R4 to allow address spaces to be reset into SYSSTC and SYSTEM, (only those belonging in SYSTEM).

If you reset a job or an address space to another service class and then wish to set the job or address space back to it's original definition as defined in the active policy, (or to another service class which has been assigned in a new policy), activating a policy will not reset the service class for the job or address space. Since the job or address space was manually reset through operator command or SDSF, you will need to manually reset the job or address space back using another operator command or SDSF.

Recommendations:

Relay the above information to system programmers, operations and others who may have the authority to reset jobs or address spaces.

Implement appropriate security in SDSF or OEM package to allow only authorized staff members to make changes. See reclassification of work into other service classes in previous section on training operators.

___ Managing Runaway Tasks:

Overview:

Most installations have set up their IEAIPSxx to contain a PGN with a domain which specifies a CNSTR value of (0,0). By switching jobs which may be in a loop or are consuming excess resources into this performance group they cause the job to be swapped out until the installation can either cancel the job or contact the owner to determine if and when the job should continue.

Documentation References:

| Document | Section or Page Reference |
|-------------|--|
| MVSPLAN | Chapter 7 |
| REDBOOK | Section 2.2.2.10 and 4.4.2 |
| MVSCMDS | Section 4.28.2 |
| CHECKLIST | Impact of reset command |
| WLMHOMEPAGE | Effective Use of WLM Controls section 6.0.6 and section 14 |

Issues and Differences:

There is no comparable keyword to CNSTR in WLM goal mode. Operations must use one of two methods to effect a similar function in goal mode. They can either use the operator command (RESET jobname,QUIESCE), or change the job to use a service class which has a resource group associated with it using the command “RESET jobname,SRVCLASS=classname”.

“Quiesce” requests the target job or address space be given the lowest possible performance characteristics. Quiesce swaps out swappable work, effectively shutting off work but only lowers the performance of non-swappable work, leaving it swapped in. The non-swappable work will run at the lowest dispatch priority, but it will consume cycles. A quiesce action will persist across any policy activation. A reset command with the resume option is necessary in order to cancel a previous quiesce command.

Information on jobs or address spaces which are quiesced is given in an RMF workload activity report. The SDSF DA panel will show under the Quiesce column any jobs which have been quiesced. You can also view the service class, workload, and the QSC=yes/no fields and other information about a job with the D J,xxxxxxx operator command.

Recommendations:

Create a service class associated with a resource group and specify a maximum service unit capacity of 1 to control non-swappable quiesced tasks. Modify any operator documented procedures with the process to be used for your installation (either a quiesce or service class change) when runaway tasks are suspected. Review checklist item on the impact of the reset command.

___ Version Management for Service Definitions:

Overview:

A management philosophy needs to be established for your installation in order to maintain change control to the WLM service definitions and policies. There is no native version management in the WLM definition, and so installation controls will be needed to ensure you have the ability to keep multiple versions of previously defined and working policies

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| WSCCHECK | Install and activate a policy |
| WLMPLAN | Coordinating Updates to a Service Definition |
| WSCOPS | PDF Presentation on Operational Considerations |

Issues and Differences:

Any updates to a WLM policy will overlay the previously stored members in the PDS if the same data set name is used. If you were to issue a save and then install and activate the new policy, the couple data set information would be overlaid. If you wanted to go back to the previous policy for any reason you would have to extract the current active policy from the couple data set and undo the changes and reinstall and activate the policy.

Recommendations:

Establish a WLM policy management process and create a backup methodology and naming convention for your saved WLM partitioned data sets. An example for the WLM partitioned data sets might be: *HLQ.WLMPOLnn.Vmmddy*.

Have any personnel authorized to use the WLM application use the following sequence when accessing the WLM application:

1. Extract the service definition from the couple data set.
2. Save the service definition in a **new** user defined PDS. This PDS will contain all subsequent changes to the service definition. All saves and exits from the WLM Administration application will store the updated definition in this new PDS. So at this stage there would be two WLM PDS's on the system HLQ.WLMPOL01.V050199, the original definition made in May 1999, and then HLQ.WLMPOL01.V050100 for new changes made in May, 2000.
3. Update the service definition
4. Install the new service definition in the couple data set.
5. Activate the new service policy.
6. Exit the WLM administration application.

The new PDS, HLQ.WLMPOL01.V050100 and production couple data sets are now in synch. The previous PDS, HLQ.WLMPOL01.V050199 would contain the prior version of the service definition which can be used as a fall back if the changes are not correct.

___ Printing the Policy:

Overview:

There is no easy way to print the WLM service policy.

The WLM ISPF application does, however, provide two print-to-file capabilities (which can be found on the *File* pull-down menu):

- The *Print* option produces a flat file with printer control characters
- The *Print as GML* option produces a file which can be used as input to IBM's *SCRIPT/VS* program product

Either of these options may be sufficient if your objective is merely to create a hard copy archive of the contents of your service policy.

However, there is often a need to provide many people (and, in some cases, whole departments) with the ability to view and print the contents of the WLM service policy. At the same time, it is probably not desirable to provide so many people with access to the WLM administrative application. The format of the WLM policy listing is also not attractive, and there have been many customer requests for a better looking document.

Documentation References:

| Document | Section or Page Reference |
|-------------|------------------------------|
| WLMHOMEPAGE | "Here's What's New" hot link |

Issues and Differences:

While, of course, it should be possible to make enough copies of the print files available to anyone who needs to use them, perhaps the following is a better way.

There is a tool called *B2H*, which can be used to convert the *SCRIPT/VS* print file version into an HTML file. The resulting file is nicely cross-referenced and formatted, making it easy to navigate online. Also, the HTML file format is very transportable; HTML files can be easily used in both the host and workstation environments. The HTML file could possibly eliminate nearly all of the need to print out a listing of the service policy, or to give a wide group of people access to the WLM application.

The script input to the *B2H* tool can also be modified before running the tool to add installation-specific information, such as the contact name for questions about the service definition, the PDS which contains the definition, or any other information the installation cares to provide.

The HTML file can then be placed on an intranet site to allow installation access to the WLM information.

Recommendations:

The WLM home page contains a hot link to the B2H ZIP archive file, which can be downloaded from IBM's VM Download Library. The HTML document file (available within the ZIP archive file) contains extensive instructions for installing and using the product.

___ I/O Priority Management and Tape Work:

Overview:

I/O priority management is an optional facility which allows prioritization when I/O requests are queued because a device is busy. When I/O priority management is enabled, I/O scheduling priority is separate from CPU dispatching priority. For each I/O request, WLM assigns a priority to the I/O request based on a value derived by WLM for the requester of the I/O service.

However, **tape I/O is not considered in I/O priority management.**

Documentation References:

| Document | Section or Page Reference |
|----------|---|
| WLMPLAN | Chapter 11.2 Specifying I/O Priority Management |
| | Chapter 16.13 Working with Service Definition Options |
| REDBOOK | Chapter 2.2.1.9 Sysplex I/O Priority Management |

Issues and Differences:

I/O priority management only applies to **non-paging DASD I/O requests**. Tape and paging I/O requests do not participate.

___ Location of WLM Application Messages:

Overview:

The WLM application issues a set of new messages (IWMAMnn) which are not documented in OS/390 MVS System Messages, Vol 5 (IGD-IZP)

Documentation References:

| Document | Section or Page Reference |
|----------|---------------------------|
| MVSMMSG | Vol 5 (IGD-IZP) |
| WLMPLAN | Chapter 17 |

Issues and Differences:

The WLM application messages (IWMAMnn) are documented in the Workload Management manual. IWM00nn messages are still found in the MVS System Messages manual.

Recommendations:

Alert programmers, operations, performance personnel and any other users of the WLM ISPF application about the location of the application messages.

I Can't Find All of the Subsystem Classification Rules I Need:

Overview:

There are at least two cases where you may not be able to find a subsystem you need listed in the available classification rules:

- You upgraded to a new release of OS/390, and IBM added a new subsystem
- You need to define a non-IBM subsystem.

So what does one do when one cannot find the subsystem classification rule one needs?

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| WLMPLAN | Chapter 9: "Defining Classification Rules" |
| REDBOOK1 | Chapter 3.3.9: "Classification Rules" |

Issues and Differences:

If the subsystem you wish to define is not listed in the set of predefined rules supplied by IBM, you may define it yourself or define one of your own. The sample dialogue screen below shows how to create a rule for a subsystem called RYO:

```

Subsystem-Type  Xref  Notes  Options  Help
-----
                Create Rules for the Subsystem Type          Row 1 to 2 of 2
Command ==>> _____ SCROLL ==>> PAGE

Subsystem Type . . . . . RYO (Required)
Description . . . . . Installation-Unique Subsystem
Fold qualifier names? . . . . Y (Y or N)

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

Action   -----Qualifier-----          -----Class-----
Type    Name      Start          Service      Report
_____ 1 TN      RY01          _____  RY01
_____ 1 TN      RY02          _____  RY02
***** BOTTOM OF DATA *****

```

Recommendations:

If you are creating a subsystem definition to support a non-IBM vendor product, be sure to consult its installation documentation for any special instructions concerning creating subsystem classification rules.

___ SMF 99 Management

Overview:

With WLM a new SMF record is produced. There are several subtypes of the SMF 99.

Subtypes

1-5 are intended for detailed audit and debugging purposes. Normally these subtypes should not be recorded. Subtype 6 is a summary record and installations may wish to record this smaller record to have an historical audit trail to help tune their WLM policy.

Installations will have to update their SMFPRMxx member to manage the SMF 99s. Failure to do so may cause an unmanageable load on the SMF data sets.

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| WLMPLAN | Chapter 15, Migration step 10 |
| REDBOOK | Section 3.2, Preparing WLM environment |

Issues and Differences:

SMF 99 contains valuable information for validating and debugging your environment. No post-processor is provided for these records. Performance management systems will need to be modified if reports based on the SMF 99 records are desired. If SMF 99 records are not collected, WLM buffers 15 minutes of records in WLM private storage. If a dump is taken within 15 minutes of a WLM problem, the data from WLM private storage can be used for analysis by IBM.

Recommendations:

The WSC has a post-processor report available for the SMF 99 subtype 6. The generation of this report would require shipment of the SMF 99 records to the WSC. A sample of the report which is created is found below.

```

HOUR MINUTE SECOND SRVCLASS GOALIMP GOALTYPE GOALVAL DP PILOCAL PIPLEX
*****
12 32 2 $SRMDUMP 0 SYS. 0 255 0.00 0
12 32 2 $SRMBEST 0 SYS. 0 255 0.00 0
12 32 2 $SRMGOOD 0 SYS. 0 254 0.00 0
12 32 2 STCHI 1 VEL. 45 253 45.00 45
12 32 2 STCMED 1 VEL. 30 251 15.00 15
12 32 2 STCCIHI 1 VEL. 40 249 1.73 1
12 32 2 STCDBHI 1 VEL. 40 247 40.00 40
12 32 2 STCCIMD 2 VEL. 30 243 30.00 30
12 32 2 STCLOW 2 VEL. 25 243 12.50 12
12 32 2 TSO 1 SHORT 250 241 5.50 5
12 32 2 OSTCDBMD 2 VEL. 30 241 30.00 30
12 32 2 TSO 2 SHORT 2000 239 0.50 0
12 32 2 TSO 2 SHORT 4500 239 5.50 5
12 32 2 STCCILO 3 VEL. 25 237 25.00 25
12 32 2 TSO 3 VEL. 30 235 7.50 7
12 32 2 BATPRDHI 3 VEL. 20 233 1.05 0
12 32 2 BATPRDMD 4 VEL. 20 231 3.33 3
12 32 2 BATTEST 4 VEL. 10 229 10.00 10
12 32 2 BATPRDLO 4 VEL. 15 227 15.00 15
12 32 2 $SRMDISC 6 DISC. 0 192 0.81 0
12 32 12 $SRMBEST 0 SYS. 0 255 0.00 0

```

Availability

___ Develop a Fall Back Plan:

Overview:

A fall back plan should be developed. This plan should include directions to operations on how to move from goal mode to compat mode, as well as what documentation should be taken before the fall back occurs.

Documentation References:

| Document | Section or Page Reference |
|----------|---------------------------|
| None | |

Issues and Differences:

Though the ability to toggle between WLM goal and compat modes can be done with an operator command, some care needs to be exercised about other users of the WLM policy. An active WLM policy is required to use scheduling environments, and application environments in compat mode. Compat mode systems also use an active WLM policy to manage enclaves via the SRVCLASS specifications in the IEAICSxx parmlib member. And other systems in the sysplex may be in goal mode. So complete freedom to change and reinstall new WLM policies may not be possible.

Also it is very important to not be running in either compat mode or goal mode with a default set of performance parameters. So clear instructions to operations on how to effect the change is a good idea.

It will also be necessary to ensure if a fall back occurs the reporting tools are able to handle the different format of performance data created. Some installations may wish to create bi-modal tools and reporting because they will have systems in and out of goal mode frequently. Other installations may wish to only create new service class reporting systems and eliminate from all reporting systems those time frames when the system was toggling between goal mode and compat mode.

Recommendations:

The scope of the fall back planning is an installation decision. Some of the items to include in this plan might be:

- How long to maintain a compatible IEAIPSxx / IEAICSxx with new or changed workloads to ensure a quick return to compat mode.
- When to begin IPLing in goal mode, and eliminate the need for operator intervention to place the system into goal mode after each IPL. This implies operator intervention is required if the system is returned to compat mode for any reason.
- How to fall back out of goal mode but still have a correct and active WLM policy for scheduling environments and application environments.

- How to fall back out of goal mode but still have a valid and active WLM policy for DDF enclave classification using the IEAIPSxx
- How to fall back out of goal mode and install a valid and active IEAIPSxx and IEAICSxx.
- How to manage the different SMF record formats created in the different modes.

___ Develop a Testing Plan:

Overview:

A testing plan needs to be created for the migration to goal mode. The plan should cover the following components:

1. Staging Plan
2. Functional Test Plan
3. Acceptance Test Plan

Obviously the scope and level of detail to which the planning phase is executed is an installation decision. WLM has many capabilities to allow a system to go into and out of goal mode non-disruptively. Installations should review these capabilities and decide which ones should be employed. WLM also provides techniques to allow the installation to gather WLM data while in compatibility mode which can be used to set the WLM policy.

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| WLMPLAN | Chapter 8, Determining Goals in Compatibility Mode |

Issues and Differences:

A migration of this type is often met with resistance and anxiety. Almost every data center problem incurred during the migration will be charged as a “goal mode” problem. The project team has to accept this. It is important to remember most data centers had problems before goal mode and they will have them after goal mode. It’s just the team must be ready to accept and respond to all problems until “goal mode” can be ruled out.

Recommendations:

1. Create a staging plan.

A staging or migration plan should cover the basic time table of the migration. It should cover the schedule of which systems will be moved to goal mode, and what shifts they will be in goal mode. The plan should include information on who to contact with issues about the migration. This plan should be communicated as widely as possible.

WLM offers great flexibility to seamlessly go into and out of goal mode. We recommend the migration to goal mode be positioned as non-disruptive to operations and end users. The installation should stress the system may be in or out of goal mode at any point in time. The WLM migration should not be viewed as a one directional migration with major implications if the system returns to compat mode. The best words of wisdom might be to view the migration as a “cha-cha” and not jumping off a cliff.

2. Functional Test Plan

The functional test plan should be one which exercises all of the systems management tasks which you have done to implement WLM. Below is a list of some sample tasks which would make up a functional test plan.

- Validate all work classifies properly (no SYSOTHER, or undefined service classes used)
- Ensure all monitors can switch between goal mode and compat mode
- Ensure all reporting programs work with the goal mode SMF records
- Ensure all accounting programs work with the goal mode SMF records
- Test all automation changes by causing the messages / actions to be triggered
- Change the WLM CDS dynamically
- Ensure security profiles are working correctly
- Test the fall back plan
- Ensure the correct subtypes of the SMF 99 records are being created

3. Acceptance Test Plan

The acceptance test plan should document the steps the installation will take to ensure the WLM goal mode policy is working as planned. This test plan should document the different time periods which will need to be exercised before the policy can be judged acceptable. For instance it may be deemed necessary the WLM definitions will have to be checked for the batch window, the weekend shifts, a month closing, a quarter closing, etc.

The acceptance test should also highlight what data will be reviewed to ensure the policy and the goal definitions are correct. For instance, for each workload the goal versus the actual goal will be tracked for each service class period. Another test might be to ensure there is a high enough completion rate in each of the service class periods. Some installations may wish to review the software impact of the migration by measuring the WLM CPU utilization before and after the migration, the activity to the WLM CDS, and other throughput metrics.

Check on WLM Maintenance:

Overview:

Prior to cutting over to goal mode a check should be made for all known WLM maintenance.

Documentation References:

| Document | Section or Page Reference |
|-----------------|---|
| WSCHOT | Hot Topics Presentation WSC Home Page www.ibm.com/support/techdocs/presentations |

Issues and Differences:

None

Recommendations:

Apply all appropriate WLM maintenance.

Performance Management and Reporting

___ Develop Naming Conventions

Overview:

Naming conventions for the different constructs used in the WLM policy are important to develop **before** actually creating the service definition. The goal for naming conventions would be to allow the name to help identify the element, such as a service class, a report class, a classification group, or a workload. Other elements such as application environments and scheduling environments should also be developed with a naming convention in mind.

Documentation References:

| Document | Section or Page Reference |
|----------|---------------------------|
| WLMMIG | Page 2 |

Issues and Differences:

In a goal mode system names given to collections of work will be visible to the end users. Previously comments describing the importance, (or lack of importance), of the work were 'hidden' in the IPS/ICS and the only external to those outside the data center was a meaningless PGN number. Now the names chosen will be read by a much wider audience. Let's give an example. Batch low work might be a candidate for a name such as BLOW for BatchLOW. Maybe not what you would really want.

Also some concerns have been raised about service class names such as BATSLOW, as the end user department assigned to this service class may not be as thrilled as you think they should be. As one manager of a large data servicer mentioned "I can't have tours of my data center and point out to the executive - there's your data running in BATSLOW. It's not the smartest marketing effort."

Actually building a naming convention is beyond the scope of this document, but it is something you should spend some time working on. Get several people, including operation to look at your scheme and see if it makes sense.

When building names remember WLM will sort the names and they will appear on the monitors in this order. So if you choose a convention such as BATHIGH, BATMED, BATLOW, on a monitor the workloads will appear as BATHIGH, BATLOW, BATMED. How much this bothers anyone is an installation decision.

Another concept to avoid in building names for service classes would be to use the goal information in the service class name. For instance, BATVEL30 for a batch workload with a velocity of 30, or TSOR8001 for a TSO workload with a response time goal of 80% in 1 second. As system configurations change, or as the WLM migration process unfolds, it may be necessary to change the service class goal. After the changes the existing service class name may be totally out of date. Changing service class names each time you needed to change a goal would impact

long term performance and capacity reporting as some translation mechanism would need to be done each time a service class needs to change.

Recommendations:

On the policies the WSC builds we try and follow a naming convention along these lines.

| TYPE | EXAMPLE | DESCRIPTION |
|-------------------------|------------------------------------|--|
| Workload | BAT_WKL | Workloads have a suffix of ‘_WKL’. Workload names are not addresses by operator commands and so the use of special characters like the ‘_’ is not disruptive. By using a suffix of ‘_WKL’ we can quickly differentiate workloads from service classes. |
| Service Class | BATREG | Try and use as descriptive a name as possible. Watch sort orders. |
| Resource Groups | LIMIT_RG | We use a suffix of ‘_RG’ for resource groups. |
| Report Classes | RJES2 RVTAM RTCPIP RCICSP | We prefix all report classes with an ‘R’ to indicate report class. This would require no service class to have the need to start with an R. A suffix of ‘_R’ would also work, but report classes may need to be specified in RMF control cards. |
| Transaction Name Group | STC_TNG | Again we use a suffix to indicate the construct is a named group. |
| Transaction Class Group | JES_TCG | Another suffix and we try and make the prefix meaningful, JES or STC, or IWEB groups. |

The characters SYS is reserved for IBM use and cannot be used as the first three characters for your own workload or service class names.

___ PM Tools Update:

Overview:

An inventory of all Performance Management tools, monitors and reports which use performance groups, (PGNs), will need to be made. With the migration to goal mode, PGNs are replaced by workloads and service classes. This will not necessarily be a 1:1 mapping. A plan will be needed to ensure all of these tools, monitors, OEM software and reporting systems are updated to handle service classes.

Performance reporting programs which use the RMF 72 records will have to be changed once the system is running in goal mode. Once the switch is made to goal mode, all of the currently active transactions are ended from an SMF point of view, and new transactions are started. At the time of the switch RMF will begin to write new subtypes of the RMF 72 records in which to record the goal mode information. Reporting programs which use the RMF 72 records will need to be updated.

Documentation References:

| Document | Section or Page Reference |
|----------|---------------------------|
| None | |

Issues and Differences:

Tools should be reviewed to determine if they provide enough granularity at the service class level. In order to support the sampling requirements for workloads defined with a velocity goal more address spaces may be combined into a single service class. Previously a monitor may have reported at the PGN level, and each PGN contained one major system component such as TCP/IP, JES, NetView, etc. Now in goal mode, many of these address spaces will be combined into a single started task service class. So it may be necessary for tools or reports to be able to display both service classes as well as report classes, or other tools may be needed to display the collection of address spaces in a service class with more granularity.

Another change which will need to be made to the performance management reporting is the scope of control of the reporting system. WLM is sysplex aware and hence the performance data generated, and the system actions, are made on a sysplex not a system basis. It may be necessary to update or create new performance reporting methods which review the performance data on a sysplex level as well as a system level.

Because of the sysplex reporting care needs to be exercised when using the RMF workload activity reports, and the amount of processor capacity used by the workload as indicated by APPL %. APPL% is the total of CPU seconds used by the service class across all systems in the sysplex without regard to processor speed. Manipulation of the RMF input can be done to get a single system view of the work, or the SYSID keyword can be used to limit reporting to just one system.

A decision will have to be made about how to migrate the reporting programs. Since there may be a period where the systems may be toggling into and out of goal mode reporting programs

may need to be written to handle the different subtypes of the RMF 72 records. Some installations may decide the migration period is a small enough of a window where the tools will only support the new subtypes of the RMF 72 records.

Recommendations:

Once you make this inventory it is probably a great to time to review if anyone is using any of the information in any of the reports or if they are being generated just to decorate a file drawer, or trash bin.

For all monitors ensure the support is provided to support service classes. Hopefully the monitors should be able to seamlessly toggle between goal mode and compat mode.

Incorporation of new metrics such as the performance index should be undertaken, and some metrics, such as those based on dispatch priority, may need to be removed. It would also be important to give some thought about how to incorporate sysplex level reporting into the performance management process, and not just report at the individual system level

___ RMF Workload Activity Goal Mode Reports:

Overview:

New RMF workload activity reports are created when running in goal mode. These new RMF reports give new and improved information to better manage and tune the systems performance. It is important to review and become familiar with generating these new reports and learning how to use some of the new fields in the report.

Installations using products such as CMF will also have new reports for goal mode systems. Knowledge of the use of these reports is outside the scope of this guide. However a presentation given at user group conferences will be available on the IBM website. See below.

Documentation References:

| Document | Section or Page Reference |
|----------|---------------------------|
| WSCRMF | Entire Presentation |
| WSCRMF | Techdocs item PRS618 |
| WLMPLAN | Defining Report Classes |

Issues and Differences:

The major difference with the RMF reports in goal mode versus compatibility mode is the goal mode activity report will be a sysplex wide report unless the SYSID parameter is used or the RMF input data is limited to a single system. The information in the report when reporting on a sysplex will be consolidated into a sysplex view, with each system reporting data having a detailed line. The major concern will be the reporting of APPL% for each service class. What is being reported are the CPU seconds from all processors. There is no CPU normalization for processors of different speed, CPU seconds used are simply added together.

Another difference with the goal mode workload activity reports is there is no summary section which sums the information for all of the service classes in the interval. This summary section would have been produced in compatibility mode, (the "ALL ALL ALL" section). The summary section in compatibility mode was often used to calculate the capture ratio for the interval. Since this summary section is not produced the user of the RMF report will have to manually sum the service class sections. Another approach which might be quicker would be to run a SYSRPTS(WLMGL(WKLD)) report which gives a summary section for each workload. This generally will require fewer sections to be summed. Again caution should be exercised when calculating or using the data from a goal mode workload activity report to ensure data from only one system image is being produced.

Report classes in goal mode are treated differently than report performance groups. Report classes are not included with the service class information in the RMF workload activity report. Instead a separate RMF control card needs to be used to generate the report class information SYSRPTS(WLMGL(RCLASS)). There is another factor to keep in mind with report classes. Since a report class can contain service classes with different goal types, the performance index is not included in the report. Enhancement have been made in z/OS R1.2 which provide report

class period reporting and more information at the report class level. See the WLM manual for more information about this.

Recommendations:

As a general rule, the WSC recommends generating the goal mode reports using the SCPER control statement `SYSRPTS(WLMGL(SCPER))`. This gives the best level of detail, and includes the response time distribution information for tuning response time goals.

If the requirement is to calculate the capture ratio, first ensure the RMF workload activity report does NOT include multiple systems and use the `SYSRPTS(WLMGL(WORKLOAD))` level control which sums the service classes within the workload. This will mean less entries to sum.

___ **Validate Resource Allocation:**

Overview:

Some method should be developed to validate the new goal mode policy. The most straight forward way to do this is to create a mapping of PGNs to service classes and then create a spreadsheet which shows the resource allocations under goal mode and under compat mode. This is done to ensure the new policy is allocating resources in the same manner as your old compat mode system. If there are differences then the analyst should review what they are and perhaps make policy adjustments.

Documentation References:

| Document | Section or Page Reference |
|-----------------|----------------------------------|
| None | |

Issues and Differences:

If your system is running well in compat mode, and the resource allocations are generally viewed to be acceptable then the migration to goal mode should ensure this allocation does not change. This validation step is probably most important when there is contention for resources. What this step will do is to ensure the business objectives set for the workloads don't cause one workload to see throughput enhancements at the expense of a workload, in a manner not acceptable to the installation.

First, the migration team will need to keep track of the mapping of performance groups to service classes. This is needed to be able to compare the migration period with some confidence the reporting buckets represent generally the same workloads. An inhibitor to doing this work will be if there isn't a direct mapping of performance group to service class. If work which was previously running in 1 performance group is now switched to multiple service classes then additional work may be needed to build the mapping after the migration. One solution might be to introduce some additional report performance groups to handle the control performance group split, split the workloads to multiple control performance groups while in compat mode, or to switch and build the validation spreadsheet using the SMF type 30 records.

Recommendations:

Review the resource allocations before and after the migration. It would be important to look at the CPU utilization, working set storage used, and throughput metrics such as MPL, response times, transactions per minute, and the relative number of transaction ending in the different periods. Reviewing the relative dispatching priority is not a bad thing to do in the early days to ensure the goal setting was done appropriately. After the migration, avoid looking at the dispatch priority as the intent of WLM to is to manage the system by dynamically changing the dispatch priority.

Below is an example of how a validation spread sheet might appear. This is only provided as a sample, and the information listed is available from RMF reports.

| Service Class | PGN | CPU C / G | WSET C / G | Ended Trx C / G | Resp Time C / G |
|----------------------|---------------|------------------|-------------------|------------------------|------------------------|
| BATHOT | 1, 100 | CPU / CPU | KBs / KBs | count / count | time / time |
| BATREG | 3, 101,102 | CPU / CPU | KBs / KBs | count / count | time / time |
| CICSPROD | 200, 201, 202 | CPU / CPU | KBs / KBs | count / count | time / time |
| CICSTEST | 210, 211, 212 | CPU / CPU | KBs / KBs | count / count | time / time |
| DB2 | 300, 301, 302 | CPU / CPU | KBs / KBs | count / count | time / time |
| STCHIGH | 5, 400,401 | CPU / CPU | KBs / KBs | count / count | time / time |
| TSO | 2, 222, 322 | CPU / CPU | KBs / KBs | count / count | time / time |
| SYSOTHER | 210, 211, 212 | CPU / CPU | KBs / KBs | count / count | time / time |
| SYSSTC | 402 403, 404 | CPU / CPU | KBs / KBs | count / count | time / time |
| SYSSTEM | 0, 500,550 | CPU / CPU | KBs / KBs | count / count | time / time |

Another type of resource which may be reviewed here is the cost of the WLM migration. Most of the cycles to provide the new goal mode support will be recorded in the CPU charged to the WLM address space. Review this CPU to get a feel for the software ITRR cost of the migration. One of the biggest factors which will influence the costs of the goal mode migration will be the setting of the CICS MAXTASK parameter. This number should not be overstated.

___ What is a PI and How do I Use it to Tune:

Overview:

A performance index, (PI), is a calculation of how well work is meeting the defined goal, and is used to compare goals of different types. A performance index of 1.0 indicates the service class period is exactly meeting its goal. A performance index greater than 1 indicates the service class period is missing its goal. A performance index less than 1.0 indicates the service class period is beating its goal.

During the migration to goal mode each service class period must be assigned a performance objective and an importance. The importance, (value 1 to 5), specified for each service class period is the relative importance of meeting the defined performance objective. With sufficient resources, all work might meet its performance objectives. When resources are constrained, WLM will donate resources from service class periods with a lower importance, (higher number), to a service class period with a higher importance, (lower number). Because the performance objective and importance values are used by WLM to direct decision making, it is critical to set these correctly if WLM is to manage the complex to the installation's satisfaction.

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| WLMSEV | Performance Index |
| WLMPLAN | Defining Service Classes and Performance Goals |

Issues and Differences:

Performance management methodologies need to be updated to include the PI metric as an installation moves to goal mode. The PI is used by WLM to direct donor/receiver resource tuning decisions.

Recommendations:

Developing a method of how to use and react to the PI depends upon the workload, and it's importance to the installation. From a WSC straw-man perspective we would suggest all work be set into two broad performance objectives:

- **Loved Work**
Work with this objective must be met for the organization to meet it's business objectives. If it is missed, a tuning activity is required to meet the objective. If tuning can not address the requirement, more resources will need to be acquired. This performance objectives will often be associated with a high importance.
- **Non Loved**
Work with this objective can include a certain amount of delay. Even though the performance is not optimized, the organization can still be successful with a certain degree of delay. Often, if this workload misses it's performance objective, the objective will be lessened versus tuning the system or acquiring more resource. These performance objectives will often be associated with a low importance.

Performance Indexes do not stand alone. They must be reviewed in terms of the resource contention on the system, as well as the amount of more important work on the processor. If

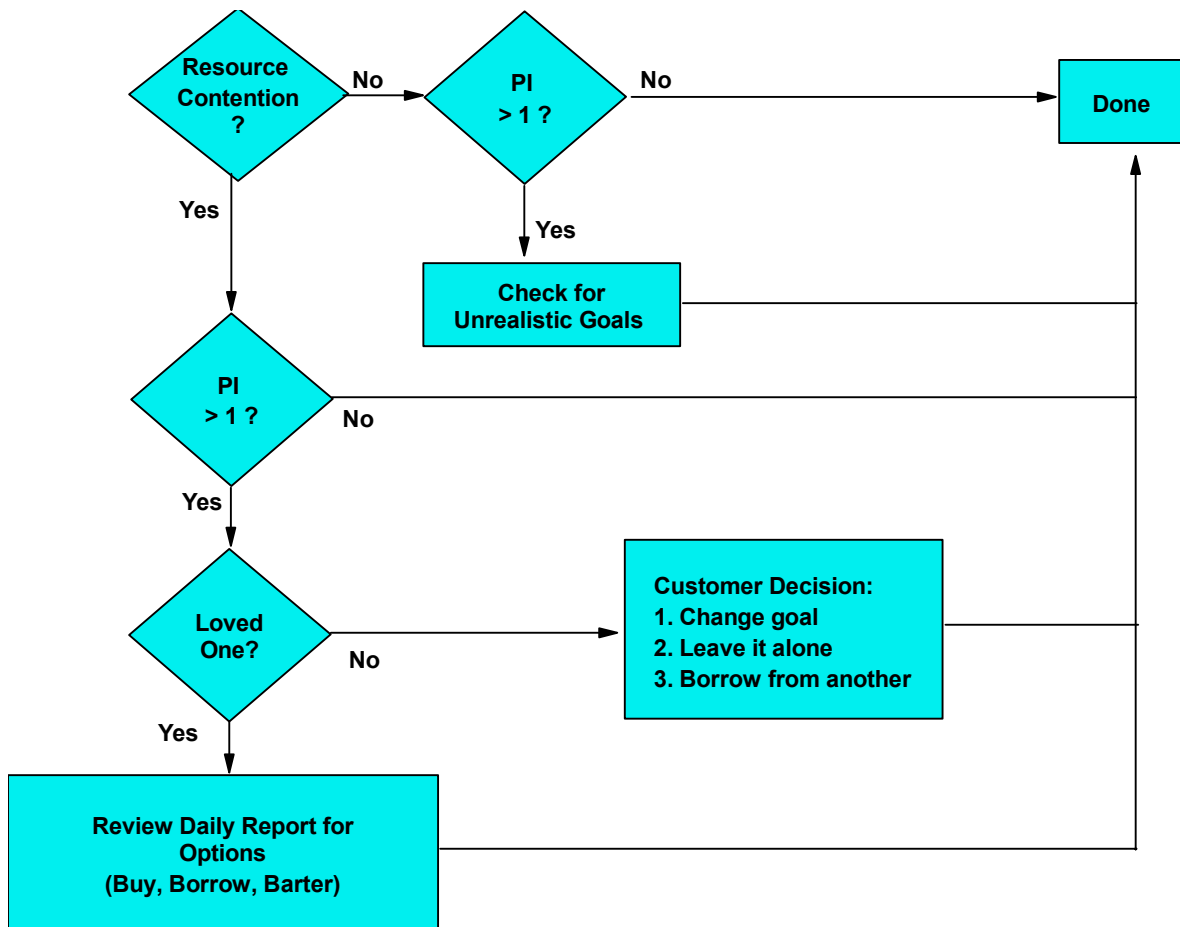
resources are not in contention, (CPU utilization is 90% or less and there is no storage, DASD, or multiprogramming delays), then the PI's should be low, or very close to one, as it should be easy to meet the objective. If the PI is not low then it may be the performance definition for the goal is unrealistic. Also keep in mind that idle time does not figure into the velocity calculations, so address spaces with lots of idle time may not meet their goals.

An unrealistic goal is one where the performance objective may not be met regardless of the amount of resources made available. The WLM policy should never include such an objective. By definition it can not be met, but WLM may initially spend resources trying to achieve the goal until it reaches a point where it no longer can provide any more resources to help.

Once the appropriate performance objectives and importance levels have been associated with all service class periods, the PI is a valuable new metric which can be used to help tune the sysplex.

The service policy should be built with the above mentioned items in mind. Once it is built to these specifications, the following flowchart can be used to refine the policy and start any required tuning actions.

PI Tuning Flowchart



Capacity Planning

___ Migrate Report PGNs to Report Service Classes:

Overview:

Report performance groups provide additional granularity for reporting performance information using the RMF type 72 data on compat mode systems. The report service class is provided on goal mode systems for the same purpose. If the installation makes use of report performance groups, it needs to be aware of the difference in the implementation of report service classes.

Documentation References:

| Document | Section or Page Reference |
|-------------|---|
| INITGUIDE | 3.3.3.7 Report Performance Groups |
| SMF | 13.34 Record Type 30 - Common Address Space Work |
| | 13.78 Record Type 72 - RMF Workload Activity and Storage Data |
| WLMPLAN | Defining Report Classes |
| WLMHOMEPAGE | WLM Latest and Greatest presentation |

Issues and Differences:

The report performance group is associated with a unit of work in the IEAICSxx parmlib member. It is possible to assign a single unit of work to multiple report performance groups. In WLM a report service class is associated with a unit of work in the classification rules of the WLM service definition. A unit of work can only be assigned to a single report service class.

The installation may decide to use classification groups to create a single group name to identify multiple work units in the classification rules. If this technique is used, be aware a report service class assigned to the group will be associated with all units of work in the group. It is not possible to assign different members of the group to different report service classes. Heavy use of report service classes may preclude the use of the group concept in WLM classification.

Report performance groups can be identified on an RMF workload activity report by asterisks in the domain and time slice group columns of the report. Report service classes are not reported on the RMF WLM workload activity report. If a report service class report is required, the WLMGL(RCLASS(option)) control statement must be added to the RMF post processor control statements. z/OS R1.2 has provided enhancement to report class reporting. See the WLM homepage for a presentation on “WLM The Latest and Greatest” or the WLM Planning manual.

Installations which did not extensively use report performance groups may need to reevaluate their use of report classes. In compat mode it was possible to provide reporting granularity by placing each address space into its own control performance group. With the design of service classes which encourages the installation to place more address spaces into a single service class the granularity of reporting is provided by report classes. Installations who did not use report performance groups should spend some time evaluating how to build and use report classes.

Recommendations:

Identify the installation reporting requirements **before** the creation of a service definition. Use report service classes whenever possible to provide the needed granularity as the number of report service classes will not have an impact on WLM performance. If report service classes will not provide the granularity required, consider using the SMF type 30 records for the required reporting.

One method an installation can use for creating report classes is to create a naming convention, using R as the first character of the report class and then append the report class name with the performance group number of the work used in compatibility mode. For example, in the IPS/ICS we have three tasks A, B and C. Each is assigned to a different pgn. Lets say task A is assigned to pgn 15, task B pgn 30, and task C pgn 40. Now lets say in goal mode we can consolidate these three tasks into one service class because of their similar characteristics. Because we still want to report on these task individually we will need to define a report class for each task because without report classes they would all be reported under the same service class. So in our classification rules when we assign a service class to these tasks we also enter a report class name for each task. We can then use the following convention: for task A we will use a report class called RPGN15, for task B, RPGN30 and for task C, RPGN40. Again these names can be anything you choose but a the R will quickly identify it as a report class and the number tells us which pgn the task use to be associated with in compatibility mode.

___ Update and Validate Capacity Planning Tools and Reports:

Overview:

The migration to goal mode requires a reassessment of capacity planning tools and reports. For example, PGNs are replaced by service classes, and there may not be a one-to-one correspondence. A detailed investigation will be needed to insure all capacity planning tools and reports are properly updated to handle the new workload scheme.

Document References:

| Document | Section or Page Reference |
|----------|---------------------------|
| None | |

Issues and Differences:

The following are some items which should be considered when planning how tools and reports will be used in goal mode:

- Make sure the tool you use has been upgraded to handle goal mode. Because of the changes reporting, (for example, from PGNs to service classes), it may be necessary to install a new version of the tool. The underlying SMF records are also changed when migrating to goal mode. New subtypes of the RMF 72 record are created in goal mode.
- Be aware of any limitations imposed by your tool's implementation of goal mode processing. For example, we are aware of tools which supported compat mode report performance groups but do not support goal-mode report classes.
- PGNs may not map one-for-one to workloads and service classes. This will almost surely require workload groupings be reevaluated.
- Importance levels are not directly equivalent to dispatch priorities. The type of goal specified for the workload and the level of goal attainment are considered when dynamically setting dispatch priorities. This will have some bearing on certain capacity planning concepts such as Saturation Design Point.
- Dispatch priorities are no longer externalized. The fields where the dispatch priority could be found in the SMF Type 30 records in compatibility mode are now zeros in goal mode. This will affect those program routines which depend on a non-zero content in those fields.
- Develop some mechanism for handling the switch between service classes and performance groups, and the implied SMF record changes. This will impact reporting during the migration period.

Recommendations:

Do a reassessment of how capacity planning tools and reports are used when migrating to goal mode. Make sure tools and reports have been updated, to handle the new WLM constructs. Also, the migration to goal mode provides an excellent opportunity to reevaluate the succinctness and usability of the reports being generated.

End Of Chapter 1

Chapter 2 - Setting Goals

The second chapter of the migration guide presents some concepts to keep in mind when setting goals for the different workloads. The chapter is outlined in the following manner.

1. Building a Policy

This section will cover the overall structural concerns of building a robust WLM policy. There is a WSC sample policy available from the WLM home page which implements many of the concepts listed in this section.

2. Workload Characteristics

There is already a great deal of information about how to set goals for different kinds of work in an OS/390 system. The checklist items are for those workloads which may present different characteristics than the “normal” workload profile which is generally used to make goal mode recommendations. These ‘special’ considerations are outlined here, and an installation will need to review, if and how they apply to their workloads.

3. WLM Release 10 Enhancements

In OS/390 Release 10 and later, new options are available to help performance administrators protect critical work. Although applicable to several other subsystem types, CICS and IMS work will particularly benefit from the enhancements described in this section.

Building a Policy

___ Default Service Class For Unknown Work:

Overview:

As work enters a system it is assigned a service class according to the classification rules established in the active policy. If no match is found in the classification rules for the unit of work, and the rules do not specify a default service class, then the work is classified to the service class named SYSOTHER and assigned a goal of discretionary. The exception to this is started task work which defaults to a service class of SYSSTC when there are no matching rules in the policy and no default STC service class.

Documentation References:

| Document | Section or Page Reference |
|-------------|--|
| WLMPLAN | System Provided Service Classes |
| REDBOOK | Section 2.2.1.10 |
| WLMHOMEPAGE | Effective use of WLM Controls section 4.0.10 |

Issues and Differences:

If you let work fall through the classification rules and into the SYSOTHER service class those tasks will run with a discretionary goal. Depending on the type of work you may have severe system problems as well as having what may be less than optimal performance for the work.

For example, a task falls through the rules and is assigned to SYSOTHER and executes with a service class of discretionary. The work then requests and gets a DB2 lock. You are now faced with a task which is delaying other work because of their inability to get the lock. Since the holder of the lock is running in a discretionary service class, the CPU service it receives will be lower than all other work in the system and the task may not complete for a very long time, and hold onto the lock for a correspondingly long time and cause problems.

Recommendations:

Create a service class (i.e. service class = unclass) which can be used as the default service class in the classification rules for those subsystems for which you have not specifically assigned a service class. Assign this service class a goal other than discretionary. Assign a unique report service class as well when using this default so you can identify when work has fallen through the rules and which subsystem the work is related to. Establish a management reporting methodology to monitor, perhaps via RMF, when work appears in this default service class. Use the unique report class to identify the subsystem rules which were not matched and then classify the work to the appropriate service class. A good policy is to ensure no work is ever found running in either SYSOTHER or the installation default service class, (UNCLASS).

___ Manage the Number of Service Classes:

Overview:

There has been a lot of confusion over the appropriate number of service classes which should be defined. The upper limit of the number of service classes in a single service definition is 99. But a more practical limit to try and define is generally given as somewhere between 20-30. Contrary to a lot of “wisdom” the reason for defining a limited number of service classes has nothing to do with WLM overhead. Limiting the number of service classes does two things, it encourages responsiveness and helps ensure enough work is present in each service class to allow good decision making.

And actually what needs to be managed is not just the number of service classes, but also the number of service class periods.

Documentation References:

| Document | Section or Page Reference |
|-----------------|----------------------------------|
| None | |

Issues and Differences:

In compat mode fine granularity of work could be achieved by placing individual work units into distinct performance groups. Because of the way work is managed in WLM this is not practical in goal mode. The biggest issue will be how to get granularity of reporting, either through use of report classes or use of the SMF 30 records.

Recommendations:

Review the number of service classes to ensure there are enough completion's to provide good information for decision making for response time based service classes. For velocity goals ensure there is enough ready work to also provide good information. Be most careful with the number of started task service classes, generally aim for 2-3 service classes, plus SYSSTC and SYSTEM.

___ Requirement For Spare Service Classes:

Overview:

As you migrate to goal mode and execute work in the system there may be times when some movement of work to other service classes may be necessary. The reasons for the need to move work from normally assigned service classes are many and varied. It can be done because the throughput requirements are not being met with the currently defined service class or perhaps the goals themselves are not properly defined and need some fine tuning, or a unit of work was not classified properly.

Documentation References:

| Document | Section or Page Reference |
|-------------|---|
| CHECKLIST | Version Management for Service Definitions Secure the WLM Policy |
| WLMHOMEPAGE | WLM Quickstart Policy Article |

Issues and Differences:

This movement or reclassification is typically done using the reset command or over typing the service class field in SDSF. However, to use this method requires the new service class to be predefined in the active policy. If this is not the case, then you need to use the WLM application to modify the policy with the new service class, and install and activate the new policy, keeping in mind the management methodology you've established for doing this.

Recommendations:

In order to avoid having to modify, install, and activate a new policy whenever a change to a new service class is required, especially during the migration period when goals for work may still need some adjusting, we recommend you define a few "spare" service classes which normally have no work assigned to them in the classification rules. You can then use the reset command to assign work to these service classes if necessary. A typical use of this method would be to create a service class with a sufficiently high velocity which can be used for "Emergency" batch jobs needing to be pushed through the system ASAP or a special TSO service class for emergencies.

Practices such as moving work from one service class to another service class not especially designed may hurt performance and capacity reporting in the long run. Refrain from moving work into a service class because it pollutes the performance data for this service class. By placing manually adjusted work into special service classes the performance specialist can get a clearer picture of what work may need longer term reclassification to a different service class and can get a better feel for where resource shortage exist. See the section on reclassification of work into other service classes in the Train Operator section in Chapter 1.

Workload Considerations

___ How Do I Manage STC's:

Overview:

In WLM goal mode the system provides a service class called SYSSTC which is the default service class for started tasks. This service class has a fixed dispatching priority of 254, a fixed and high (MPL), multiprogramming level. Another WLM defined service class, SYSTEM, is used for system tasks which have the privileged and high dispatch attributes. Work in SYSTEM runs at a dispatch priority 255, and has a high and fixed MPL.

Started tasks are generally characterized as having large amounts of idle time, and few if any completions. Generally started tasks are server type tasks, and do something only when asked, and as a result will frequently enter wait states. Work in an idle period is not evaluated by WLM, and so much of the time a started task will not be giving information to WLM about the best way to treat this work.

Documentation References:

| Document | Section or Page Reference |
|-------------|---|
| REDBOOK | Section 4.4.3.2 |
| WLMHOMEPAGE | Effective use of WLM Controls section 8.0 |
| | WLM Quickstart Policy Article |
| WLMPLAN | Using a SYSPLEX with Systems at OS/390R4 Later |
| DB2ADMIN | DB2 Administration Guide: Recommendations for WLM |

Issues and Differences:

Customers have several options when deciding on how to classify started task work. They can classify many of the small CPU intensive started tasks to SYSSTC. This provides the option of not forcing WLM to manage these tasks from a CPU point of view. This will ensure the system work, needed to run the system, is generally treated very favorably.

Because of the high dispatching priority of work in SYSSTC, some care should be exercised about CPU intensive tasks running at this priority. The set of tasks which fits this bill will vary from system to system, not just by installation. Exercise caution with started tasks which use more than 10% of a CP, or have an unfortunate habit of looping. These types of started tasks may be better suited to a installation defined service class rather than SYSSTC. In a managed service class, WLM will has more control to ensure system started tasks and high priority online environments coexist.

Work in service class SYSSTC does not get to take advantage of WLM storage management policies. The performance analyst should review which started tasks have need for more aggressive storage controls.

Classification of started tasks is generally one of the more onerous parts of the migration processes. Use caution. Put the majority of server tasks in SYSSTC. Numerous performance problems have been caused by placing critical system functions in other service classes without enough CPU resource.

Recommendations:

Let all system tasks which want to go to SYSTEM go to SYSTEM. This is determined by the program properties table (PPT) and the SCHEDxx member of parmlib. If the SYST attribute is specified or the address space itself is created with the high dispatch attribute, then the task will be assigned the the appropriate service class. Don't second guess. Identify your key server started tasks, and if the CPU requirements are reasonable and the storage situation is good, put them in SYSSTC. Use the SPM rules in the STC classification rules to enforce this convention. Check after the migration to see how much CPU is being used by the SYSSTC service class. If it's a lot, or if you start to see some online delay, then adjust SYSSTC by taking some of the heavier CPU users out of the service class. Review the CPU usage patterns across multiple periods, such as peak hours, and the batch window. Review it after the migration to get a comfort level the SYSSTC service class is behaving reasonably, (not CPU spiking, and not paging).

For the rest of the started tasks we recommend you create as few service classes as possible. Try for 2. Maybe slice them this way:

- CPU intensive and important
- All other started task work

Give the service classes a velocity goal. Give the important one a high velocity and a high importance, give the other a medium velocity and a high importance.

You should not execute your transaction managers such as CICS and IMS in SYSTEM or SYSSTC. Also an exception to the SPM rules is DB2. These address spaces are marked in the PPT with the SYST attribute and by default will be classified into SYSSTC. Only the IRLM task should go into SYSSTC. All the other DB2 address spaces need to be manually classified into a high importance/high velocity service class and specified BEFORE the SPM rules. See the DB2 Administration Guide for more information on WLM classification of DB2 and the WLM planning manual on the SPM rules.

We strongly recommend you review the information available in the documentation on the SPM qualifier type as it relates to started task rules. Another good check is to review the PPT entries in SCHEDxx to get an understanding of what other products in the system are defined with the SYST attribute as these customization actions can impact the default classification of started tasks. The tasks of JES2AUX, TRACE, SYSBMAS, PCAUTH, VMCF, TNF, PORTMAP, ROUED, ANTSAS00 and perhaps a few others will not be classified to any service class by the SPM rules so you need to assign these to the appropriate service class which is generally, SYSSTC.

Handling Different Populations of TSO:

Overview:

TSO performance groups on compat mode systems often have multiple periods defined. The initial period will have a high dispatch priority while subsequent periods will drop in priority. Goal mode provides different service objectives for TSO workloads. These include percentile or average response times and velocities. Since WLM makes dynamic tuning decisions depending on the objectives specified, it is critical to define the correct type of service objective for all TSO workloads. This includes the default application development communities as well as the CADAM, CATIA, systems programmer and UNIX Shell communities.

Documentation References:

| Document | Section or Page Reference |
|-----------------|--|
| InitRef | 40.2 Statements/Parameters for IEAIPSxx |
| WLMPLAN | Defining Service Classes and Performance Goals |

Issues and Differences:

Some installations have used the UNT=R parameter in the performance group definitions for the TSO workloads. This parameter is no longer available in goal mode, so the DUR parameter must be changed from a time value to a service unit value for these workloads.

It is not uncommon for a small group of users to be associated with a unique control performance group. The performance group may provide a higher set of performance objectives than the default TSO performance group. During a migration to goal mode, this performance group must be evaluated. In order for WLM to provide responsive performance assistance to a service class period with a response time objective, there must be a sufficient number of ended transactions. If the small group of users in the performance group do not generate enough ended transactions, (usually several per minute), WLM may not respond to performance issues with the new service class in a timely manner.

The default TSO service class typically contains multiple periods. The first period will use a percentage response time objective with a high importance. Subsequent periods will have less aggressive objectives, (often changing to a velocity objective), and a reduced importance level. The use of a response time objective raises some interesting issues with some environments accessed through TSO. These environments include CADAM, CATIA, HSM recalls, and the UNIX Shell.

CADAM and CATIA workloads often share the transaction workload between the terminal control unit and the MVS system. As such, it is difficult to define the requirements for a host processor transaction for these workloads using a response time value based on service units.

HSM provides space management and can migrate data sets to less active volumes. If a TSO user touches a migrated data set, service units are not charged to the TSO user while HSM retrieves the data set, but the elapsed time to retrieve the data set does apply to the response time

as measured by WLM. This will cause the user to stay in a period for a long period of time because the duration is not exceeded because the TSO transaction is not accumulating service.

If the UNIX shell is entered directly from the TSO prompt, users will notice a period of time from the response to a shell command until the UNIX state changes from a RUNNING state to the INPUT state. It is important to understand this time delay, (which can be up to 15 seconds), is part of the transaction response time as measured by WLM.

Each of these environments can impact the response time metrics achieved on the system.

Recommendations:

The default TSO service class should have multiple periods. The minimum should be specified, two will be sufficient for most installations. The first period should use a response time objective with a high importance. This will provide good performance for the trivial TSO transactions. The majority of TSO transactions often complete in first period and require a minimal amount of resource. This will keep the vast majority of TSO users happy at minimal cost. The later periods should use a velocity objective and the importance will probably be reduced. A discretionary goal is not recommended for TSO periods for the default TSO service definition as ultimately there is a user on the other side of the terminal. Instead, a low velocity is probably a better way to limit the amount of resource consumed by complex TSO.

If a small community of TSO users must have their own service class and their requirements do not generate enough ended transactions for the use of a response time objective, a velocity may be a better choice for first period. If this group of users is of sufficient importance, they may have multiple periods with descending velocities but a higher importance level for all periods, than normal TSO.

Because the CADAM and CATIA workloads normally require fast response times and the host transaction characteristics are hard to quantify, these workloads are often best served with a single period service class with a high velocity and high importance, always balanced by the customer requirement for the workload. If these workloads impact the performance of other workloads on the OS/390 system, multiple periods can be defined with a reduced velocity but higher importance for each subsequent period.

The TSO community using the UNIX shell will often best be served with a service class made up of multiple periods and using velocities instead of a response time objective. This is due to the delay between command response and transition from running to input mode as discussed above.

The above recommendations are made assuming each TSO community runs a single type of work and can be identified. If the TSO communities can not be isolated and the non typical TSO users make up a small percent of the transaction rate, then the entire community can often be supported with a single TSO service class. This service class should be the default TSO service class discussed above with a percentile response time objective for first period. The normal, application development type TSO workload normally has a well behaved first period and will provide a high percentage of consistent response times. Using a lower percentile value, such as

70% or some other value, may allow the other workloads to not have a significant impact on the PI for the period. If the different communities are consolidated into a single service class, the installation should expect to spend some time fine tuning the TSO service class during the goal mode migration.

___ TSO Emergency ID:

Overview:

Installation often have a set of emergency ids used for logging on TSO when there are CPU loops. This id is often defined at the very highest dispatching priority, and is set with a single period definition. This capability can still be provided in goal mode.

Documentation References:

| Document | Section or Page Reference |
|-----------|---------------------------|
| CHECKLIST | TSO Workloads |

Issues and Differences:

There are two ways to provide a high priority TSO environment for emergency use. If a performance boost is needed and the TSO user is logged on then an operator command can be used to move a TSO userid(s) to SYSSTC while the emergency exists. Another choice is to move the TSO user to a SPARE service class with a high velocity and a high importance. If a user defined service class is used, (not SYSSTC), then the service class should be defined as a single period, there should not be any need to cause emergency work to transition periods.

Care will be needed to understand the performance impact emergency TSO access may make. If the issue being worked on is not related to the high priority online environment, but deals with other workloads then putting TSO work which may require sustained amounts of CPU at SYSSTC may in fact impact the onlines. In this case the use of a spare service class, with business objectives below the online may be a better alternative.

Recommendations:

Review the current requirements for a TSO emergency id. If no such use for this type of TSO support exists now it will not be needed just because of a goal mode migration. If such a capability is needed then ensure all the parts of the TSO logon path receives good performance including TCAS.

___ Discretionary Work and Capping:

Overview:

New support was introduced in R6 which changed the processing of discretionary workloads. With this new support, work which is over achieving can now be capped, hence freeing CPU which can be used for discretionary work. Workloads which are eligible for being capped are service classes with velocity goals of ≤ 30 or response time goals of >1 minute. The performance index, (PI), of the capped work will be kept between .7 and .81

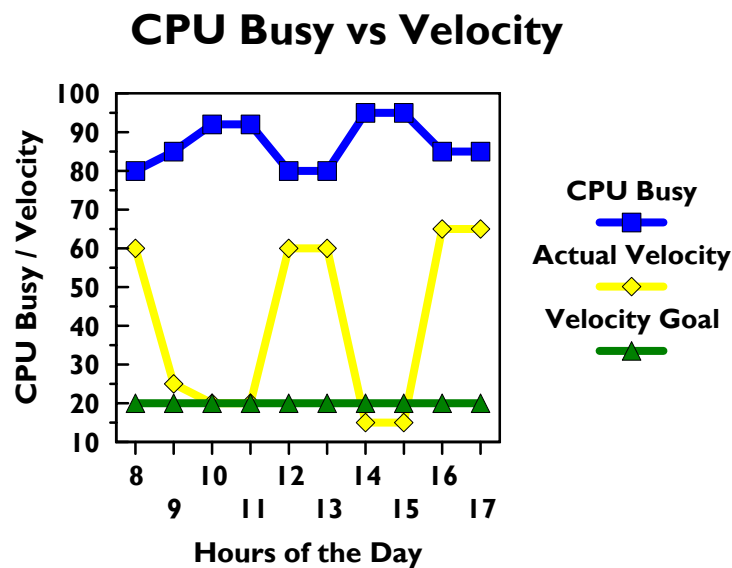
Documentation References:

| Document | Section or Page Reference |
|----------|--|
| REDBOOK | Section 2.2.1.10 WLM/SRM Discretionary Goal Management |

Issues and Differences:

Installations need to be aware of the inherent fluctuation in velocity goals. It's easiest to explain with an example. If there is a production online and production batch in a system it may be necessary to describe the production batch velocity goal using peak hour information to ensure the production batch workload has some throughput during peak hours, but doesn't impact the online workloads.

Using the workloads as represented below, the prod batch would be set to a velocity goal of 20 to handle the morning and afternoon peaks. Many systems have workload arrival patterns where the online workloads recede during a lunch period, during which time production batch fills the void. It is during these time periods the production batch work gets more access to CPU and the batch work "catches up".



If a system is designed in this manner it may not be a good idea to allow discretionary work to be defined in the system. The production batch work's velocity will rise when there is more CPU and hence this work may now be eligible to be capped in favor of discretionary work if present. This may cause the overall throughput of the production batch work to suffer, as we allow discretionary batch to run.

Recommendations:

Plan the use of discretionary work carefully. It may be necessary to raise some production batch work velocity values to slightly above 30, (31 or 32), to disable the capping from impacting workloads. Or rather than using discretionary use all velocity values with low importance work having very low importance (5), and a low importance (5). Caution should be used in classifying into discretionary such tasks as DDF work and batch jobs which require DB2 locks. If these tasks were to obtain a lock and not be able to get dispatched due to higher importance work saturating the system then serious consequences could occur.

There is material in WLM literature which states systems run better with discretionary work present in the system and so this type of work should always be defined. The WSC knows of no instance where this has been proven to be true, though we understand this statement is being made due to the perceived impact of mean time to wait dispatching algorithms. What should be stressed is enough ready work should be allowed in the system (In and Ready) to allow the processor to be 100% busy. This allows gaps in arrival rates of higher importance work to be filled with lower importance work, keeping throughput high.

The benefits of mean time to wait dispatching may have been outstripped by improvements in the I/O subsystems, and the improved processor and storage controller caching capabilities. And if mean time to wait dispatching priorities are beneficial then this support would need to be extended to velocity goals and response time goals, where workloads with differing business importance could also benefit.

___ The Use of Velocity Goals:

Overview:

A difficult performance goal to effectively manage is execution velocity. Velocity is a measure of acceptable delay. Most analysts don't normally evaluate workloads with the perspective of acceptable delay, rather they discuss work in terms of responsiveness or throughput. This checklist item will provide some guidelines for choosing velocity goals and applying them to your workloads.

Documentation References:

| Document | Section or Page Reference |
|-------------|---|
| WLMPLAN | Chapter 8: Defining Velocity Goals |
| REDBOOK | Chapter 4.2.10: Using Execution Velocity Goals |
| WLMHOMEPAGE | Documents: "Velocity Goals: What You Don't Know Can Hurt You" |

Some of the following information is extracted from the white paper by John Arwe (referenced in the table above), formally titled "MVS Workload Manager Velocity Goals: What You Don't Know Can Hurt You"; it will hereinafter be referred to as the *Velocity White Paper*.

Issues and Differences:

- Keep in mind when considering execution velocity goals (hereinafter referred to merely as *velocity* goals) the apparent disconnect between its name and what it actually measures: a **velocity goal** defines **how much delay is acceptable** for work when it is ready to run. Velocity is an artificial construct which attempts to provide a metric for comparing non-transactional oriented work to work which has more easily measurable and definable response time goals.
- It may be unrealistic to expect any work to attain a velocity even approaching 90-100%.
- The same work will almost surely achieve different velocities if it were run multiple times. Factors such as CPU speed, number of available engines, and workload mix all affect the resulting velocity of a piece of work.
-

A rather extreme example is shown in the following table taken from the velocity white paper. It shows the effect of moving a multitasking DB2 workload and SYSSTC from a uni-processor 9021 to a 5-way:

| Processor Model | 9021-711 | 9021-952 |
|--------------------------|-----------------|-----------------|
| Number of CPUs | 1 | 5 |
| DB2 achieved velocity | 13.3 | 74.9 |
| SYSSTC achieved velocity | 17.4 | 74.7 |
| CPU Busy % | 89.75 | 88.6 |

- Because of the likelihood of the variability in achieved velocity, defining fine gradations in velocity goals is probably a futile effort. It is better to define **ranges** of velocity goals.
- Remember there exist the additional controls of importance levels and resource groups to “rein in” potential “resource hogs”. For example, it is perfectly acceptable to define a low importance but a high velocity goal to CPU-intensive work which must run and yet not be allowed to dominate the system.

Recommendations:

- Use response time goals, (which are more easily quantifiable and consistently measurable), whenever possible. However, some work, (such as long-running jobs which generate few samples), can only be defined using velocity goals.
- Consider moving “well-behaved” production applications which are consistently small consumers of CPU cycles into SYSSTC.
- For the remaining work, create a few “bands” of velocities. A good starting point might be to define three: low, medium, and high with goals of 10, 40, and 70 respectively. Observe how well this fits your workload mix, and adjust accordingly.
- Consider using low importance levels and/or resource groups for work with the potential to dominate system resources.

___ Multiple Period Batch:

Overview:

Multiple periods are possible in goal mode as well as in compat mode. Up to 8 performance group periods are possible. In most cases installation should ensure enough work is running in each period to ensure WLM treats the service class period appropriately.

Documentation References:

| Document | Section or Page Reference |
|----------|--------------------------------------|
| WLMPLAN | Chapter 8, Using Performance Periods |

Issues and Differences:

In compatibility mode it was possible to have multiple period batch with later periods having a higher dispatching priority than earlier periods. Below is an example of what an IPS specification for this would look like:

```
PGN=3 , ( DMN=3 , DP=M3 , DUR=10000 )          /* BATCH FIRST PERIOD */
          ( DMN=32 , DP=M2 , DUR=300000 )       /* BATCH SECOND PERIOD */
          ( DMN=33 , DP=F3 )                    /* BATCH THIRD PERIOD */
```

This type of specification causes last period to have a higher priority than first period. Installations will use this technique for a variety of reasons. One reason a specification like this is done is when the complexity of batch work cannot be determined by external attributes such as job class. Jobs with the same job class, and even the same job name, may have vastly different resource requirements. This happens in service bureau type organizations, where all clients run the same streams of batch work, and the resource requirements for the different streams is related primarily to the size of the client. So for instance one customer may run job XYZ against a file with 100,000 records, while another client would run the same job against a file with 10 million records.

Now as work enters the system it ages through to last period. Other work will enter the system and if enough work is running in the higher periods this work may dominate the processor and later period work, at lower dispatch priorities, may not finish because it doesn't get any cycles. This work has to finish even if it's more complex. So once the work proves it's very complex by falling through to last period the work is flushed from the system by giving it a high dispatch priority. This specification allows this type of work to complete and free resources being held by the work, such as data sets, tape drives, and the initiator.

Recommendations:

Continue using multiple period batch definitions. If you have the IEAIPSxx definitions like the one outlined above unpredictable results may occur in goal mode. Installations will need to review carefully how important this throughput technique is to the batch processing. Resource groups will not be a solution because resource groups apply to the service class and are not definable at the service class period level.

Note:

Prior to R10 it was possible to define a service class period where the later periods had a higher importance or a more stringent goal. New service definitions created for R10 or later will not allow the specification. See the section on R10 Enhancements in this chapter.

_____ PVLDP - They're Just Different:

Overview:

PVLDP is used to assign the privileged dispatch priority for work on a WLM compat mode system. It is specified in the IEAIPSxx member of parmlib and is no longer externalized in WLM goal mode. PVLDP is involved with enqueue promotion as well as it is the dispatch priority of initiators prior to job selection. These dispatch priorities will be dynamically calculated by a goal mode system.

Documentation References:

| Document | Section or Page Reference |
|-----------------|---|
| InitRef | 40.2 Statements/Parameters for IEAIPSxx |
| WLMPLAN | 9.2.1.2 System-Provided Service Classes |

Issues and Differences:

Enqueue promotion is the process used by OS/390 to aid in the resolution of enqueue delays. Enqueues are normally issued to allow only one unit of work to have control of a resource while the resource is being changed. It is possible for the unit of work which owns the enqueue to have a low dispatch priority and hold the resource for an extended period of time due to the inability to gain access to the CPU. If this condition occurs, OS/390 will detect it and can promote the dispatch priority of the enqueue holder to the value specified by PVLDP. The unit of work can run at the higher priority until it releases the enqueue or consumes the number of service units specified by the ERV parameter in IEAOPTxx.

Compatibility mode systems required the installation to define the value of PVLDP. WLM goal mode will dynamically calculate the value. This calculated value will be higher than a percentage of the work currently using the processor. This should guarantee the holder access to the CPU although the relative dispatch priority of the enqueue holder could be different from the compat mode system.

The PVLDP value is the dispatch priority of initiators prior to job selection. This will change when a system migrates to goal mode. Goal mode initiators are assigned the SYSSTC service class prior to job selection. This can also change the relative dispatch priority of initiators on goal mode systems versus the compat mode environments. This needs to be understood if the installation causes a large increase in path length for the initiator prior to job select. Examples of increased path length are complex RACF account code validation or compute intensive SMF exits. Some customers have found the need to keep the PVLDP value below the dispatch priority of CICS or other high performance address spaces to make sure the initiators do not pre-empt these workloads. This capability is not available in goal mode prior to APAR OW55344. This APAR now allows the installation to specify an option for initiator dispatch priority. See the APAR for details.

Recommendations:

The change from a static, externally specified dispatch priority to a dynamically calculated value is not expected to cause problems for installations moving to goal mode. However, the

relationship of the dispatch priority of initiators and enqueue promoted work may be different to the dispatch priority of other workloads on the goal mode system versus their relationship on the compat mode system.

The normal path length through the initiator is very short. Installations should minimize their use of any functions which can increase the path length of the initiator prior to job select. This is especially important for machines or LPARs with a small number of processors. If initiators executing in SYSSTC pose an issue for you then set the keyword value in IEAOPTxx appropriately as documented in the R10 Enhancement section.

Some customers have reported complex RACF accounting validation routines or compute intensive SMF exits have taken the processor(s) away from workloads such as CICS and other high importance workloads. Rework of their complex exits typically have resolved the problem.

Enclave Information:

Overview:

This section was created to help document the various places where you can reference information on enclaves, what they are, who uses them, how they function, and what accounting information, (SMF), is changed as a result of using enclaves. An enclave represents a “business unit of work”. It is managed separately from the address space in which it executes and can be assigned a performance objective in compatibility mode as well as goal mode.

Documentation References:

| Document | Section or Page Reference |
|-----------------|--|
| REDBOOK | Section 2.3 |
| WLMHOMEPAGE | Documents / pre-emptible SRBs/ WLM, DB2 and Enclaves |
| WLMPLAN | WLM Planning Manual at the R9 level or higher |
| SMF | Type 30 ,72 and 97 records |
| RMF MONITOR III | Overview section - detail reports |

Issues and Differences:

The SMF accounting process for enclaves needs to be understood. Enclave transaction counts and resource usage are recorded in the SMF 72 record for the enclave’s service class/pgn and report class/rpgn for accounting and charge back. Existing fields include data from enclave transactions.

There are no SMF type 30 records for enclaves themselves. Transaction counts and resource usage are recorded in the SMF 30 records of the owning address space. New SMF30 fields record enclave active time, transaction counts, CPU time in seconds, CPU service units, I/O time and I/O counts.

Resources consumed by the DDF transactions are no longer reported in the DDF address space’s SMF 72 record unless the enclaves are in the same service class/pgn as the DDF address space. The net result of DDF using enclaves is there will be many more transactions with a increase in CPU time/service in the service class/pgn where enclaves are running, and a corresponding decrease in the DDF address space SRB cpu time/service in the SMF 30 record.

OS/390 R9 introduced a new construct called multisystem enclaves. Before R9, the scope of an enclave was limited to a single system. Multisystem enclaves allows a work manager to extend the scope of an enclave to multiple systems in a parallel sysplex. Along with this also came the introduction of a new SMF record. CPU time used by enclaves on other systems is charged back to the originating jobs SMF 30 record. A new SMF record, record type 97, identifies CPU time used by enclaves on behalf of jobs on other systems. See the WLM planning manual.

Recommendations:

Familiarized yourself with the new SMF records. z/OS R1.3 provides the ability to reset enclaves via the SDSF R1.2. See the section on DDF Service class in this chapter for more information.

When to use I/O Priority Management

Overview:

I/O priority management is a service definition option which, when enabled, allows the I/O priorities to be managed separately from the dispatching priorities according to the goals of the work. This option is disabled by default in the service definition but must be enabled in order to exploit dynamic alias management of parallel access volumes (PAVs).

Documentation References:

| Document | Section or Page Reference |
|----------|---------------------------------|
| WLMPLAN | Chapter 8 and 11 |
| REDBOOK | Section 2.2.1.9 |
| WSCSHORT | WSC Short Subjects Presentation |

Issues and Differences:

As part of the calculation for the achieved velocity goal of a service class I/O priority management includes disconnect time as a component of I/O using. As a result, if your DASD subsystem contains high disconnect times, the velocity values reported will be higher than if I/O priority was disabled. Since the Performance Index, (and subsequent WLM decisions), are based on the goals set versus goals achieved you may need to create or adjust your velocity goals if you chose to enable this option.

Disconnect time is included as a using sample. More using samples hence the higher the velocity. However, in most DASD I/O high disconnect time is an indicator of an I/O issue, and may lead to performance problems because it is taking longer to get data. WLM will report this as improved performance, (a higher achieved velocity), at the same time the work is actually suffering.

Recommendations:

If an analysis of your DASD subsystem shows high disconnect times and little or no IOS queue time then we recommend you set this option to NO. If you have high IOS queue times and also have high disconnect times or wish to enable dynamic alias management then you will have to evaluate the benefit of turning I/O priority management on in your installation. You will need to set and monitor your service class velocities carefully. RMF workload activity reports will show the velocity calculated if I/O samples are included, look for the field listing velocity migration values. You may also want to consider using static PAVs, which do not require I/O priority management be enabled, versus dynamic alias management of PAVs.

NOTE: An APAR has been taken by IBM to eliminate the disconnect times from the velocity calculation used by WLM. The APAR is OW47667 and is applicable to R8 systems or higher. Once this APAR is applied the general recommendation is to turn on I/O priority management.

___ Using DB2 Subsystem Classification Rules:

Overview:

DB2 subsystem classification rules are used for DB2 sysplex parallel query support **only**. Other DB2 work, such as DB2 Stored Procedures or DB2 Distributed work, (DDF) do not use DB2 rules. The DDF workloads or stored procedure workloads use the rules associated with their controlling work managers.

Documentation References:

| Document | Section or Page Reference |
|----------|--|
| WLMPLAN | 9.0 Defining Classification Rules for Each Subsystem |
| REDBOOK1 | 5.1.1 - 5.2 DB2 Sysplex Query Parallelism |

Issues and Differences:

With Sysplex Query Parallelism, DB2 work must be classified potentially twice. Once on the originating system, (the coordinating system), and again on the remote system, (the assistant system(s)). On the originating system the work manager which introduces the query will be used to classify the queries which run on the originating system. So if DDF introduces the split query, DDF classification rules will be used on the originating system to classify the work, and this classification is used to run any of the split queries running on the originating system. The same is true for a TSO, Batch, CICS, or IMS work unit which issues the query request. The work manager for this work will use the appropriate classification rules and this will determine the WLM service class definition for work on the originating system.

Query work which is eligible to be split into multiple queries, and is eligible to have some of the queries run on remote systems within the parallel sysplex will need to have a new WLM classification done for the split query on the remote system. It is this remote query which would use DB2 rules. This happens because in a way DB2 on the remote system is the work manager on the system 'sponsoring' the remote split query. As a result the classification would need to be done using DB2 rules.

Recommendations:

Ensure all remote query work is classified appropriately. If not classified appropriately, remote split query work will be running in the default service class SYSOTHER. SYSOTHER work gets a discretionary goal, so the remote split query will run at the lowest dispatch priority in the system.

___ DDF Service Class:

Overview:

DDF is a feature of DB2 which allows a DB2 application to access data at other DB2 systems and at any remote relational database systems which supports DRDA. Introduced in MVS V5.2, and enhanced in OS/390 R3, an OS/390 dispatching construct called an enclave was created. DB2 V4.1 uses this construct and in the classification rules, under subsystem DDF, you can assign these enclaves to different service classes with different performance objectives.

Documentation References:

| Document | Section or Page Reference |
|-------------|---------------------------|
| REDBOOK | Section 2.3 and 4.4.6 |
| WLMHOMEPAGE | Preemptible SRBs |
| WLMHOMEPAGE | Enclaves the inside story |

Issues and Differences:

In compatibility mode if you have not set up your IPS/ICS appropriately the enclaves will run at the priority of the DB2 DIST address space. In goal mode, if you do not create any classification rules for the DDF subsystem, the DDF enclaves will be assigned to the service class of SYSOTHER, not the service class of DB2 DIST. This is not how it worked in compat mode. Unclassified work in compat mode ran at the dispatch priority of the DB2 DIST address space.

Enclave service class assignment can only be done through the classification rules in the policy*. Unlike other work, there are **NO** operator controls to reset the service class of an enclave once it is classified. The only way to change the service class of an enclave is to change the classification rules and install and activate a new policy. The new classification will take place when an enclave matching the classification rules starts. There is no operator control available to change the service class of a currently executing enclave.

*Note: WLM z/OS R1.3 and SDSF R2.12 provide the ability for users to now reset, quiesce and resume enclaves by overtyping the SrvClass field with a different service class or use the action code field. See APAR PQ50025 and PQ65547. Care should be taken when resetting enclaves as this new capability also allows you to reset dependent enclaves as well as independent enclaves. Dependent enclaves execute under the same service class as the task that called them. So for example, if a TSO user issues a SQL call to DB2 and an enclave is created for that request, the enclave executes under same service class as the TSO user. Thus if you quiesce or reset the enclave you are performing that action against the TSO user not just the SQL transaction.

The monitoring of the enclaves for accounting purposes has changed and depends on whether they are dependent or independent enclaves. For information on accounting services for enclaves review section 2.3.5 in the redbook.

Recommendations:

For setting up the service classes of DDF transactions we recommend you initially start by creating a default service class with 2 periods. Specify the first period as a response time goal and the second period as a velocity goal. Review the DDF statistics over time and then if warranted or if more individual control is needed add additional service classes as appropriate.

Be careful if you define DDF work as discretionary because DDF transactions can request DB2 locks and may cause performance problems if they don't receive enough service to complete in a timely fashion.

When setting up service classes for DDF work try and limit the number of service classes you define. Because some DDF applications produce infrequent arrivals review the RMF reports for the service classes and watch the number of completions, keeping in mind WLM is more responsive when there are more completions in an interval due to the sampling based nature of its resource adjustment algorithms. Review your fallback plans as well and be sure to review section 4.4.6.1 in the redbook to make sure your requests are classified appropriately in compatibility mode as well.

___ SPAS Service Class:

Overview:

DB2 stored procedure address spaces (SPAS) exploit WLM to manage the performance of individual client requests. WLM is involved with DB2 SPAS in three ways:

1. Managing performance objective of the DB2 enclaves
2. Controlling the number of active DB2 SPAS address spaces
3. Controlling the distribution of the stored procedure requests across the server address spaces.

Documentation References:

| Document | Section or Page Reference |
|-------------|---------------------------|
| REDBOOK | Section 5.2 |
| WLMHOMEPAGE | Enclaves the inside story |

Issues and Differences:

The classification of client requests depends on how the stored procedure is accessed. This process is described in the manuals but further clarification is often necessary. When the request comes in remotely, via DB2 DIST, this type of request is referred to as independent enclaves and is assigned a service class using the DDF rules.

Which qualifiers to use in classification depends on how the request is made to DB2 SPAS. If the first SQL statement issued by the DDF client is a CALL statement, use the procedure name (PR) work qualifier in your classification rules. If the first statement after the connect is not a CALL to a stored procedure, for instance an SQL SELECT, followed by a Stored procedures CALL statement, then DDF attributes such as plan name (PN) would be used to classify the work. The resulting service class would be use to process both the SQL SELECT statement and CALL.

Users of PeopleSoft applications often make extensive use of the correlation information (CI) for classification.

Recommendations:

Review with your database programmers and administrators how the DDF requests to the SPAS will be made so you can use the appropriate classification qualifiers. For your fallback plans be sure to review section 5.2.2.2 in the redbook to make sure your requests are classified appropriately in compatibility mode as well.

___ UNIX Service Class:

Overview:

When programs issue the fork or spawn C function or use the OS/390 callable service, WLM provides a new address space. These address spaces or processes as they are referred to in UNIX can be associated with a user, be a daemon process, which typically perform continuous or periodic system wide functions, or be a kernel process.

Controlling OMVS through WLM has implications across four subsystem type rules ; STC, OMVS, and TSO and JES.

Documentation References:

| Document | Section or Page Reference |
|-------------|---|
| REDBOOK1 | Section 5.4 |
| WLMHOMEPAGE | Effective use of WLM Controls section 9 |
| UNIXSSPLAN | Chapters 1,3, 16 |

Issues and Differences:

The subsystem classification of OMVS related tasks will differ depending on how the UNIX work enters the system. When using UNIX services in a TSO session by issuing the OMVS or the ISHELL command, as opposed to an rlogin, the installation should check the goals associated with first period TSO. Work in this type of TSO environment can run for a long time (20 seconds or so), without accumulating any service, resulting in high average response times for first period TSO. OMVS and BPXOINIT are system tasks and therefore use the STC classification rules. Daemons and kernel processes are typically started at system initialization time, but can also be initiated from a shell, from a cataloged procedure, or by executing the BPXBATCH program. This type of work is via the OMVS rules, except for BPXBATCH, which uses JES rules. User forked transactions are also classified using the OMVS classification rules.

Recommendations:

By default, (via the PPT), the OMVS and BPXOINIT address spaces will be classified into the SYSTEM service class. So using the SPM qualifier in the STC rules is the recommended approach rather than specifying a unique service class for these tasks. The FTP task is started under the OMVS rules, however.

Some processes are forked by the UNIX system services initialization process, BPXOINIT, and are classified via the OMVS rules. Those which execute under the userid of OMVSKERN should be intercepted by using the userid qualifier under the OMVS rules and assigned a velocity goal higher than other forked processes.

By default, the fork and spawn services set the jobname value, (which uses the transaction name qualifier), to the userid of the issuer and then adds a number, (1-9), to the end. You can use this qualifier under the OMVS rules to assign these user processes a velocity goal which is lower than the OMVSKERN userid forked processes. Due to the variable nature of UNIX workloads,

multiple period service classes, using velocities, are generally appropriate. You may switch to a percentile response time goal for first period as a tuning action at a later time if appropriate after reviewing the service class reports in RMF

For TSO, use a percentile response time goal rather than an average response time goal keeping in mind an OMVS environment can affect the goals for other TSO users based on response times. Consider isolating TSO OMVS users into a separate service class if there is going to be a sufficient number of these users.

If any UNIXs task will be started using the BPXBATCH interface be sure to setup appropriate JES rules.

Review your fallback plans as well and be sure to review section 5.4.3 in the redbook and the UNIX Systems Service Planning manual to make sure your UNIX transactions are classified appropriately in compatibility mode as well as goal mode.

Since the introduction of UNIX can be very variable it is a good idea to place checks in the performance reporting system to ensure work is running in the correct service classes, and no undefined work is running.

IWEB Service Class:

Overview:

The IWEB work requests include all requests from the world-wide-web being serviced by the Internet Connection Server (ICS), Domino Go Webserver, or IBM http Server for OS/390. As of OS/390 Release 6, these requests also include those handled by the Secure Sockets Layer (SSL). As of OS/390 Release 7, this also includes transactions handled by the Fast Response Cache Accelerator.

Documentation References:

| Document | Section or Page Reference |
|-----------------|---|
| WLMPLAN | Chapter 9 |
| SC24-4826 | HTTP Server Planning, Installing, and Using |

Issues and Differences:

Special Note for the Fast Response Cache Accelerator.

In OS/390 Release 7, there is a function called the Fast Response Cache Accelerator for high performance handling of cached static web pages. You must classify this work as described below. Otherwise, it will be assigned the default service class for IWEB work.

The transactions handled by the Cache Accelerator are all joined to a single, long-lived enclave. This enclave would be assigned a unique transaction class (as specified on the Webserver FRCAWLMParms directive). This transaction class should then be assigned to a service class with a single period and a velocity goal in the service policy under the IWEB subsystem type. Neither response time goals nor multiple periods are appropriate for this work, as WLM is not aware of the individual Cache Accelerator requests. (Because each individual transaction is so trivial, it would cost more resource to manage them than to just process them.) In RMF reports, you will see zero ended transactions for the Cache Accelerator service class (assuming you have no other work running in this service class), but you will see some amount of accumulated service for this single enclave.

Recommendations:

The transaction class qualifier (TC) is probably the most useful qualifier because of its flexibility. Transaction class is the arbitrary class name you specify in the ApplEnv directive in the Web server configuration file. You can use the filtering function in the Web server to assign transactions to transaction classes based on the requested URL. Then in turn, the transaction classes can be assigned unique service classes via the WLM policy using the transaction class qualifier.

For the IWEB rules (as well as others) it is highly useful to specify unique report classes for transactions that use the same service class. This allows you to better isolate various transactions and their performance and greatly help with any problem determination tasks and capacity planning.

___ CB (WebSphere Application Server) Service Class:

Overview:

WebSphere Application Server V.4.0.1 for z/OS uses the workload manager (WLM) to manage the performance of application server regions in z/OS. Each WebSphere transaction is classified by WLM according to the 'CB' service classification rules, is dispatched as a WLM enclave in a server region, and managed according to its service goal objectives.

Documentation References:

| Document | Section or Page Reference |
|--------------|---------------------------------------|
| REDBOOKS | |
| WLMPLAN | Chapter 9 |
| WSCHOME PAGE | Techdocs number TD100887 and TD100889 |

Issues and Differences:

WebSphere applications are deployed within a WebSphere generic "server." One or more "Server Instances" must be defined on one or more systems within the WebSphere "node." Each Server Instance consists of one "Control Region" and one or more "Server Regions". The Control Regions are started by an MVS operator command as MVS "Started Tasks", and Server Regions are started by WLM as needed.

The classification can be based on the following classification criteria:

- Server name (CN) – This is the "generic" server name.
- User ID assigned to the transaction (UI)
- Transaction class (TC)

The Server instance name (SI) can also be used for classification, but this is usually not very useful because you cannot control which server instance runs a transaction.

Server Region Classification

You should classify the WebSphere Application server regions to a high STC importance service class so that they can be initialized quickly when WLM determines they are needed. The service class chosen will be the WLM goal when Java Garbage Collection (GC) is running which can be CPU intensive. Place these server regions in the service class hierarchy to ensure that higher importance work such as WebSphere transactions, CICS, or IMS is not adversely impacted.

Number of Server Regions: See "Managing the Number of Application Server Regions with WebSphere V4.0.1 for z/OS and Workload Manager" TD100887 on Techdocs at <http://www.ibm.com/support/techdocs/TD100887> for tips on managing the number of WebSphere Application server regions.

Control Region Classification

There is a certain amount of processing in the WebSphere application control regions to receive work into the system, manage the HTTP transport handler, classify the work and do other

"housekeeping." Therefore, control regions should also be classified in "SYSSTC" or a high importance and velocity goal.

Recommendations:

For the CB rules (as well as others) it is highly useful to specify unique report classes for transactions that use the same service class. This allows you to better isolate various transactions and their performance and greatly help with any problem determination tasks and capacity planning.

NETV Subsystem:

Overview:

Netview 1.4 can now support enclaves. These are classified under the subsystem of NETV. APAR OW54858 provides more information regarding this function. Also the WLM web site at URL: <http://www-1.ibm.com/servers/eserver/zseries/zos/wlm/pdf/wlmNETV.pdf> has presentation material about this function.

Documentation References:

| Document | Section or Page Reference |
|-----------------|---|
| WLMHOMEPAGE | See the article under the section titled "Useful Information" |
| IBM SUPPORT | APAR OW54858 |
| SC31-8874 | Tivoli Netview For z/OS Installation |

Issues and Differences:

Please note that if you choose to implement this capability and do not follow the proper procedures for the NETVIEW tasks, they will default to the SYSOTHER service class. The web page and documentation has more information you need to review.

Recommendations:

If appropriate for your installation, implement the new services provided by the NETV subsystem. See the chapter in this guide on adding subsystems titled "I Can't Find All of the Subsystem Classification Rules I Need".

_____OSDI (Oracle) subsystem:

Overview:

Oracle release 8.1.7.3 and above now supports enclaves. You need to be sure you are at the OSDI level of Oracle. Management of Oracle transactions can be specified by defining a new subsystem named OSDI. Oracle Net has support for a keyword parameter named ENCLAVE. You can specify ENCLAVE(SESS) or ENCLAVE(CALL). The recommendation is to specify ENCLAVE(CALL) as this classifies a transaction every time a request arrives from the client much like DDF acts with THREADS(INACTIVE) and RELEASE(COMIT). Multiple periods and response time goals are appropriate for these types of transactions. If you specify ENCLAVE(SESS), a single long running enclave will be created. In this situation a single period velocity goal is appropriate. See Oracle documentation for more info.

Documentation References:

| Document | Section or Page Reference |
|-----------------|--|
| ORACLE | Appropriate documentation for Oracle release 8.1.7.3 and above |

Issues and Differences:

Prior to this new enhancement, the service class goal of the Oracle address spaces determined the dispatch priority of the transactions executing within the address spaces. With the exploitation of enclaves, you now have more granular control over the individual Oracle requests.

Recommendations:

Upgrade your version of ORACLE to the version that supports enclave classification. Add the OSDI subsystem to your WLM service definition. See the chapter in this guide on adding subsystems titled :”I Can’t Find All of the Subsystem Classification Rules I Need”

OS/390 R10 Enhancements:

Overview:

In OS/390 Release 10 and later, new options are available to help performance administrators protect critical work. Although applicable to several other subsystem types, CICS and IMS work will particularly benefit from these enhancements:

- 1) Long-term storage protection (available on Release 10 with APAR OW43810 installed)
- 2) Long-term CPU protection (available on Release 10 with APAR OW43855 installed)
- 3) Region Management -- exemption from transaction response time management (available release 10 with APAR OW43812 installed).
- 4) Additional classification criteria (qualifier types) in the classification rules.

Documentation References:

| Document | Section or Page Reference |
|-----------------|---|
| WLMPLAN | Chapter 12 |
| WLMHOMEPAGE | Protecting Your Loved Ones-WLM Goal Mode Enhancements R10 |

Issues and Differences:

CPU protection

Problem:

In goal mode, there is no way to guarantee the dispatch priority of "loved ones" is higher than other work. Usually this is only an issue for customers with very tightly monitored CICS/IMS regions.

Solution: Allow loved ones to be identified so that their dispatching priority is *always* above that of less important work.

CPU protection is always by service class

Storage protection

Problem:

In goal mode, there is no way to guarantee "loved ones" are storage isolated proactively.

Solution:

Allow loved ones to be identified and assign protective storage isolation.

Regions with low overnight activity are protected from page faults when interactive users return in the morning. Regions are always protected from sudden changes in storage demand (runaways). Storage protection is assigned in the CICS/IMS subsystem type classification rules. Since individual transactions cannot be managed, protection is for the entire service class.

Both CPU and Storage Protection are:

Available to any address space.

Allow protection of low activity regions that cannot be managed as servers

Allow protection of regions running conversational transactions

Guarantee preferential resource access

Available for CICS/IMS transaction service classes. You don't have to know transaction topology (which regions serve which transaction service classes), if the transaction topology changes, WLM adjusts to it. You protect the transaction service classes and WLM protects the regions serving them.

Region management

Problem:

In goal mode, CICS/IMS transaction management is all or nothing. You are unable to manage just the production regions by transactions and the test regions at the region level. Also sampling overhead is incurred even if data sampled is not used for workload management.

Solution:

Allow regions to be identified by how you wish them to be managed. Each region can be managed either by transaction response times or by velocity goals. When using region management the velocity regions are sampled less frequently. A new column on classification rules allows you to specify either transaction or region. Transaction is the default.

If you use region management by specifying region in the classification rules, you can also code service classes in the CICS/IMS rules and add transaction response time goals here. Then the RMF reports will give you information on the distribution of work. The CICS/IMS service class rules will not be used in this case for workload management since you have specified region as your method of management. The velocity goal service classes specified in the STC rules (or JES if batch jobs) will be used. This can be valuable in planning your service class definitions if and when you migrate to transaction management.

New classification attributes

New qualifiers have been added to the allow more specific classification of tasks. The following are the new qualifiers:

System name, system name group, sysplex name, scheduling environment name, and subsystem collection name.

The system name and system group name allow you to assign goals based on execution system.

The sysplex name allows you to assign goals based on execution sysplex. The sysplex name makes it easier to use a single service definition across multiple sysplexes.

The scheduling environment name is an enhancement which was requested by customers currently using scheduling environments.

The subsystem collection name is a general way to talk about "subsystem name" that includes the effect of subsystem clusters. Used with JES2 XCF group name, JES3 JESplex name or the subsystem collection name associated with the originator of the DB2 query. A single job class can now have multiple definitions within a JES2 JESplex or JES3 Global/Local plex.

Migration issues:

Exploitation of any new function above causes a new functionality level in the service definition (functionality level 011). Functionality level 011 also imposes new restrictions on the service definition. Service classes used for CICS and/or IMS transactions cannot be used by other subsystem types. Multiperiod service classes must have periods whose importances do not decrease, that is, later periods must not be more important than earlier ones.

Recommendation:

Generally speaking, the features of CPU protection and Storage protection should only be used if your individual circumstances require it. These functions are limit WLM's ability to dynamically manage all of the workloads. Most customers should not see a need for this specific type of protection. Use of region management and the new classification attributes however, may well provide more flexibility in your policy and allow you to manage your workload more effectively.

END OF DOCUMENT