SG24-2089-00

OS/390 Release 4 Implementation

December 1997

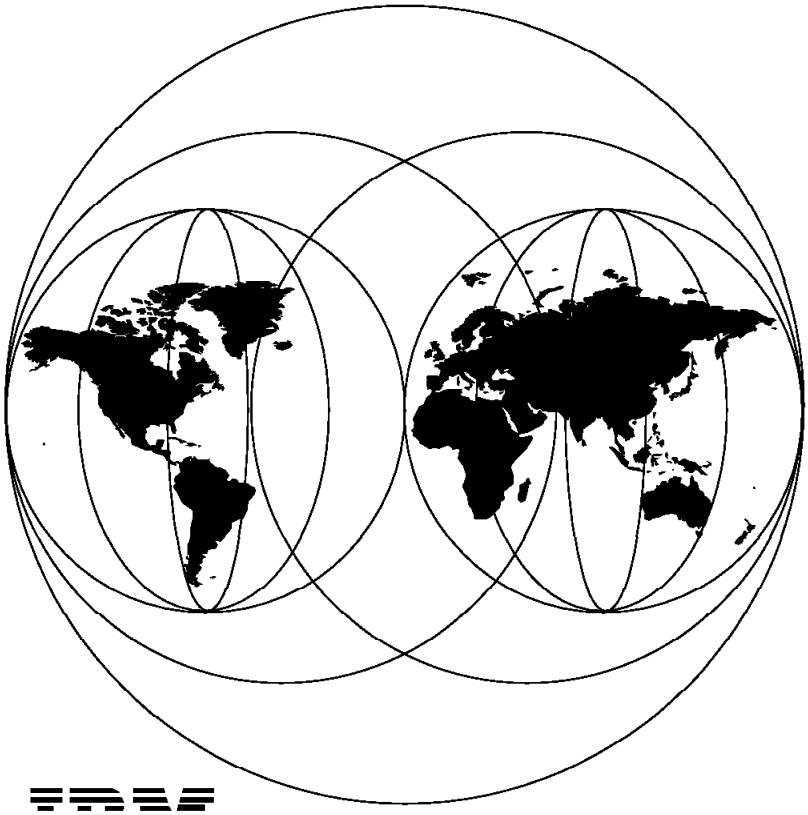

IBM

International Technical Support Organization
Poughkeepsie Center

International Technical Support Organization

SG24-2089-00

OS/390 Release 4 Implementation

December 1997

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special Notices” on page 241.

First Edition (December 1997)

This edition applies to Release 4 of OS/390, Program Number 5645-001 and to all subsequent releases and modifications until indicated in new editions.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HYJ Mail Station P099
522 South Road
Poughkeepsie, New York 12601-5400

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Figures

Tables

Preface

This redbook describes the new functions available with OS/390 Release 4. These new functional enhancements are for the following components: workload manager, JES2, SDSF, dynamic LPA, system logger, JES3, standalone dump, transaction dump, TSO/E, and DFSMShsm.

The enhancements to the workload manager are for batch jobs, which includes batch initiator management and resource affinity scheduling.

This redbook will help you to install, tailor and configure the new functions of OS/390 Release 4.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization Poughkeepsie Center.

Bob Haimowitz	ITSO Poughkeepsie
Paul Rogers	ITSO Poughkeepsie
Juha Vainikainen	IBM Finland
Joaquim e Silva	Standard Bank, South Africa
Ron Northrup	IBM South Africa
Redelf Janßen	IBM Germany
Dale Norman	IBM UK
Michael Ferguson	IBM Australia
Benno Aladjem	IBM France

Thanks to the following people for their invaluable contributions to this project:

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 259 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Home Pages at the following URLs:

  For Internet users `http://www.redbooks.ibm.com`
  For IBM Intranet users `http://w3.itso.ibm.com/redbooks`

- Send us a note at the following address:

  `redbook@vnet.ibm.com`

Chapter 1. OS/390 Release 4

OS/390 is a set of MVS base, open, client/server and applications-enabling functions orderable with one program number and delivered as one product.

OS/390 consists of base elements that deliver essential operating system functions. Most of these elements are products that have been available for some time and you will have been running some of them. To distinguish an element from its prior or preceding product, the term *root product* can be used. The OS/390 level of an element can be any of the following:

- A repackaging of the root product
- The root product with some additional function
- The root product unchanged

In addition to the base elements, OS/390 has optional features that have an affinity to the base. There are two types of features. One type is shipped with OS/390 whether you order the feature or not. These features support dynamic enablement. If you order the feature they are shipped enabled, otherwise they are shipped disabled. If you later want to use a disabled feature, let IBM know and enable it dynamically through parmlib. The second type of feature does not support dynamic enablement and is not shipped with OS/390 unless you specifically order it in addition to the base.

The OS/390 base elements and optional features deliver a set of services that provide an integrated set of MVS, UNIX, LAN, distributed computing, and application enablement services, as shown in Figure 1.

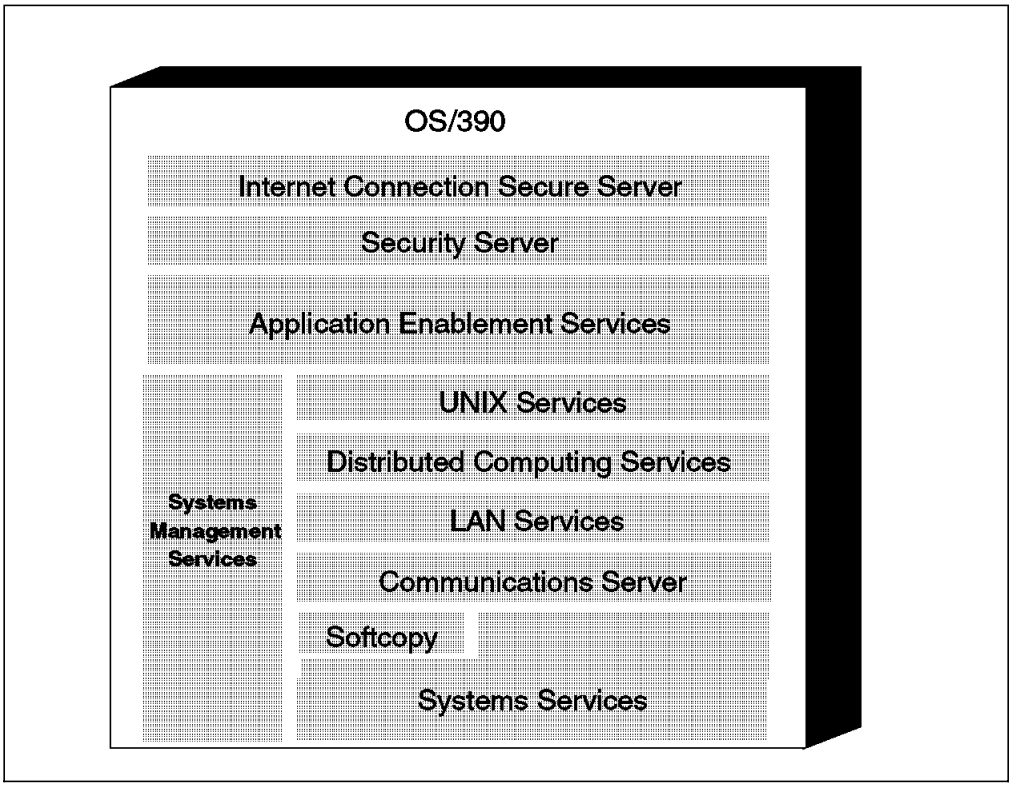

Figure 1. Overview of OS/390 Release Structure

1.1 Workload Manager (WLM)

OS/390 Release 4 provides new functions for the WLM component. These are:

- Scalable Web server
- Adaptive resource management
- WLM base enhancements

1.1.1 Scalable Web Server

This function supports scalability improvements for the MVS Web Server (ICS). The objective is to provide intelligent sysplex distribution of Web requests, through cooperation between WLM and the Domain Name System (DNS) server. When you place a DNS within an OS/390 sysplex, the DNS invokes WLM sysplex routing services to determine the best system to service a given client request. This provides functional equivalence with VTAM Generic Resources support that is currently available within SNA networks.

1.1.2 Adaptive Resource Management

This enhancement to the MVS workload management algorithms expands the scope and quality of WLM decision making by maintaining a local PI evaluation. The WLM/SRM logic that identifies service classes requiring attention is driven by the performance index (PI) for each service class, periodically calculated to reflect how well each class has been performing with respect to a user-defined goal for the class. Multiple indices are maintained for each class. A local PI is maintained for each system, and a sysplex PI reflects the aggregate for all systems in the configuration. The initial implementation of logic using the indices, which runs on each system in the sysplex, focuses first on the sysplex PI; classes meeting their goals from a sysplex perspective are deemed to be acceptable. The local PIs on each system are only examined if all classes of work are meeting their goals based upon the sysplex PIs; then additional logic is performed to adjust the highest importance service class not meeting goals on the local system(s). The net effect of this implementation is that conditions may exist where a high importance class may be meeting its goals from a sysplex perspective, but is performing very poorly on one or more of the individual systems. Though this satisfies the overall WLM sysplex philosophy, this logic should not permit a very important service class to consistently miss its goals on any individual system.

1.1.3 WLM Base Enhancements

The WLM base enhancements are designed to make migration to goal mode less disruptive. When a WLM policy is activated, either by an operator command or through the WLM administrative application, all historical information previously accumulated is discarded. Data structures are created describing the newly established goals, and then data accumulation begins in the context of the new service classes, service class periods, and service class period goals. Until such time as sufficient information is accumulated, significant algorithmic adjustments are not made. OS/390 Release 4 retains as much information as is easily possible across a policy activation, minimizing the extent of policy activation disruption.

Control of system service classes permits more administrative control over the system service classes internally created by WLM. It allows definition of WLM work classification rules that reference the SYSSTC service class, and enables reference to the SYSTEM and SYSSTC service classes via the RESET operator

command. Some restrictions are enforced to prevent total chaos. However, this new capability permits operational correction of errant classification rules that currently require an IPL to correct.

1.1.4 Batch Management

Batch management introduces initial goal-oriented management of JES2 batch. This support is in three functional areas:

- JES2 initiator address space management

  This support provides dynamic management of the quantity of batch initiators. Address spaces are created and destroyed based on the amount of work, the availability of system resources (such as CPU and storage), the existence of system constraints (such as auxiliary storage and ASIDs), the mix of work already in progress, and the ability of the system to achieve the WLM-defined goals for the various workloads being managed. This support removes existing JES2 manual controls over the quantity of initiator address spaces, including initialization parameters and operator control commands. A new level of job scheduling enhancements is required to compensate for the loss of manual control mechanisms, such as the JES2 operator command to force job initiation and minor improvements to existing job scheduling logic.

- JES2 resource affinity scheduling

  This is new support for managing asymmetric sysplex resource configurations. The initial application of this support is toward the scheduling of batch jobs, ensuring that jobs are only initiated on systems having the necessary system and/or user-defined resources. The primary objective is to obviate the need for explicit end user declaration (in JCL) of the specific system(s) having the resources needed for each job. With this new capability, the end users can declare (explicitly or implicitly) the specific resources needed by each job, leaving it up to system software to identify where those resources are available. This support also applies to the creation of DSOM and DCE server address spaces under the control of WLM, ensuring that servers are started only on systems where the necessary physical resource managers are active. Users may also use this capability to narrow the domain of systems eligible for supporting DCE and DSOM, thereby allowing additional customer control over server placement decisions.

- JES2 job sensitivity

  The WLM response time calculation for batch jobs is from the time a job is submitted until the job completes execution. This is the desired definition, except for cases where:

  - A job is held by an operator command
  - A job is held by the submitting end user (TYPRUN=HOLD)
  - A job is held because its associated job class is held by an operator command

  In Release 4, the response time calculation for these three situations is changed by starting the response time clock when such jobs are made eligible for scheduling.

1.2 OpenEdition Enhancements

OS/390 Release 4 provides new enhancements for OpenEdition which include:

- File caching
- Fork and spawn performance
- File system
- Dump exploitation
- HFS files

1.2.1 File Caching

This new capability allows files that are selected by the user to be cached in virtual storage in a dataspace associated with the kernel. A considerable reduction in path length, multiple concurrent use, and faster access improve performance for frequently used read-only files.

Significant performance benefits may be obtained by caching the data. Many files that are read only and are accessed with high frequency will benefit, such as C header files. Internet and intranet applications will benefit, because numerous files are accessed frequently but rarely updated.

1.2.2 Fork and Spawn Performance

OpenEdition fork and spawn performance is improved as the workload manager (WLM) is called to create and schedule work for fork and spawn callable services, instead of APPC, as done in the previous release of OS/390.

Performance may improve and goal-oriented services provided by WLM can be used for resource control.

1.2.3 File System

Previously, the OpenEdition Run Time Library (RTL) had added most of the C sockets interface from MVS TCP/IP. This included SRB mode support specifically for DB2, for a subset of socket functions, and basic asynchronous I/O capability for authorized programs. Now, for OS/390 Release 4, more function provides this capability for additional uses. All the sockets callable services are enabled for SRB mode callers, and the ASYNCIO callable service is extended to support asynchronous I/O for general users, which should improve the efficiency when using the converged sockets of OpenEdition.

1.2.4 Dump Exploitation

Prior to OS/390 Release 4, SVC dumps under the control of a non OpenEdition program did not include the OpenEdition CTRACE data. If OpenEdition is involved, this data may be critical in determining the problem. An OpenEdition SVC dump exit is now provided, which makes sure that OpenEdition CTRACE data is dumped for any non-OpenEdition SVC dumps that include a newly created address space or the OpenEdition kernel address space.

Note: Serviceability may improve for debuggers solving problems that involve non OpenEdition programs that call and use OpenEdition functions.

1.2.5 HFS Files

Currently, in order to run as program controlled or as APF-authorized, the executable file must reside in a PDS. Now OpenEdition allows these HFS executable files to be marked as able to run APF authorized or as a program controlled executable.

In addition, an extended file attribute is available to prevent newly spawned processes from running in a shared address space.

These enhancements have more flexibility regarding where to place programs that reside in the HFS file system.

1.3 Dynamic Link Pack Area

A new facility, dynamic link pack area (dynamic LPA), is available with OS/390 Release 4. It allows products to add additional modules to the link pack area after the LPA has been created. The modules can be added from either a PDS or a PDSE.

A re-IPL is currently required if a product containing LPA modules is installed for use. A product that utilizes dynamic LPA can be installed and used without the re-IPL.

Modules can be added to and deleted from dynamic LPA, and can have their status displayed. There is both an operator and a programming interface to manipulate dynamic LPA.

1.4 Enhancements in TSO/E

Starting with OS/390 Release 4, source code for TSO/E is no longer available in the optional material. With OS/390 Release 4, TSO/E support for Parallel Sysplex systems is enhanced by the following functions.

1.4.1 TSO/E SEND Command

The TSO/E SEND command is enhanced to allow the TSO user to send messages to other users independent of which system within the sysplex the receiving user is connected to. The message can also be sent to a specific operator, or the master console operator, or a specific operator console.

The operator SEND command now allows the operator to send a message to all users in the sysplex, or all users on a specific system in the sysplex, or all users logged on to a subset of the system in the sysplex.

1.4.2 TPUT Interface

The TPUT interface has been enhanced to forward any message issued via the TPUT to the proper system where the user is logged on to in the sysplex.

1.4.3 PARMLIB Command

The PARMLIB command is enhanced to allow the user to direct the PARMLIB command to all systems in the sysplex, or a single system, or a subset of the systems.

1.4.4 Send Statement

A new parameter on the SEND statement of the IKJTSOxx PARMLIB member allows the specification of the userlog data set size, eliminating the need to code a user exit.

1.4.5 CALL and TEST Commands

The CALL and TEST commands are enhanced to use the standard load module search sequence if no data set name is specified.

1.5 RMF Enhancements

RMF introduces the following enhancements in OS/390 Release 4.

1.5.1 Data Set Level Reporting

As an extension of the current spectrum of resource-oriented reports, RMF will offer reporting of data set usage. This new capability is one of the key requirements on many RMF users to help them identify device problems on a data set level. Three new Monitor III reports provide information on how individual data sets on a specific device are being utilized. By using this data, the customer can easily identify data sets that should be moved to another device to avoid contentions or bottlenecks.

1.5.2 Postprocessor Cache Reporting Enhancements

The RMF Postprocessor Cache support is enhanced by two new reports:

1. The Postprocessor Summary Report provides a comprehensive view about all control unit and DASD data belonging to the cache subsystem.
2. The Exception and Overview Report enables customer-defined reporting and provides data for spreadsheet converter processing.

1.5.3 Performance Monitoring of OS/390 Enhancements

The Performance Monitoring of OS/390 is enhanced by:

- A set of analysis functions providing “drop in” performance analysis simplifying the task of a performance analyst and enhancing his/her effectiveness through:
  - “Point and shoot” navigation
  - PerfDesks for performance analysis actions
- TCP/IP support allowing an alternative host connection to APPC.

1.6 SDSF Enhancements

SDSF is an efficient way to control JES2 job processing and devices. It allows you to monitor jobs while they are running and browse output without printing it. You can also browse the system logs, including the sysplex-wide operations log. SDSF does this with an interactive panel interface that provides immediate, up-to-date information about jobs, output, printers, initiators and other system resources.

OS/390 Release 4 SDSF includes all of the function present in OS/390 release 3 SDSF. OS/390 Release 4 SDSF provides enhancements in the following areas.

1.6.1 Scheduling Environment Display

OS/390 Release 4 Workload Manager (WLM) allows you to create groups of resources and give them a name. This *scheduling environment* is associated with a job through a new JCL statement.

SDSF provides a Scheduling Environment panel which simplifies the WLM support for scheduling environments. The Scheduling Environment panel gives a MAS-wide (or optionally, sysplex-wide) display that shows all the scheduling environments that are defined, a description of each, and the systems on which the scheduling environments are available.

1.6.2 Resource Display

The Resource panel can be accessed by typing an action character on the Scheduling Environment display. It shows the resources defined through WLM and the status of each (ON, OFF or RESET) for the selected scheduling environment. The Resource display can also be accessed as a primary display, with a command or from a pull-down, in which case all the resources in the MAS (or optionally, the sysplex) are displayed. You can modify the state of each resource on the system.

1.6.3 Job Class Display

OS/390 Release 4 SDSF allows you to manage JES2 and WLM classes using the new Job Class panel. The display shows the classes defined in the MAS, the status of the class, and the number of jobs waiting and active in each class. Action characters and overtypeable fields can be used to hold or release a class or modify characteristics of a class.

1.6.4 WLM-Managed Classes

With OS/390 Release 4, there are WLM initiators as well as JES2 initiators.

These are created dynamically by WLM and are displayed by SDSF on the DA display.

1.6.5 Job Information Pop-up

The new Job Information pop-up display shows information that might be responsible for delaying the processing of a job.

The Job Information pop-up displays information such as:

- Job name and ID
- Scheduling environment name

- Job class limit indicator
- Job class held indicator
- Estimated (average) time until execution
- System or systems which satisfy the resource requirements

1.6.6 Support for the $SJ Command

OS/390 Release 4 JES2 allows you to start a job that is in a class managed by WLM. SDSF has a corresponding action character, J, that can be issued from the ST and I displays.

1.6.7 Maximum Return Code

OS/390 Release 4 SDSF has a new column on the H , O, and ST displays which shows the highest return code returned for a job.

1.7 JES3 Enhancements for OS/390 Release 4

There are significant availability enhancements in this new version of JES3. These enhancements include:

- JES3 64K job number support

  The job number enhancement provides support to increase the maximum number of jobs in a JES3 complex from the current limit of 32,767 to 65,534.

- JES3 configuration improvements

  This provides support to improve system availability in an OS/390 Parallel Sysplex system by reducing and shortening outages needed to make JES3 Configuration changes. This support is comprised of 3 parts:

  – Dynamic update support
  – Refresh with hotstart support
  – Faster restart support

1.7.1 Dynamic Update Support

In OS/390 Release 4, JES3 supports dynamic update for adding or modifying initialization statements for the following statements:

- SNA RJP
  – RJPWS - SNA RJP work station characteristics
  – CONSOLE - SNA RJP consoles
  – DEVICE - SNA RJP devices (readers, printers and punches)
- VTAM Attached FSS printers
  – FSSDEF - Functional Subsystem Definition
  – DEVICE - for VTAM attached FSS printers

1.7.2 Refresh with Hotstart

One of the major problems with JES3 is the fact that it takes a warm start to change many of the JES3 initialization parameters. The warm start is very disruptive because not only must the JES3 global address space be brought down, but all processors in the JES3 complex must be IPLed.

In OS/390 R4, JES3 is changed to read the initialization stream during a hotstart without IPL to allow many of the parameters to be changed. A new start type called HOTSTART with REFRESH is created to read the initialization stream and process certain initialization statements (for example, FSSDEF, RJPWS, and non-execution DEVICE).

1.7.3 Faster Restart

There are occasions when you must restart the global processor, such as when you need to implement certain configuration changes or to apply service. While the global is down, jobs in execution may experience delays when they request global services such as spool space allocation, opening SYSOUT data sets and so on. The HOTSTART needed to make the configuration changes as described above will cause such delays.

JES3 RESTART processing is being enhanced to reduce the time it takes the JES3 global to reinitialize.

1.8 Security Server

The following enhancements are available with the 0S/390 Release 4 version of the Security Server.

1.8.1 Security Server and DB2

The Security Server for OS/390 Release 4 is provides a new function which lets you control access to DB2 objects using RACF profiles. This function is provided via a fully supported exit load module called the RACF/DB2 external security module. To take advantage of this new support, the module is designed to receive control from the DB2 access control authorization exit point, a new function in DB2 Version 5. The benefits allow for:

- Administration and auditing of access control from a single point of control
- Ability to define security rules before a DB2 object is created
- Ability to have security rules persist when a DB2 object is dropped
- Ability to control access to DB2 objects with generic profiles
- Flexibility to control access to DB2 objects for single or multiple subsystems with a single set of RACF profiles
- Ability to validate a user ID before permitting it access to a DB2 object
- Elimination of DB2 cascading revoke

1.8.2 OpenEdition DCE Security Server

In OS/390 Release 4, a DB2-based security registry is introduced for the DCE Security Server. This enhancement offers an optional replacement for both the in-storage registry and the HFS registry files. It allows commit and backout of updates to the registry, transaction logging, and log recovery. It is now possible to migrate an HFS registry to a DB2 registry and vice versa.

1.9 Language Environment

OS/390 R4 is provides a new function called Run-Time Library Services (RTLS). This function allows use of multiple levels of the Language Environment run-time library during application execution. It is no longer required to STEPLIB and maintain JCL to access different levels of LE run-time libraries. Instead, with the use of new system parameters in parmlib and new run-time options, it is now possible to dynamically control which level of LE run-time libraries are to be used on an individual application basis. This function also provides an alternative to placing the LE run-time library in the linklist concatenation.

1.10 OpenEdition DCE

The IBM OpenEdition DCE products are based upon the Distributed Computing Environment (DCE) technology, an integrated set of open, industry leading, distributed computing technologies licensed from the Open Group. DCE currently consists of a base set of services, including security and directory services, and a distributed file system component.

In Release 4, the OpenEdition DCE base services introduces a cell directory server on OS/390. This provides a complete set of DCE services on OS/390. This helps reduce the number of systems requiring administration and thereby lower administrative cost. Further, you can have the option of keeping all critical DCE cell servers on the OS/390 platform. This is important for running middle tier support servers on S/390.

OpenEdition DCE base services provides the following additional support and enhancements in OS/390 Release 4:

- Exploitation of current C/C++ compiler features and optimization is being provided by DCE, while allowing existing customer applications to run without modification.
- Improvements to the reliability, availability, and serviceability characteristics of DCE have been made, for example, through enhancements to the error-checking code which can improve the user′s ability to determine what actions are required in the event of a problem.
- Significant performance improvement can be achieved by a variety of enhancements, including for example, reducing the path length of routines that convert data from ASCII to EBCDIC.
- The number of languages that are National Language Support (NLS)-enabled is increased to 32. Messages continue to be translated into Japanese. OSF DCE 1.1 internationalization support is being activated in OS/390 DCE, which allows users to write applications that contain clients that can interoperate with servers running in different locales. These NLS items will make OS/390 OpenEdition DCE more usable on an international scale.
- Parallel Sysplex system enablement is being provided as a first step to allow clients to exploit DCE in a Parallel Sysplex system environment. Changes are being made to the client runtime support so that applications can be developed now which run without change when the server support is later provided.
- Introduction of hardware crypto support provides better performance over the current software encryption/decryption, if the ICSF software and ICRF hardware are installed. This support allows DCE to take advantage of the

hardware encryption/decryption function, and improved performance is realized particularly when dealing with packets of greater than 60 bytes.

1.11 Communications Server for OS/390 Release 4

Significant enhancements to the TCP/IP function used in the OS/390 OpenEdition environment provide improvements in:

- Performance
- Reliability
- Serviceability
- Network Management
- National Language Support

1.11.1 Improved Performance and Reliability

There is a new TCP/IP protocol stack for improved performance and reliability. The TCP/IP function in the communications server for OS/390 has been designed for native MVS to improve its performance and reliability. It capitalizes on OS/390 enhanced MVS services and has multiprocessing capability. These improvements are delivered in stages. The first is included in the communications server for OS/390 Release 4 and is directed at users of the OS/390 OpenEdition environment for UNIX-based applications.

Interim to delivering the final stage of the TCP/IP stack, support for non-OpenEdition applications will continue to be provided by the TCP/IP stack currently available in Communications Server for OS/390 Release 3 (TCP/IP V3R2). Both stacks will be included in Communications Server for OS/390 Release 4, and both can be run concurrently on the same MVS image. Communications Server for OS/390 Release 4 users will benefit from improved performance for their OpenEdition applications, while being able to continue running their non-OpenEdition applications.

In a subsequent release, these benefits will be available to both non-OpenEdition and OpenEdition applications on a single, converged, TCP/IP protocol stack in OS/390.

1.11.2 Serviceability with VTAM Device Drivers

The communications server for OS/390 Release 4 provides more synergy between the SNA and TCP/IP functions. The TCP/IP function will capitalize on even more of the proven device support offered by VTAM. The following functions have RAS improvements in OS/390 Release 4:

1.11.2.1 TCP/IP

TCP/IP now uses the MVS CTRACE for improved serviceability as follows:

- The trace facility is changed to trace more than 256 bytes per record.
- The default trace buffer size is 64K-128K to allow larger trace records and more trace information to remain in storage.
- TCP/IP takes advantage of the minimal trace option offered by CTRACE to always trace a subset of information that is always available.
- The IP packet trace continues to be traced to GTF.

- The new TCP/IP OpenEdition stack supports dynamic stack configuration to allow the deletion of definitions while TCP/IP is active. This enhances the S/390 servers′ ability to provide 24 hour, 7 days a week availability.

1.11.2.2 Configuration Enhancements

- Configuration of the TCP/IP address space is now independent of other TCP/IP stack processing.
- The VARY commands now allow dynamic changes to the system operation and network configuration without stopping and restarting the TCP/IP address space.
- Use of the OBEYLIST is migrated to RACF for configuration change authorization, APF authorization, or OpenEdition superuser status.

1.11.3 Improved Network Management and Routing

There are some network management protocol enhancements to provide a more robust network security environment.

- Support for Version 2 of SNMP is provided by TCP/IP OpenEdition.
- The Distributed Programming Interface (DPI) 2.0 library is provided for OpenEdition. This programming interface can be used to implement dynamic management information base (MIB) variables. A new DPI subagent is provided as part of the Communications Server for OS390 Release 4. DPI 2.0 allows easier subagent development than DPI 1.1.
- OpenEdition SNMP commands are provided.

1.11.3.1 OpenEdition Ehancements

The Tracert, Netstat, and Ping commands are enhanced for use in the OpenEdition environment.

ROUTED, a mechanism for dynamic routing in TCP/IP, is enhanced to run in the OpenEdition environment.

The TCP/IP Version 3 for OpenEdition MVS Applications Feature, previously a separate feature, will be integrated into the base TCP/IP product.

1.11.4 National Language Support (NLS)

New TCP/IP functional components, applications, and the stack will provide NLS basic enablement for language other than US English (except for bidirectional support).

1.12 Support for IBM Network Station

TCP/IP supports the IBM Network Station by delivering the code that runs on the IBM Network Station from the OS/390 via TFTP. This support is also available for TCP/IP 3.2 via a PTF.

Chapter 2. Workload Manager (WLM)

Workload management addresses the need for managing workload distribution, load balancing, and the sharing of finite computing resources among competing workloads. Workload Manager (WLM) *goal* mode refers to the management of resources according to the goals specified in an active workload management policy. Workload Manager *compatibility* mode refers to the management of resources according to the values in the IEAICSxx and IEAIPSxx members of SYS1.PARMLIB.

This chapter describes both the existing WLM functions in 2.2, "Introduction to Workload Manager" on page 15, and the enhancements made to the Workload Manager in OS/390 Version 2 Release 4.

2.1 WLM Release Overview

Changing from resource-based performance management to goal-oriented workload management has meant that a number of existing ideas relating to performance management have had to change. These changes in management philosophies, together with extensive product changes required to exploit and report on WLM activities, have been the primary reasons for the staged approach to the delivery of Workload Manager function.

The following sections summarize the WLM evolution so far.

2.1.1 WLM on MVS/ESA SP 5.1.0

The initial MVS Workload Manager release introduced the concept of managing performance from an administrative perspective, in conjunction with dynamic system resource management algorithms that control work in execution. Preliminary cooperation between the base control program and the key subsystem work managers was also introduced, enabling sysplex-wide specification and reporting of installation goals for work managed by IMS, CICS, TSO, JES2, JES3, MVS/APPC, and OpenEdition.

To ensure a level of comfort with such a change, switching to the new MVS Workload Manager environment included the capability to define the new goal-oriented externals (WLM policy) in parallel with the existing system externals (IEAIPSxx and IEAICSxx parmlib members). The ability to dynamically change to the new orientation (goal mode) and back (compatibility mode) if necessary, made it possible to experiment with the new function with relative safety.

2.1.2 WLM on MVS/ESA SP 5.2.0

The second MVS Workload Manager release contained a collection of services and functions supporting CICS and DB2 workload management, supporting new MVS parallel sysplex capabilities, and supporting other software vendors that provide MVS systems management tools.

Through cooperation with VTAM, sysplex-wide session placement decisions for DB2 and CICS can be turned over to WLM when the customer activates the VTAM Generic Resources support. Once enabled, sessions can be spread across multiple MVS system images in a manner transparent to end users. For

distributed access to a DB2 data sharing configuration, DB2 enabled classification and reporting of a DB2 transaction, where a transaction equated to the commit scope of the client. These improvements allowed MVS to manage and report on this new work consistent with other types of MVS work.

In support of DB2 data sharing, a new feature was introduced to spread distributed DB2 work across multiple DB2 instances using information provided by WLM. New WLM services were introduced, making it easier to integrate system management tools into the WLM environment. Using key features of these services, it is now possible to develop workstation-based administrative applications to enhance or replace the supplied WLM TSO/ISPF application.

2.1.3 WLM on OS/390 Release 3

The third MVS Workload Manager release expanded the range of work environments influenced by WLM to include sysplex distribution of connection requests for DCE, DSOM, TSO, and APPC/MVS. New function was provided for management of server address spaces supporting the DCE, DSOM, and DB2 Stored SQL Procedures environments, including dynamic server space startup, shutdown, and operational controls. The scope of WLM resource control was expanded to include sysplex-wide management of DASD I/O priorities, setting the stage for more comprehensive I/O performance management in future releases. Minor functional enhancements were made to address requirements received from customers exploiting the first two releases.

For further information, refer to *OS/390 Release 3 Implementation*, SG24-2067.

2.1.4 WLM on OS/390 Release 4

This fourth release of MVS Workload Manager provides the following enhancements:

- Scheduling environments, which provide a way to define resource requirements for different types of work, and to ensure that work is assigned only to systems that satisfy those resource requirements.

  Note: In OS/390 Version 2 Release 4, batch jobs submitted through JES2 are the only work that currently use scheduling environments.

- Batch initiator management, in which WLM can dynamically manage the number of initiators for one or more JES2 job classes to meet specific performance goals.
- A new work classification qualifier for job priority.
- Batch response times, reflected in the RMF workload activity report for jobs submitted with the TYPRUN=HOLD keyword parameter on the JOB statement, now exclude the amount of time spent in the held state.
- You are now able to use the SYSTEM and SYSSTC service class names within installation-defined work classification rules and on the RESET operator command.
- A new mechanism, in concert with the Domain Name System component of TCP/IP, is able to provide intelligent distribution of TCP/IP network connections across multiple systems in a sysplex.
- Policy activation is less disruptive, because only information that has changed is updated in the policy history.

2.2 Introduction to Workload Manager

MVS workload management addresses the need of a single-system sysplex or a multisystem sysplex for:

- Managing workload distribution
- Load balancing
- The distribution of computing resources to competing workloads

Workload management introduces a sysplex-wide view of performance administration, monitoring, and management of workloads running under MVS. The MVS Workload Manager uses XCF services to communicate between systems in a sysplex. There is no requirement for a coupling facility hardware unit; ESCON or parallel channel-to-channel (CTC) connections can be used.

An ISPF application, as shown in Figure 3, allows the service administrator:

- To enter the workload manager service definition into the system.
- To install the definition in a data set called the Workload Manager couple data set (WLM CDS). The MVS Workload Manager runs in a sysplex, with a Workload Manager function couple data set (WLM CDS) shared among all the systems. All systems are informed of any change in the active service definition. The workload manager couple data set allows high availability: in case of a failure in the path or in the disk itself, the alternate data set is automatically activated.
- To activate a definition for the sysplex

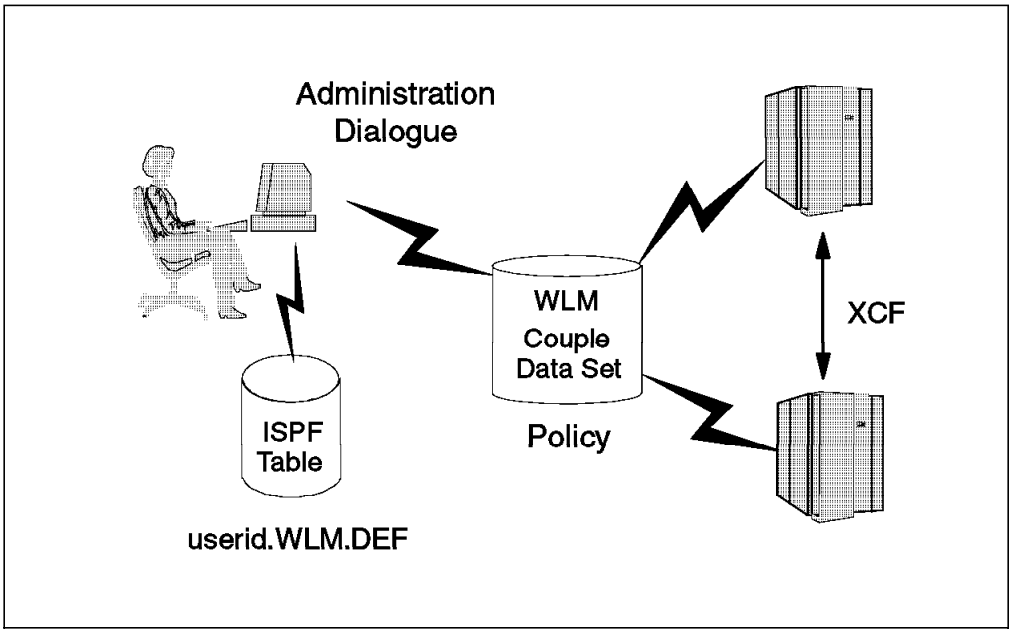

Figure 2. ISPF Application to Create WLM Policy

2.2.1 WLM Couple Data Set

The Workload Manager (WLM) couple data set holds the WLM policy, which allows you to define service goals for your workloads. The WLM couple data set can be formatted using the couple data set format utility, IXCL1DSU. Sample JCL to invoke the format utility, along with explanations of the information you need to specify when formatting each type of couple data set, is provided in *Setting up a Sysplex*.

Note: The WLM couple data set can also be formatted interactively using the WLM administrative application. See *OS/390 MVS Planning: Workload Management*.

See 2.11, “Migration Considerations” on page 53 for the details of using an OS/390 Release 4 WLM couple data set.

The WLM policy can be updated using the WLM ISPF Application on any system within a multisystem configuration. Using the dialog, the WLM service definition information is read from the WLM couple dataset. The required changes are made, and the Install option of the dialog is used to write the new information back to the WLM CDS. The updated policy can then be activated, either from the dialog or via an MVS command from the system console.

Following successful activation, that system then contacts each of the other active WLM address spaces within the sysplex, and they read the newly written active policy information. WLM then activates the new policy on each system. (activating a new policy on one system in the sysplex results in all other systems activating the same new policy simultaneously).

2.2.2 Workloads

A *workload* in WLM terms, is a set of related work items. A workload can consist of a collection of similar work types, such as all TSO users. Alternatively, a workload can be a grouping along business lines, such as “Loans,” which migh include all the CICS transactions, TSO user ids, and batch jobs, used by staff in the Loan department. WLM workloads are used solely for reporting purposes. Each workload consists of one or more Service Classes.

2.2.3 Service Class

A *service class* is similar in concept to a performance group. A service class represents a grouping of work with similar resource requirements and performance requirements, which can be server address spaces, batch jobs, or CICS/IMS transactions. Just as performance groups can be subdivided into smaller entities called Performance Group Periods, similarly, Service Classes can be subdivided into Service Class Periods. Just as in compatibility mode, with Service Classes, some work types are suited to breaking down into periods (such as TSO), while others (such as critical production batch), may not be. If a service class has more than one period, each period must have a duration, defined in numbers of service units, as in compatibility mode. Each service class can have one or more *periods*, and each service class period has a number of attributes.

2.2.4 Service Class Attributes

The *service class* or *service class period* is the level at which you define your objective to OS/390. Using all of the following attributes, you tell OS/390 what level of service is required for each work item and how important it is that the service levels are met.

2.2.4.1 Duration

The duration specifies the length of a period in terms of a number of service units. If the work does not complete within this number, the work moves to the next service class period. Do not specify a duration for the last period, or if the service class has only one period. Multiple service class periods are not available for CICS or IMS transactions.

See *OS/390 MVS Planning: Workload Management* for more information on subsystem support of multiple periods.

2.2.4.2 Goals

Each service class or service class period must have one of four different types of goals assigned to it. The goals may be expressed in four ways, as shown in Figure 3 and 2.2.4, "Service Class Attributes."

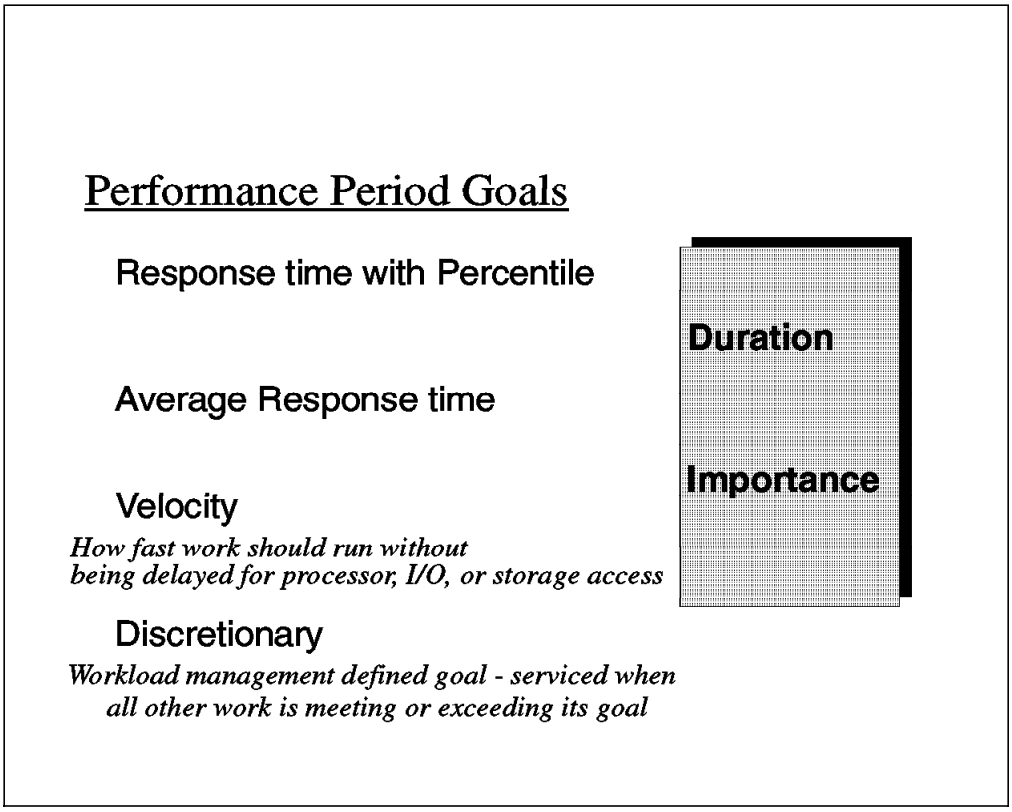

Figure 3. Workload Manager Definitions

The goals for each performance period may be expressed as:

Response time average Specify what you want the average response time to be.

Response time percentile Specify a response time and the percentage of total transactions you want to be below the specified response time, (for example, 90% of transactions are to complete in less than 1/2 second).

Velocity Specifies the extent to which delays to this class of work are acceptable.

Discretionary Work with this type of goal only receives service if all other service classes are meeting or exceeding their goals.

All work items are assigned a goal, as shown in Figure 4.

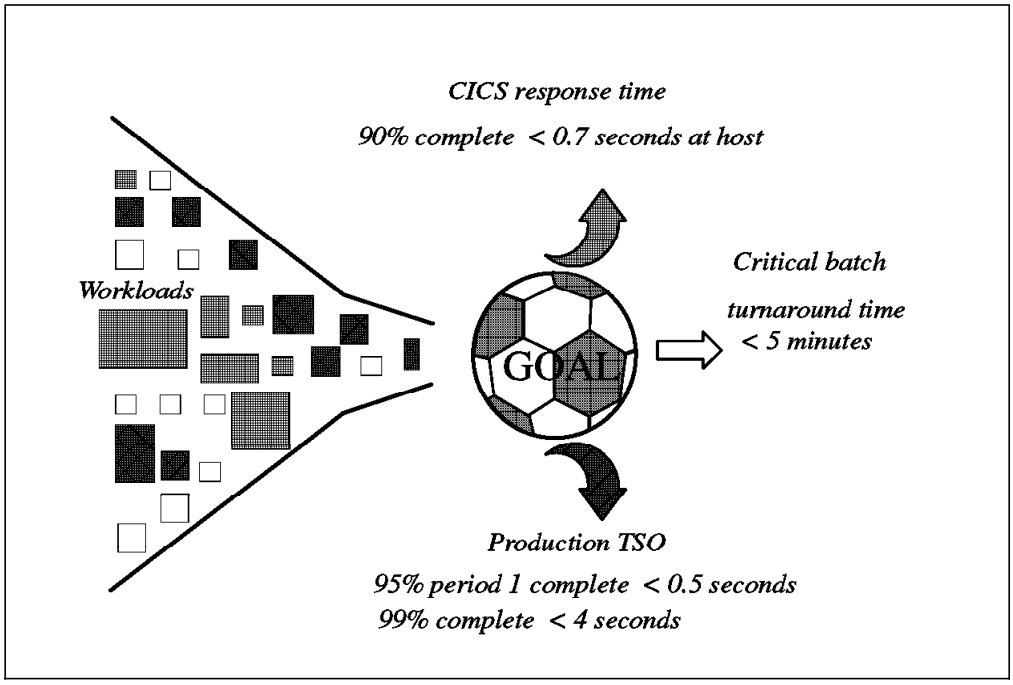

Figure 4. Assignment of Business Goals to Work

2.2.4.3 Importance

As well as goals, service classes are given an *importance* value. This is the importance of the goal being met, not the absolute importance of the workload.

Note: This should not in any way be associated with the dispatching priority. It is possible that an installation may place more importance on a particular class of batch work having a one hour turnaround than on CICS getting a half-second response time. CICS will more than likely have a higher dispatching priority and will probably get far more resource. However, if both the batch work and CICS are missing their goals, Workload Manager tries to help the batch work before it tries to help CICS if the batch service class has a higher importance.

Importance is specified as a value between 1 and 5, with 1 being the most important. Also, note that importance is a relative value, not an absolute one. Given two service classes missing their goals, their *relative* importance determines which one Workload Manager tries to help first.

2.2.4.4 Performance Index

When you define goals for a workload, you also need some feedback on whether the goals are being met or not. The achievement of each goal is measured by a ratio called the *performance index* (PI). The ratio is such that a value equal to 1 means that the service class is achieving its goal. A value less than 1 indicates the goal is being overachieved, while a value greater than 1 indicates that the service class is missing the goal. The definition of the ratio depends of the type of goal, as shown in Table 1:

Type of Goal	Performance Index Calculation
Average response time	Actual response time / Goal response time
Response time percentile	Actual percentage / Goal percentage
Velocity	Objective Velocity / Actual Velocity
Discretionary	Set to 0.81

Table 1. Performance Index Calculation

Two calculations are done. Both the sysplex-wide PI (global performance index) and individual system PIs (local performance index) are determined. Each service class period may have two types of PIs:

Sysplex PI A PI that represents its global performance across all the systems in the sysplex.

Local PI A PI that represents its performance in each local system of the sysplex. A service class may have more than one local PI if that service class has work executing on multiple systems in the sysplex.

Through services provided by SRM, performance data is periodically exchanged between sysplex systems in goal mode. The information is sent across the systems and allows SRM to construct an approximate view of the status of the sysplex through the calculation of sysplex PIs. The aggregate view enables SRM algorithms such as policy adjustment to make trade-offs within each individual system. This exchanged data includes:

- Response time information for each service class period, to allow each local SRM to calculate the sysplex PIs for each service class period.
- CPU service consumption information for each resource group, to enforce globally the capping and the protection algorithms.
- Information identifying the active service policy

RMF reports the actual response time, the velocity attainment, the performance index, and the goals.

Figure 5 on page 20 shows an example of the sort of information RMF provides at the service class period level.

```
REPORT BY: POLICY=CICSWORK   WORKLOAD=BATCH      SERVICE CLASS=OURBATHI   RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=3

TRANSACTIONS    TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ---SERVICE----   --SERVICE RATES--   PAGE-IN RATES    ----STORAGE----
AVG     23.80   ACTUAL           1.16.930   SSCHRT 719.5   IOC      1031K   ABSRPTN      8153   SINGLE     0.0   AVG     1480.08
MPL     23.80   EXECUTION        1.13.486   RESP    22.1   CPU     13388K   TRX SERV     8153   BLOCK      0.0   TOTAL   35230.9
ENDED      80   QUEUED              3.443   CONN     2.4   MSO     43427K   TCB         781.2   SHARED     0.0   CENTRAL 35230.9
END/SEC  0.27   R/S AFFINITY            0   DISC     3.2   SRB    394127    SRB          22.7   HSP        0.0   EXPAND     0.00
#SWAPS      1   INELIGIBLE             0   Q+PEND   1.5   TOT     58240K   RCT           0.1   HSP MISS   0.0
EXECUTD     0   CONVERSION           707   IOSQ    15.0   /SEC   194117    IIT          13.8   EXP SNGL   0.0   SHARED     0.00
                STD DEV          1.35.195                                   HST           0.0   EXP BLK    0.0
                                                                            APPL %      272.5   EXP SHR    0.0

VELOCITY MIGRATION:   I/O MGMT  29.8%    INIT MGMT 29.8%

        ---RESPONSE TIME---  EX   PERF   AVG   --USING%--  --------------- EXECUTION DELAYS % ----------------  ---DLY%--    %
        HH.MM.SS.TTT         VEL  INDX  ADRSP   CPU   I/O  TOTAL   I/O  CPU                                       UNKN IDLE QUIESCE
GOAL                        40.0%
ACTUALS
*ALL                        29.8%  1.3   24.2  12.5  16.6   68.7  50.7 18.0                                        2.2  0.0  0.0
SC66                        24.8%  1.6   19.9   8.0  16.4   73.9  56.0 18.0                                        1.7  0.0  0.0
SC67                        53.4%  0.7    4.3  33.5  17.4   44.5  26.3 18.1                                        4.6  0.0  0.0
```

Figure 5. Extract of RMF Service Class Period Report

2.2.4.5 Donor and Receiver Algorithm

In order to meet the goals and service rate group requirements, the policy adjustment routine determines:

Receiver Which service class periods demand more resources, that is, which have the highest sysplex PI and local PI above a certain less-than-one threshold. The SRM refers to the class period it is accessing helping as *receivers*.

The policy adjustment routine also determines resource bottlenecks (delays) in the receivers.

Donor The class periods that donate resources to the receivers are referred to as the *donors*. Donors donate these resources to receivers.

Periodically (several times per second) the following cycle is done.

1. Select all service class periods whose global performance index is greater than 1.
2. Order the service class periods by goal importance so the most important are on top of the list. Make a second list by performance index with the worse performers at the top.
3. Search in the list for the first service class that is delayed for a resource in order to obtain a resource donor.

   A donor is a significant user of a resource needed by a receiver and:

   - Is less important than the receiver or
   - Is more important than the receiver, but is over-achieving its goal

   If a receiver and a donor are found, the receiver's priority is increased and the donor's priority is decreased for the resource, and then the cycle ends.

   If SRM is unable to find a donor/receiver couple in the list, or the list is empty, the process is repeated in each member of the sysplex using local performance indexes.

Note: Before OS/390 Release 3, the only resources that could be donated and received were processor storage and CPU. With OS/390 Release 3, I/O is a resource for the donor/receiver algorithms.

The donor and receiver are rarely a single address space. WLM usually refers to a receiver service class period, not a receiver address space.

WLM also ensures that the reallocation of resources from the donors to the receivers does more good than harm. The term used is:

Net value A receiver is only helped by a specific donor if there is projected to be sufficient *net value* to the resource allocation. This net value assessment makes sure that when using a donor to help a receiver results in more projected benefit to the receiver than more projected harm to the donor.

2.3 Workload Manager with OS/390 Release 4

Once OS/390 Release 4 is installed, it is possible to use previous levels of:

- WLM ISPF table data sets
- WLM couple data sets
- WLM service definitions
- Activated WLM service policies

To take advantage of new features provided in OS/390 Release 4, the latest level of the WLM administrative application may be used to change current service definitions. The first time a change is made, the WLM ISPF tables are dynamically converted to the format required for OS/390 Release 4. This changes the functionality level of the service definition being modified to level 4.

Note: Once this occurs, the WLM administrative application from older OS/390 releases are no longer able to use that WLM service definition. There is a warning when the conversion takes place.

2.4 Batch Initiator Management

Workload manager now has the capability to dynamically manage the number of batch initiator address spaces serving specific classes of batch jobs that have been identified and submitted through JES. These jobs become assigned to one or more WLM service classes, allowing WLM to start new initiators on demand, to meet the performance goals of that category of work.

The system resources manager (SRM) has been changed to use OS/390 Release 3 address space management algorithms in the batch environment. See 2.13.1.1, “Server Address Space Management” on page 62. These enhancements rely on additional work queuing delay information which is provided by JES using new programming interfaces supplied by WLM.

JES uses the WLM-provided services during initialization to connect to WLM via the IWMCONN service. This initial contact with WLM during initialization occurs by specifying the following options:

WORK_MANAGER(YES) This parameter identifies the role that the connecting address space processes.

SUBSYS(JES) This parameter identifies the set of WLM classification rules to be used.

SUBSYSNM(name) This parameter, when combined with the SUBSYS-TYPE value, uniquely identifies the subsystem instance.

This information physically binds the JES address space to WLM, enabling use of WLM work classification services by JES and establishing WLM address space termination recovery.

Changes have also been made to the initiator code that allows WLM to intercept job scheduling requests. This makes it possible to dynamically reassign initiator address spaces to work queues that require additional throughput.

You can have as many job classes under WLM-managed mode as you choose. Once there, you have the flexibility, through operator command or SDSF, to dynamically switch any class back to JES-managed mode.

Prior to OS/390 Version 2 Release 4, job initiation delay was driven solely by by the number of initiators you had configured. With this new capability, you have the option of turning that responsibility over to WLM and the SRM. Under this new setup, initiators are created as they are needed, depending on:

- The presence of backlog work in WLM managed classes
- Adequate capacity to do more work
- The service class goals associated with jobs in the backlog

2.5 Service Classification

There is one set of classification rules in the service definition for a sysplex. They are the same regardless of what service policy is in effect; a policy cannot override classification rules. You should define classification rules after you have defined service classes, and ensure that every service class has a corresponding rule.

2.5.1 Default Service Classes

When work arrives in the system and there is no associated service class defined in the WLM service definition′s classification rules, a default service class is assigned based on the type of work. There are three default service classes:

SYSTEM This is used for the high dispatching priority (X′FF′) system address spaces. WLM recognizes these as special and treats them accordingly. If there are unique situations where you want to have more control over them, you can do the following:

- Define a service class for them
- Set up a classification rule in the STC subsystem type which will assign them, by transaction name, to a service class other than the default.

OS/390 Version 2 Release 4 enhancements permit you to assign work directly into the SYSTEM service class as part of your work classification rules. However, you can only do this within the STC subsystem type.

SYSSTC This is a special service class for started tasks. Work in this service class is assigned an importance level just below the special system address spaces, in terms of dispatching.

OS/390 Version 2 Release 4 enhancements permit you to assign work directly into the SYSSTC service class as part of your work

classification rules. You are able to do this within the following subsystem types:

- APPC
- JES
- OMVS
- STC
- TSO

SYSOTHER This service class acts as a catcher for all other work that falls through the classification rules.

Notes:

1. You are not able to control *MASTER* or INIT in a service class.
2. Once you choose to control an address space, you do so until it is restarted, which may require an IPL. This means your control over the address space continues even after a WLM mode switch or a service policy switch. Subsequent SET ICS,RESET or VARY WLM,POLICY= commands treat the address space as an ordinary started task.

2.5.2 Classification Rules

New support is provided in WLM to permit the use of the JES job priority within the classification rules. Priority references within classification rules must be within the numeric range of 0 to 15, which corresponds to the range of job scheduling priorities currently supported by both JES2 and JES3.

To minimize the number of classification rules that may be required, new logical operators are supported as follows:

> Specifies greater than

< Specifies less than

See Figure 8 on page 24 for use of these logical operators.

2.5.2.1 WLM Work Qualifiers

Figure 6 is the list of work qualifiers and their abbreviations. A new qualifier, job priority (PRI) is for qualifying batch work by JES.

AI	-	Accounting information	**PN**	-	Plan name
CI	-	Correlation information	**PR**	-	Procedure name
CN	-	Collection name	**PRI**	-	Job priority
CT	-	Connection type	**SI**	-	Subsystem instance
LU	-	Logical Unit name	**SPM**	-	Subsystem parameter
NET	-	Netid	**TC**	-	Transaction class/job
PF	-	Perform	**TN**	-	Transaction name/job
PK	-	Package name	**UI**	-	Userid

Figure 6. WLM Work Qualifiers

2.5.2.2 JES Work Qualifiers

WLM classification services are used by JES to assign a service class to each job after it is converted. This is regardless of the job class or eligibility for WLM-managed initiators. The service class is assigned based on WLM classification rules, which can include attributes as shown in Figure 7.

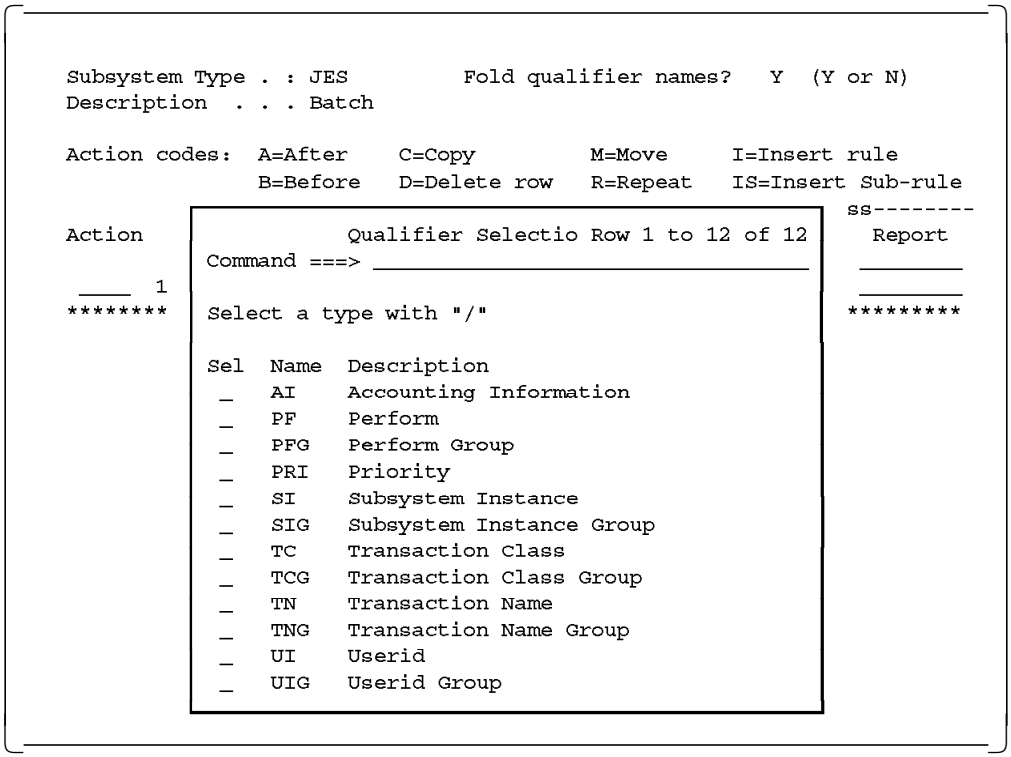

Figure 7. JES Work Qualifiers

Figure 8 contain the batch classification rules and for job class C shows that jobs with a priority greater than 9 are assigned a service class of BTCH_HI.

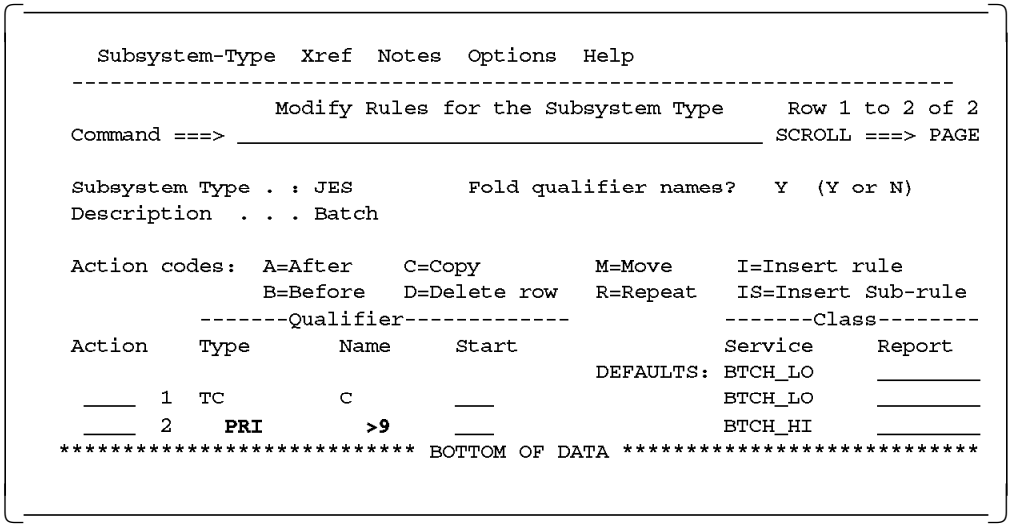

Figure 8. Batch Classification Rules

2.6 Batch Job Selection

Figure 9 illustrates the flow of a job from the reader through the execution phases.

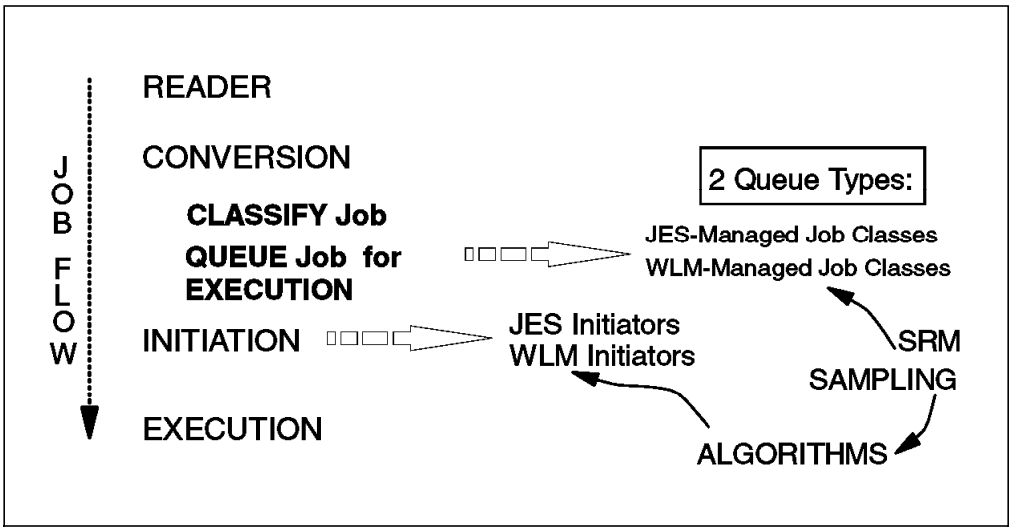

Figure 9. Batch Job Flow Phases

2.6.1 Classify Job

Following conversion, JES calls the WLM IWMCLSFY service to classify the job. Passed to WLM work classification are many of the job attributes extracted by the conversion phase, as shown in Figure 10.

```
IWMCLSFY USERID=$DWORK2,  Owning userid
      TRXCLASS=JQXJCLAS,  Job class
      TRXNAME=JQEJNAME,   Job name
      ACCTINFO=(R6),      Accounting string
      ACCTINFL=$DWORK+4,  Address of acctg length
      PERFORM=(R5),       Address of Perf group
      PRIORITY=$DWORK,    Priority
      CONNTKN=(R4),       WLM connect token
      SERVCLS=JQASTOK,    Service class token (rtn)
      SRVCLSNM=JQAWSCN,   Service class name  (rtn)
      MF=(E,$CLSFY)       Address of list form
```

Figure 10. JES2 Job Classification

On return from WLM classification, JES receives:

- A four byte classification token
- A service class name

Prior to OS/390 Release 4, the classification of jobs for WLM purposes was performed during job initiation. Job processing by JES prior to initiation was driven solely by JES job class associations and the related job priorities. This early classification enables introduction of goal-oriented logic within JES, driven by the WLM view of job goals and importance levels.

The classification token contains three key pieces of WLM information:

- An index to the WLM service class
- An index that identifies the report class name
- An indicator used by WLM to confirm that a policy activation has not occurred since the last classification was performed.

JES maintains the service class name and service class token until the completion of job execution. This information is passed to the initiator during job select processing.

2.6.2 Queue Job for Execution

Batch jobs are selected for execution by initiators, each one running in a separate address space. Initiators are controlled by JES or by WLM. Which one controls the initiators is determined by the job class and by the MODE= parameter (JES or WLM) on the JES2 JOBCLASS statement. JES2 now maintains two different queue organizations for all jobs awaiting execution, as shown in Figure 11:

1. All jobs are queued by job class, priority, and the order in which they finished conversion. This is the queue from which JES2-managed initiators select jobs for execution.
2. Jobs awaiting execution in WLM managed job classes are also queued by their WLM assigned service class in the order they were made available for execution.

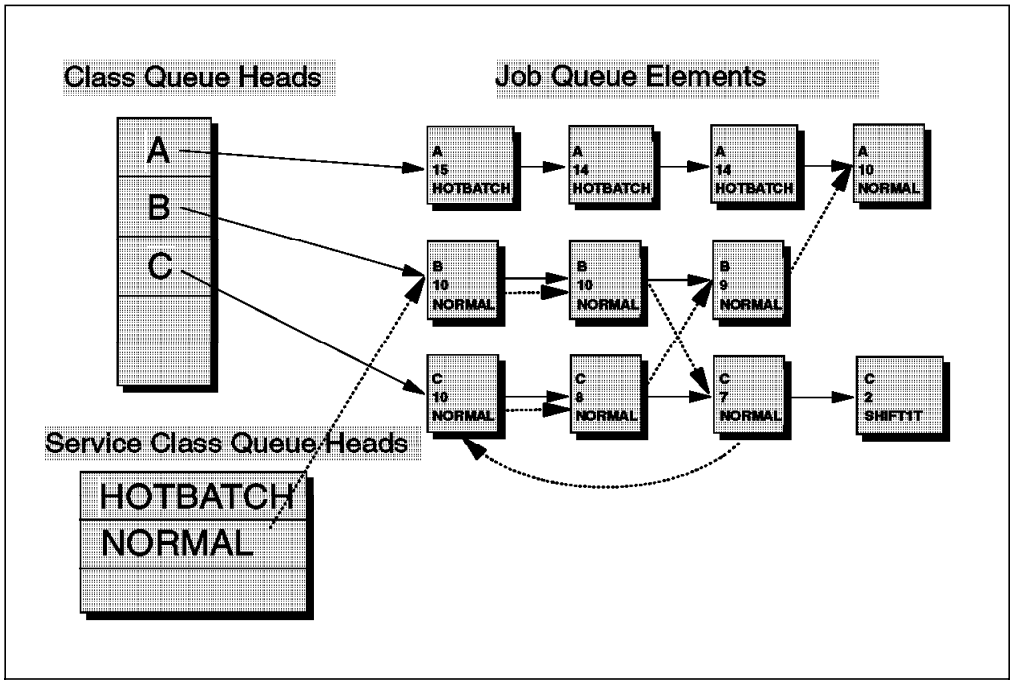

Figure 11. JES2 and Service Class Queues

2.6.3 WLM Goal Mode and Batch

For the batch work to be managed automatically by WLM on a system, that system must be in *goal* mode.

WLM impacts both the JES products and the way in which jobs are queued for execution scheduling, and it provides new operational capabilities, as follows:

- The SRM is changed to enable application of address space management algorithms to the batch environment, relying upon additional work queueing delay information provided by JES through new programming interfaces supplied by the WLM component.
- The initiator component is changed to permit WLM interception of job scheduling requests, enabling dynamic reassignment of initiator address spaces to work queues requiring additional throughput.

WLM goal types need special considerations for batch jobs:

Response Time A queue delay is included in the response times of jobs in all classes, therefore response time increases as the job queues increase. This may be appropriate for jobs that have specific turnaround times.

Velocity A queue delay is also included in the velocity of jobs in WLM-managed classes. As the number of jobs waiting for WLM initiators grows, the velocity of jobs in the service class decreases.

Discretionary These may be appropriate for “overnight batch” jobs, but only if it is acceptable when jobs are not run because of insufficient resources.

If your installation has large job execution queues, then the queue delay can dominate the response time or velocity. This may have the effect of the SRM not addressing other delays because to do so would significantly affect the performance index.

2.6.4 Preparing for Goal Mode Batch Management

The introduction of initiator batch management allows for job classes to be initiated by JES and WLM. In preparation for this support, the following steps must be considered:

- The WLM couple data set must be reformatted for Release 4. For details, see 2.11, “Migration Considerations” on page 53.
- For classes to be managed by WLM, the job classes must be defined to JES as MODE=WLM. The JES2 JOBCLASS initialization statement has a new optional keyword, MODE. For details, see 3.3.1, “JOBCLASS Statement Changes” on page 78.

  Values for this keyword are either JES (the default) or WLM. By specifying MODE=JES or omitting this keyword altogether, you keep the jobclass in JES-managed mode. With this setting, JES2 manages the batch initiators for that job class, in exactly the same way as it has in prior releases. By specifying MODE=WLM for selected classes, you place those classes into WLM-managed mode. See 2.6.6, “Using WLM Managed Initiators” on page 29.
- Service classes for the WLM-managed batch must be defined in a WLM policy.

  This policy must be activated, as shown in Figure 14 on page 31.

2.6.5 Defining Batch Service Class Goals

For batch work, you can define response time goals or velocity goals or discretionary goals. For a service class with multiple periods, you cannot switch from a velocity goal in one period to a response time goal in a later period.

Execution velocity goals define how fast work should run when ready, without being delayed for processor, storage, I/O access, and queue delay. Execution velocity goals are intended for work for which response time goals are not appropriate, such as started tasks or long-running batch work.

Discretionary goals are for low priority work for which you do not have any particular performance goal. Workload management then processes the work using resources not required to meet the goals of other service classes.

2.6.5.1 Batch Velocity Goals

Velocity goals define the acceptable amount of delay for work when work is ready to run. Velocity goals are intended for subsystems which use address spaces or enclaves to represent individual work requests or for long-running batch jobs.

The formula for velocity is:

$$\frac{\text{using samples}}{\text{using samples} + \text{delay sample} * 100}$$

Where:

using samples If you are running with I/O priority management on, using samples is:

- The number of samples of work using the processor.
- The number of samples of work using non-paging DASD I/O resources in a state of device connect or device disconnect.

Note: The I/O samples are derived from actual time measurements.

If you are running with I/O priority management off, using samples is:

- Only the number of samples of work using the processor.

delay samples If I/O priority management is on, delay samples is:

- The number of samples of work delayed for the processor.
- The number of samples of work delayed for storage which includes paging, swapping, multiprogramming level (MPL), server address space creation delays, and initiation delays.
- The number of samples of work delayed for non-paging DASD I/O which includes IOS queue, subchannel pending, and control unit queue delays.

Note: The samples for subchannel pending and control unit queue delay are derived from actual time measurements.

If you are running with I/O priority management off, delay samples is:

- The number of samples of work delayed for the processor.
- The number of samples of work delayed for storage.

Note: MPL is the SRM-controlled function that adjusts the number of address spaces allowed to be in central storage and ready to be dispatched.

Importance levels should stay the same or decrease as the transactions move from one performance period to the next. Remember that importance applies only if a goal is not being met during the duration of the period.

```
Service Class = BATCHX
Period 1
    Velocity      = 40
    Importance    = 4
    Duration      = 3000 SU
Period 2
    Velocity      = 30
    Importance    = 4
```

Initiation delays cause the velocity value to decrease. You should adjust your velocity goals accordingly. To estimate a velocity goal, if all the batch jobs in the service class are WLM-managed job classes:

- Run the job class
- Examine the initiation delay data
  - RCAETOTDQ and RCAETOTU in IWMWRCAA
  - R723CTDQ and R723CTSA in SMF 72, Subtype 3
- Include these delays in the velocity formula for an estimate of the new, lower velocity. Plugging this delay data into the velocity formula gives you:

  ```
  <<RCAETOTU> over <RCAETOTU + RCEATOTDQ>> * 100
  ```

  RMF will do this calculation for you. Look for the INIT MGT field in the RMF Monitor I workload activity report.

2.6.5.2 Batch Response Time Goals

If you are setting batch response time goals for goal mode, consider the following:

- The number of batch transactions equals the number of jobs.
- Defining a response time goal may not be appropriate for some types of batch work, such as jobs with very long execution times.
- Jobs that are appropriate for a response time goal should have at least three transaction completions per 20 minutes of elapsed time. If there are too few completions, use a velocity goal.

2.6.6 Using WLM Managed Initiators

The steps that must be in place before you are able to make use of WLM-managed initiators are covered in 2.11, “Migration Considerations” on page 53. Simply by coding a WLM controlled class on the jobcard places the initiation of the job under the control of WLM. Initiators are started as required. Figure 12 on page 30 shows a SYSLOG when a job is initiated by WLM:

```
IWM034I PROCEDURE INIT STARTED FOR SUBSYSTEM JES2
APPLICATION ENVIRONMENT SYSBATCH
PARAMETERS SUB=MSTR
IEF196I         1 //INIT     JOB MSGLEVEL=1
IEF196I         2 //STARTING EXEC INIT
IEF196I
IEF196I ************************comments deleted**********************
IEF196I
IEF196I  STMT NO. MESSAGE
IEF196I         2 IEFC001I PROCEDURE INIT WAS EXPANDED USING SYSTEM
IEF196I LIBRARY SYS1.PROCLIB
IEF196I         3 XXIEFPROC   EXEC   PGM=IEFIIC,DPRTY=12
ICH70001I RONN     LAST ACCESS AT 08:13:13 ON WEDNESDAY, JUNE 4, 1997

$HASP373 RONNWMIB STARTED - WLM INIT  - SRVCLASS STC_HI   - SYS SC63
PROGRAM : IEFBR14  IN STEP : STP1     - COND CODE : 0000
TOTAL CPU (SECS)  BLKS READ/WRITE      TOTAL SWAPS
            0.01
$HASP395 RONNWMIB ENDED
```

Figure 12. SYSLOG of a Job Started by a WLM Initiator

The highlighted sections are of primary interest:

- Message IWM034I is issued at the start of the process.
- WLM-initiated jobs are identified by WLM INIT as the job starts.

As shown in Figure 14, the $D JOBCLASS command has a new field, MODE, which makes it easy to determine whether the class is JES-or WLM-managed. Also, see 4.5.1, “Invoking the Job Class Panel” on page 106 for the SDSF display of job classes.

```
$DJOBCLASS(W)

$HASP837 JOBCLASS(W)         ACCT=NO,PGMRNAME=NO,
$HASP837                     TIME=(000450,00),
$HASP837                     REGION=0002M,
$HASP837                     COMMAND=DISPLAY,BLP=YES,
$HASP837                     AUTH=(ALL),MSGLEVEL=(1,1),
$HASP837                     COPY=NO,HOLD=NO,IEFUJP=YES,
$HASP837                     IEFUSO=YES,JOURNAL=YES,
$HASP837                     LOG=YES,MODE=WLM,
$HASP837                     OUTDISP=(,),OUTPUT=YES,
$HASP837                     PERFORM=000,PROCLIB=00,
$HASP837                     QHELD=NO,RESTART=YES,
$HASP837                     SCAN=NO,SWA=ABOVE,
$HASP837                     TYPE26=YES,TYPE6=YES,XBM=,
$HASP837                     XEQCOUNT=(MAXIMUM=*,
$HASP837                     CURRENT=0)
```

Figure 13. Display WLM-Managed Job Class

Note: When you first try this you may find that the job does not initiate even though it is in the correct class. Remember that activating a WLM service policy does not place a system in goal mode; to achieve this, you must also issue the F WLM,MODE=GOAL command on each participating system.

2.6.7 Activating a WLM Policy

Figure 15 illustrates how to 1 activate a WLM service policy 2 place each participating system into goal mode and 3 verify the results.

```
V WLM,POLICY=RONNP1                    1

IWM001I WORKLOAD MANAGEMENT POLICY RONNP1 NOW IN EFFECT


F WLM,MODE=GOAL                        2

IWM007I SYSTEM SC63 NOW IN WORKLOAD MANAGEMENT GOAL MODE

ROUTE SC64 F WLM,MODE=GOAL

IWM007I SYSTEM SC64 NOW IN WORKLOAD MANAGEMENT GOAL MODE


D WLM,SYSTEMS                          3

IWM025I  18.40.34  WLM DISPLAY 846
  ACTIVE WORKLOAD MANAGEMENT SERVICE POLICY NAME: RONNP1
  ACTIVATED: 1997/06/03  AT: 18:40:12  BY: RONN      FROM: SC63
  DESCRIPTION: Test OS/390 R4 WLM Enhancements
  RELATED SERVICE DEFINITION NAME: WLMR4_SD
  INSTALLED: 1997/06/03  AT: 18:37:08  BY: RONN      FROM: SC63
  WLM VERSION LEVEL: LEVEL004
  *SYSNAME*  *MODE*  *POLICY*  *WORKLOAD MANAGEMENT STATUS*
  SC63       GOAL    RONNP1    ACTIVE
  SC64       GOAL    RONNP1    ACTIVE
```

Figure 14. Operator Commands to Activate Goal Mode Policy

Note: Be aware that a system is returned to compatibility mode if the IEAIPSxx and IEAICSxx members are referenced in parmlib. To avoid this, remove the IPS= and ICS= pointers from IEASYSxx at this stage.

2.6.8 Controlling Batch Job Selection

Although the Workload Manager dynamically controls its initiators, there are several methods you can use to manage them with JES2 commands.

2.6.8.1 Job Class Limits

You can limit the number of jobs in each class that can execute simultaneously in the MAS by using the XEQCOUNT=MAXIMUM= parameter on the JOBCLASS statement. This value can be altered with the $T JOBCLASS command.

Note: There is no facility to control the number of jobs in execution by class on an individual system.

2.6.8.2 Stop or Start Job Selection

You can stop or start the selection of batch jobs by JES-or WLM- managed initiators on an individual system with the $P XEQ or the $S XEQ command.

The scope of the JES2 $PI command is limited to the control of JES2-managed initiators. New JES2 commands are introduced to control JES2 job scheduling on a system basis. The $P XEQ command halts further scheduling on the system where the command is executed, and causes the related JES2 and WLM

initiators to drain and terminate. Scheduling can then be resumed by using the $S XEQ command.

Note: The initial setting when JES2 first initializes is to allow selection. If you want to prevent this, place a $P XEQ command in your JES2 initialization deck or better still, have your automation changed to do it.

2.6.8.3 Job Selection Using Affinity

You can control where jobs execute by using member affinity or scheduling environments. Member affinity is specified by using the /*JOBPARM user control card and specifying SYSAFF=system-name.

For scheduling environments, see 2.7, “Resource Affinity Scheduling” on page 35.

2.6.8.4 Forcing Job Selection

To immediately start executing a batch job that is in a WLM-managed job class, use the $S Job command. The Workload Manager selects the job for immediate execution on an eligible system, based on the SYSAFF and SCHENV parameters.

Note: The $S JOB command causes the job to be started even if the job is in hold.

2.6.8.5 Job Selection Using User Exits

With the use of JES2 user exits, you can change the classification of jobs entering the system by enforcing installation controls. This can be done by using the following exits:

Exit 2 Job statement scan

Exit 6 Converter/Interpreter text

Exit 20 End-Of-Job-Input processing

Exit 49 Job Queue Work Select

2.6.9 Pre-Execution Job Delays

With OS/390 Release 4, support is provided for controlling how and when batch jobs are scheduled. For a job that passes through the conversion phase of processing, that job′s response time or pre-execution delays begins. Any scheduling delays depend jointly upon what the end user specifies in JCL and upon installation actions by system operators and automated operations packages. These new delays are included in job response times for reporting purposes. These job delays are included in RMF workload activity reports. These job delays are shown in Figure 15 on page 33.

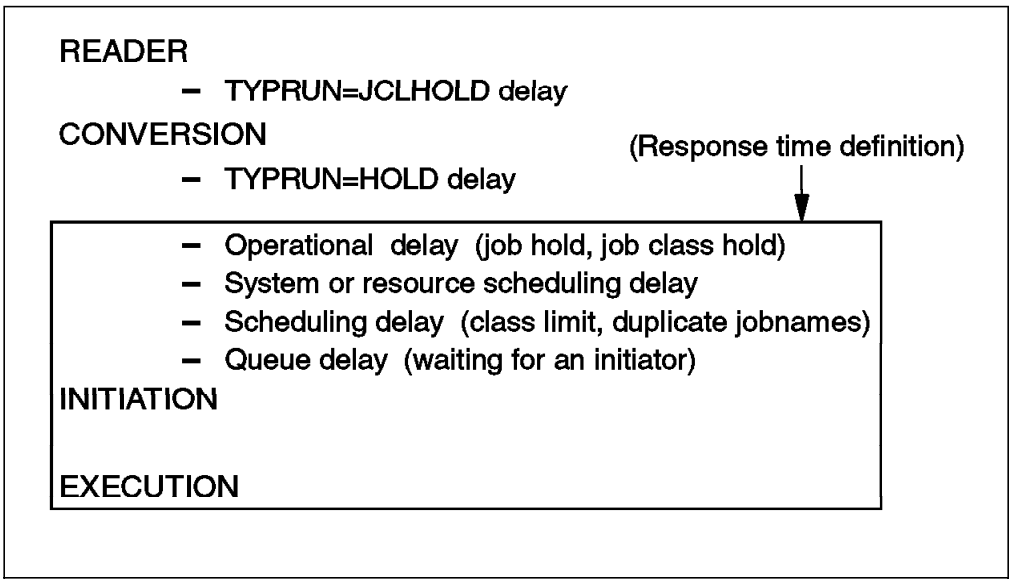

Figure 15. Pre-Execution Job Delays

Pre-execution job delay times does not include TYPRUN=HOLD or JCLHOLD delays. It does include the following:

- Operational delays (jobs or job classes held by operator command)
- System or resource shortage delays
- Scheduling delays due to class limits and duplicate jobnames
- Waiting for an initiator

Figure 16 on page 34 represents the different stages in the life of a job (output and purge processing have been omitted because they are not pertinent to the discussion).

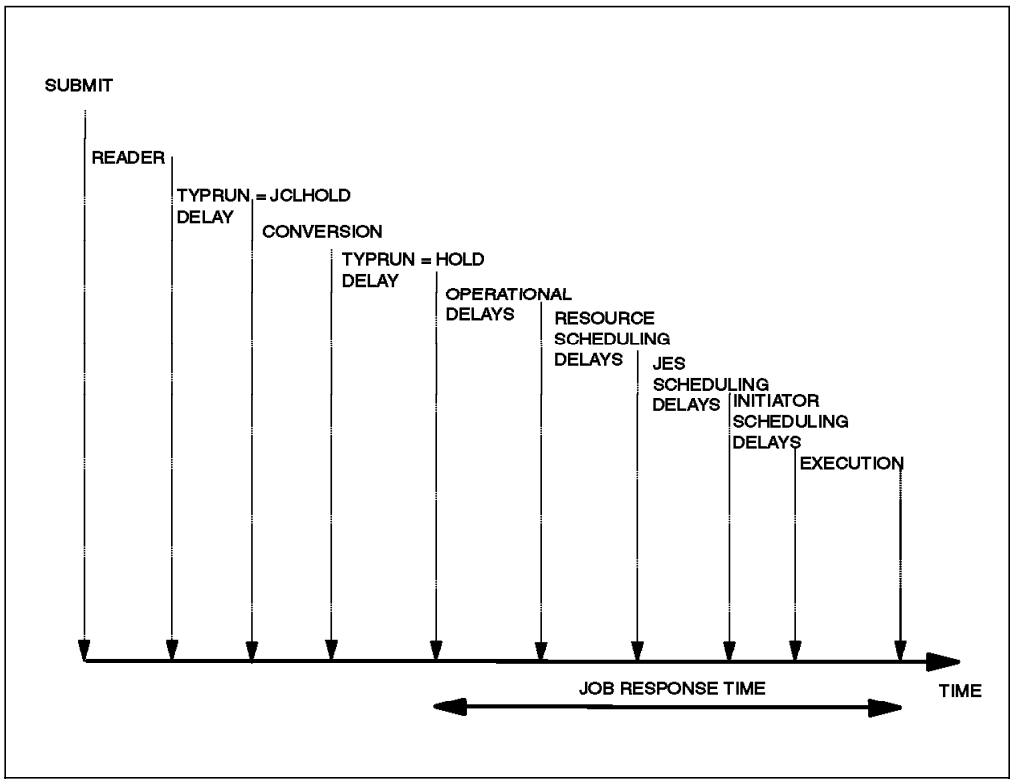

Figure 16. Life Cycle of a Job

Following is a description of the job phases and how they relate to scheduling delays:

Reader
The JCL is scanned by JES. Job information is stored on the spool and the job nows exists in the system.

TYPRUN=JCLHOLD delay
If this parameter is specified, the job waits for conversion until it is released. This delay is not monitored (or it becomes part of the job response time).

Conversion
The JCL is partially validated and transformed into internal formats required for job initiation.

TYPRUN=HOLD delay
If this parameter is specified, the job remains in a class queue until it is released. This delay is not monitored (or it becomes part of the job response time).

Operational delays
This phase includes a set of delays induced by operator actions:

- An individual job may be held by the operator.
- A job class is held.
- The entire job queue is held.
- Jobs having data on a given JES spool volume are held.

The installation causes these delays. They are therefore accounted for in the job response time.

Resource scheduling delays	These delays are produced by the unavailability of a scheduling environment.
JES scheduling delays	These delays are due to the way JES works: • Duplicate jobname serialization • Job class execution limits (which is new with this release).
Initiation delays	These are caused by waiting for an initiator either in JES mode or WLM mode.
Job execution time	The job is now executing and the job response time continues until the job completes execution.

2.7 Resource Affinity Scheduling

The initial announcement of Workload Manager did not have as one of its primary functions to dynamically balance batch work across multiple systems in a sysplex. OS/390 Release 4 begins to provide workload balancing, but only with JES2 batch work. New function has been added to WLM which makes it possible to define and control abstract resource names, mapped against real system resources (for example, systems in a particular DB2 datasharing group, or systems dedicated to online production workloads or those with cryptographic capabilities). JES2 provides scheduling mechanisms that only initiate jobs on systems where the required resources are available.

Resource scheduling lets an end user declare which resources are required for a job, and have the job scheduled only on systems having the required resources. Using resource scheduling, you can control:

- Where batch work is initiated (where a job's resources exist, and also where you do *not* want a job to execute)
- Where batch work is initiated (when jobs are run)
- A combination of these items

This eliminates the need for an end user to:

- Specify a specific system by name
- Specify a job class for job initiation

This allows, primarily for batch work, to have scheduling and controls that ensure that jobs are scheduled only to the systems having access to the required resources, such as:

- Devices
- Hardware features such as encryption
- Specific products
- Specific subsystems

2.7.1 Scheduling Environments

A *scheduling environment* is a list of resource requirements that you define using the WLM administration application. When an MVS system meets all of the requirements in the scheduling environment requested by the batch job, then that system becomes eligible to process the job. If all of the requirements are not met, then the job cannot be processed on that system.

2.7.2 Resource Names

Resource names are really indicators on an MVS system that show the presence or absence of a real resource. The resource can be an actual physical entity, such as a data base or a hardware feature (vector processor or cryptography). It can also be an intangible quality, such as a certain time of day or a certain day of the week. However, understand that the resource names are abstract and have no actual meaning or direct relationship to physical resources.

Resource names have no real meaning and may have no direct relationship to any physical resources in existence. You can establish a resource name, such as *DB2A*, and then schedule jobs where a given instance of a DB2 subsystem either exists or does not exist. However, the name DB2A could equally be XYZ and the concept would be the same and work equally well.

A resource name could be created that establishes time periods, such as *OVERNITE*, which would identify jobs not eligible to run during the period prime shift.

A resource name could be associated with a logical scheduling control, such as *BATCH*, where the resource would identify systems within a sysplex eligible for processing the batch workload.

Note: Resource names are installation- dependent.

For details on resource name definitions, see 2.8.1, “Defining Resources for Scheduling Environment” on page 40.

2.7.3 Resource States

Resource states are used by workload management in conjunction with scheduling environments to ensure that work is scheduled only on a system with the appropriate resources to process that work. See *OS/390 MVS Planning: Workload Management* for more information on resources and scheduling environments.

The definition of a job resource requires that the resource be one of two states; ON or OFF. Resource states are defined as follows:

ON If the required state is *ON*, then the resource state must be set to ON in an MVS image for the resource requirement to be satisfied. This is classified as “positive affinity.” A positive affinity means that the job can only run on a system where that resource name has been explicitly turned ON.

OFF If the required state is *OFF*, then the resource state must be set to OFF in an MVS image for the resource requirement to be satisfied. This is classified as “negative affinity.” A negative affinity means that the job can only run on a system where that resource name has been explicitly turned OFF.

RESET This state is the initial setting for all resource states after an IPL.

For example, you could define a resource name to be ZYXW with a required state of ON. If on system SYS1 the corresponding ZYXW resource state is set to ON, then the resource requirement is satisfied. WLM does not care what ZYXW means or whether the ON setting really does signify the existence of some real resource. You could use ZYXW as nothing more than a toggle switch, setting it

ON for whatever reason you choose. As long as the settings match, the requirement is satisfied.

For details on resource name and state definitions, see 2.8.1, “Defining Resources for Scheduling Environment” on page 40.

2.8 Defining Scheduling Environments

The following figures describe how to define scheduling environments and resource requirements, how to set the resource states on each individual MVS system, and how to associate a scheduling environment name with a batch job.

The following steps are necessary before using scheduling environments:

- Define one or more scheduling environments, and all of the resource requirements that are listed in those scheduling environments, in the workload management service definition.

  Each resource requirement consists of the name you choose for the resource and a required state of either *ON* or *OFF*.

- On every system in the sysplex that processes batch jobs according to resource requirements, the individual resources must be set to one of the following states:
  - ON
  - OFF
- Each batch job that is submitted for execution must specify the name of the scheduling environment to be used when determining which systems the batch job can be scheduled on.

A new environment, called a scheduling environment, is created by WLM to manage newly defined resources, as shown in Figure 17.

```
------------------------------------------------------------------------
Functionality LEVEL004         Definition Menu         WLM Appl LEVEL004
Command ===> ____________________________________________________________

Definition data set  . . : 'RONN.SANDBOX'

Definition name  . . . . . WLMR4_SD  (Required)
Description  . . . . . . . WLM Definitions - OS/390 R4

Select one of the
following options. . . . . ___  1.  Policies
                               2.  Workloads
                               3.  Resource Groups
                               4.  Service Classes
                               5.  Classification Groups
                               6.  Classification Rules
                               7.  Report Classes
                               8.  Service Coefficients/Options
                               9.  Application Environments
                              10.  Scheduling Environments
```

Figure 17. WLM Defintion Menu With New Scheduling Environment

This support introduces:

- WLM administrative definitions to define the resource names that can be requested by jobs. For details, see 2.7.1, “Scheduling Environments” on page 35.
- A new job card JCL external, SCHENV=, for the end user to define a resource affinity for a job. For details, see 2.9, “Using Job Scheduling Environments” on page 45.
- New operator commands for control over which systems have access to the resources.
- Support in SDSF to display resources requested by jobs and which jobs are being delayed by resources.
- Job scheduling support to direct the execution of jobs to systems having the necessary resources available.

This new capability is fully functional in both WLM goal mode and in WLM compatibility mode.

The task of defining scheduling environments and resource requirements is done using the WLM administrative application. There is a new option (10) on the Definition Menu called the primary WLM panel, as shown in Figure 17 on page 37. If no previous scheduling environments exist, selecting this option takes you into the panel shown in Figure 18.

```
                          Decide what to create?
Command ===> ____________________________________________________________


No scheduling environments exist.  Would you like to create a
scheduling environment or define resources for scheduling
environments?

Select one of the following options.
 1  1.  Create Scheduling Environment
    2.  Create Resource(s)
```

Figure 18. Define First Scheduling Environment

Selecting Option 1 allows you to create a scheduling environment when you enter the data shown in Figure 19 on page 39. Within this panel, enter the following information:

Scheduling Environment Name This field is required and is the 1- to 16-character name of the scheduling environment. You can have up to 999 unique scheduling environments defined in a sysplex.

Description This field is optional, and you can use up to 32 characters to describe the scheduling environment name.

```
  Scheduling-Environments  Notes  Options  Help
------------------------------------------------------------------------------
                     Create a Scheduling Environment          Row 1 to
Command ===> ____________________________________________________________

Scheduling Environment Name    BATCH_DFLT______  Required
Description  . . . . . . . . . Default Batch_____________________

Action Codes: A=Add  D=Delete

                           Required
Action  Resource Name      State      Resource Description

__                         _______
****************************** Bottom of data ************************
```

Figure 19. Creating a Scheduling Environment

When you press Enter and then press PF3, the scheduling environment is created.

Figure 20 shows the entry just created and now shows the creation of many scheduling environments. This panel is displayed from Option 10 of the Definition Menu once scheduling environments have been created.

```
  Scheduling-Environments  Notes  Options  Resources  Help
------------------------------------------------------------------------------
                 Scheduling Environment Selection List        Row 1 to 7 of 7
Command ===> ____________________________________________________________

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name  Description
  1     BATCH_DFLT                   Default Batch
  __    CICS_DEV                     Development CICS active
  __    NO_CICS_DEV                  Development CICS inactive
  __    NO_CICS_PROD                 Production CICS inactive
  __    OVERNIGHT                    Over Night (1800 - 0700)
  __    PRIMESHIFT                   Prime Shift (0800 - 1700)
  __    SANDBOX                      Test Sysplex: SC63 and SC64
****************************** Bottom of data ********************************
```

Figure 20. Defined Scheduling Environments

The Action Codes listed in Figure 20 show the functions that can be performed from this panel. Keying in a **1** (create) and pressing Enter causes the Create a Scheduling Environment panel to be presented, as shown in Figure 21 on page 40. Creates can be done from this panel without going back to the primary Definition Menu.

In Figure 21 on page 40, a scheduling environment name and description are entered.

```
  Scheduling-Environments  Notes  Options  Help
------------------------------------------------------------------------------
                    Create a Scheduling Environment          Row 1 to 1 of 1
Command ===> ____________________________________________________________

Scheduling Environment Name     PROD_CICS_ATMs __  Required
Description  . . . . . . . . .  Automated Teller Machines _______

Action Codes: A=Add  D=Delete

                          Required
Action  Resource Name     State      Resource Description

__                        _______
****************************** Bottom of data ********************************
```

Figure 21. Create a Scheduling Environment

Press PF3 and you are returned to the previous panel, now updated with the newly defined scheduling environment. The newly created scheduling environment is shown as the highlighted fields in Figure 22.

```
  Scheduling-Environments  Notes  Options  Resources  Help
------------------------------------------------------------------------------
                    Scheduling Environment Selection List      Row 1 to 10 of 10
Command ===> ____________________________________________________________

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name  Description
  __    BATCH_DFLT                   Default Batch
  __    BATCH_PROD                   Production Batch
  __    CICS_DEV                     Development CICS active
  __    CICS_PROD                    Production CICS active
  __    NO_CICS_DEV                  Development CICS inactive
  __    NO_CICS_PROD                 Production CICS inactive
  __    OVERNIGHT                    Over Night (1800 - 0700)
  __    PRIMESHIFT                   Prime Shift (0800 - 1700)
  3     PROD_CICS_ATMS               Automated Teller Machines
  __    SANDBOX                      Test Sysplex: SC63 and SC64
****************************** Bottom of data ********************************


        Scheduling environment PROD_CICS_ATMS was created. (IWMAM654)
```

Figure 22. Scheduling Environment Created

2.8.1 Defining Resources for Scheduling Environment

The next activity is to define the resource requirements that are to be assigned to the new scheduling environment. To start the process, enter **3** (modify) next to PROD_CICS_ATMS, as shown in Figure 22, and tab the cursor to the Resources field at the top of the panel and press Enter.

A popup window appears with the option to Process Resources, as shown in Figure 23 on page 41. Select this option by keying a 1 and pressing Enter.

```
  Scheduling-Environments  Notes  Options  Resources  Help
-------------------------------------- ------------------------- -----
                  Scheduling Environment  1  1. Process Resources   10 of 10
Command ===> ____________________________                          _____

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name  Description
  __    BATCH_DFLT                   Default Batch
  __    BATCH_PROD                   Production Batch
  __    CICS_DEV                     Development CICS active
  __    CICS_PROD                    Production CICS active
  __    NO_CICS_DEV                  Development CICS inactive
  __    NO_CICS_PROD                 Production CICS inactive
  __    OVERNIGHT                    Over Night (1800 - 0700)
  __    PRIMESHIFT                   Prime Shift (0800 - 1700)
  3     PROD_CICS_ATMS               Automated Teller Machines
  __    SANDBOX                      Test Sysplex: SC63 and SC64
****************************** Bottom of data ********************************
```

Figure 23. Create a Resource Definition List

The Resource Definition List panel is presented, as shown in Figure 24. To create the resources, type an A in the Action field, enter a resource description, and press Enter.

```
  Resources  Notes  Options  XREF  Help
------------------------------------------------------------------------------
                        Resource Definition List            Row 1 to 12 of 12
Command ===> ____________________________________________________________

Action Codes: A=Add  D=Delete  X=XREF  /=Menu Bar

Action  Resource Name     In Use  Resource Description
  A     BATCH_DFLT        YES     Default Batch
  __    BATCH_PROD        YES     Production Batch
  __    CICSD01           YES     Development CICS 01
  __    CICSD02           YES     Development CICS 02
  __    CICSD03           YES     Development CICS 03
  __    CICSP01           YES     Production CICS 01
  __    CICSP02           YES     Production CICS 02
  __    CICSP03           YES     Production CICS 03
  __    OVERNIGHT         YES     Over Night (1800 - 0700)
  __    PRIMESHIFT        YES     Prime Shift (0800 - 1700)
  __    SC63              YES     LPAR 1/2 Within Sysplex SANDBOX
  __    SC64              YES     LPAR 2/2 Within Sysplex SANDBOX
****************************** Bottom of data ********************************
```

Figure 24. Define Resource

The Define Resource panel is displayed, as shown in Figure 25 on page 42, and here you enter the resource name and description:

Resource Name This field is required, and is the 1- to 16-character name of the resource.

Resource Description This field is optional, and you can use up to 32 characters to describe each resource.

```
                          Define Resource
Command ===> ____________________________________________________________

Resource name  . . . . . . . . . Crypto_____  Required
Resource Description . . . . . Integrated Cryptographic Feature
```

Figure 25. Create Resource Name

When you press Enter and then press PF3, the panel shown in Figure 26 is displayed.

```
  Resources  Notes  Options  XREF  Help
------------------------------------------------------------------------------
                         Resource Definition List            Row 1 to 13 of 13
Command ===> ____________________________________________________________

Action Codes: A=Add  D=Delete  X=XREF  /=Menu Bar

Action  Resource Name     In Use  Resource Description
  __    BATCH_DFLT        YES     Default Batch
  __    BATCH_PROD        YES     Production Batch
  __    CICSD01           YES     Development CICS 01
  __    CICSD02           YES     Development CICS 02
  __    CICSD03           YES     Development CICS 03
  __    CICSP01           YES     Production CICS 01
  __    CICSP02           YES     Production CICS 02
  __    CICSP03           YES     Production CICS 03
  __    CRYPTO                    Crypto resource
  __    OVERNIGHT         YES     Over Night (1800 - 0700)
  __    PRIMESHIFT        YES     Prime Shift (0800 - 1700)
  __    SC63              YES     LPAR 1/2 Within Sysplex SANDBOX
  __    SC64              YES     LPAR 2/2 Within Sysplex SANDBOX
****************************** Bottom of data ********************************


                  Resource CRYPTO was created. (IWMAM665)
```

Figure 26. Resource Name Created

After you are returned to the Resource Definition List panel which shows the newly defined resource, CRYPTO, press PF3 and you are returned to the Modify A Scheduling Environment panel shown in Figure 27 on page 43.

```
  Scheduling-Environments  Notes  Options  Help
------------------------------------------------------------------------------
                    Modify A Scheduling Environment           Row 1 to 1 of 1
Command ===> ____________________________________________________________

Scheduling Environment Name  : PROD_CICS_ATMS
Description  . . . . . . . . . Automated Teller Machines

Action Codes: A=Add  D=Delete

                          Required
Action  Resource Name     State      Resource Description
  A_                      ________
****************************** Bottom of data ********************************
```

Resource changes have occurred. (IWMAM664)

Figure 27. Panel to Create Resource State

Type in an A (add) and press Enter. The Resource Definition lis is displayed, as shown in Figure 28.

```
  Resources  Notes  Options  XREF  Help
------------------------------------------------------------------------------
                          Resource Definition List          Row 1 to 14 of 14
Command ===> ____________________________________________________________

Selection For Scheduling Environment PROD_CICS_ATMS

Action Codes: A=Add  S=Select  X=XREF  /=Menu Bar

Action  Resource Name     In Use  Resource Description
  __    BATCH_DFLT        YES     Default Batch
  __    BATCH_PROD        YES     Production Batch
  __    CICSD01           YES     Development CICS 01
  __    CICSD02           YES     Development CICS 02
  __    CICSD03           YES     Development CICS 03
  __    CICSP01           YES     Production CICS 01
  __    CICSP02           YES     Production CICS 02
  __    CICSP03           YES     Production CICS 03
  S     CRYPTO                    Crypto resource
  __    OVERNIGHT         YES     Over Night (1800 - 0700)
  __    PRIMESHIFT        YES     Prime Shift (0800 - 1700)
  __    SC63              YES     LPAR 1/2 Within Sysplex SANDBOX
  __    SC64              YES     LPAR 2/2 Within Sysplex SANDBOX
****************************** Bottom of data ********************************
```

Figure 28. Select Resource to Create State

Now select the resource requirements that you want to associate with the scheduling environment PROD_CICS_ATMS by typing in an S next to the resource names you want to select and press PF3. Figure 29 on page 44 is displayed.

```
  Scheduling-Environments  Notes  Options  Help
------------------------------------------------------------------------------
                    Modify A Scheduling Environment           Row 1 to 1 of 1
Command ===> ____________________________________________________________

Scheduling Environment Name  : PROD_CICS_ATMS
Description  . . . . . . . . . Automated Teller Machines

Action Codes: A=Add  D=Delet  +-----------------------------------------+
                              | State for a resource must be specified. |
                              | Specify either ON or OFF. (IWMAM676)    |
                         Re   +-----------------------------------------+
Action  Resource Name     State       Resource Description
  __    CRYPTO            ON_____     Crypto resource
****************************** Bottom of data ********************************
```

Figure 29. Set Resource State

You are led back to an updated Modify A Scheduling Environment panel which displays a popup window requesting that the resource state be set. You specify ON or OFF as follows:

Resource State This field is required. For each resource name, you must specify either ON or OFF in the Required State field:

ON Specifies that the resource name must be set to ON for a given system, in order for batch jobs associated with this scheduling environment to be eligible for processing on that system.

OFF Specifies that the resource name must be set to OFF for a given system, in order for batch jobs associated with this scheduling environment to be assigned to that system.

```
  Scheduling-Environments  Notes  Options  Resources  Help
------------------------------------------------------------------------------
                  Scheduling Environment Selection List      Row 1 to 10 of 10
Command ===> ____________________________________________________________

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name  Description
  __    BATCH_DFLT                   Default Batch
  __    BATCH_PROD                   Production Batch
  __    CICS_DEV                     Development CICS active
  __    CICS_PROD                    Production CICS active
  __    NO_CICS_DEV                  Development CICS inactive
  __    NO_CICS_PROD                 Production CICS inactive
  __    OVERNIGHT                    Over Night (1800 - 0700)
  __    PRIMESHIFT                   Prime Shift (0800 - 1700)
  4     PROD_CICS_ATMS               Automated Teller Machines
  __    SANDBOX                      Test Sysplex: SC63 and SC64
****************************** Bottom of data ********************************

        +-----------------------------------------------------+
        | Scheduling environment PROD_CICS_ATMS was changed. (IWMAM656) |
        +-----------------------------------------------------+
```

Figure 30. Browse Resource Requirements

After you press PF3 and go back to the panel shown in Figure 30, it shows that the scheduling environment has been changed.

Selecting the Browse option (4) from the Scheduling Environment Selection List panel and pressing Enter produces the panel shown in Figure 31, for the completed scheduling environment PROD_CICS_ATMS and its resource requirements.

```
  Scheduling-Environments  Notes  Options  Resources  Help
------------------------------------------------------------------------------
                  Scheduling Environment Selection List      Row 1 to 10 of 10
Command ===> ____________________________________________________________

Act
     Browse                                          Line 00000000 Col 001 072
     Command ===>                                          SCROLL ===> PAGE
Act  ***************************** Top of Data ******************************
  _
  _  Scheduling Environment Name. . PROD_CICS_ATMS
  _  Description. . . . . . . . . . Automated Teller Machines
  _
  _   Resource Name     State  Resource Description
  _   -------------     -----  --------------------
  _   CRYPTO            ON     Crypto resource
  _  ************************** Bottom of Data ******************************
  4
  _
***
```

Figure 31. Display of Resources and States

2.8.1.1 Resource Names

There can be more than one resource name listed in a scheduling environment. Conversely, the situation can also exist in which there may be no resource names listed in a scheduling environment. If the resource requirements in an existing scheduling environment are no longer needed, that scheduling environment could remain empty. This would mean that every system would be eligible to process jobs that were associated with that scheduling environment. You may have up to 999 unique resource names defined in a sysplex.

2.9 Using Job Scheduling Environments

You may now assign a job to a scheduling environment, thus ensuring that it executes on certain members in the sysplex which have the required resources or specific environmental states. The states may be associated with hardware resources such as a vector facility, an address space (DB2) or an abstract state such as offpeak (time-of-day), or a certain day of the week. The environment is identified either by the *SCHENV=* keyword parameter on the job statement, or through an installation exit, as shown:

```
//RONNATM1 JOB (POK,999),NORTHRUP,CLASS=W,
//             MSGCLASS=T,NOTIFY=&SYSUID,
//             SCHENV='BATCH_ATM',TYPRUN=HOLD
     ...
```

For this example, assume the following defintions have been made:

- The scheduling environment is BATCH_ATM
- The resource names for scheduling environment BATCH_ATM are:

```
Scheduling Environment Name. . BATCH_ATM
Description. . . . . . . . . . Production Batch For ATMs

  Resource Name     State  Resource Description
  -------------     -----  --------------------
  ATM_CICS          ON     ATM CICS Region
  DBRP              ON     DB2 Production Datasharing
  PRD_SYS           ON     Production Online System
```

A scheduling environment name on a job card associates this batch job with the same system as the one which processes the CICS ATM online work. This is shown in the preceding resource description, which displays the BATCH_ATM scheduling environment.

The SCHENV= parameter is either coded by the job submitter or enforced automatically by an installation exit.

Note: If the scheduling environment name is unknown, the job fails with a JCL error during conversion, accompanied by an appropriate error message as follows:

```
IEFC452I RONNATM - JOB NOT RUN - JCL ERROR
                ...
//RONNATM  JOB (POK,999),NORTHRUP,CLASS=W,MSGCLASS=T,NOTIFY=&SYSUID,
//             SCHENV='ATM_SUITE',TYPRUN=HOLD
                ...
IEFC690I SCHEDULING ENVIRONMENT ATM_SUITE DOES NOT EXIST IN THE WLM
         SERVICE DEFINITION
```

In order for a system to be eligible to process a batch job with the scheduling environment of BATCH_ATM specified on the job card, all of the resource requirements (CRYPTO, DBRP and PRD_SYS) on that system would have to be set to ON.

Scheduling environments may be available on all or only selective systems in the sysplex depending on the criteria defined in the WLM service definition. As shown previously, the WLM administrative application is used to define the scheduling environments and make them available or unavailable on each system, based on the ON or OFF state of resources.

2.9.1 Job Flow with Scheduling Environments

The scheduling environment specified for batch jobs is validity-checked by the converter. If the scheduling environment is not defined to the workload manager, the jobs fail with a JCL error. There is no default scheduling environment. The submittor can either specify it on the job statement parameter, or it can be assigned using an installation exit.

JES detects the availability of scheduling environments on each member and allows an initiator to select jobs only if the specified scheduling environment is available. This is true for both JES-managed initiators and WLM-managed initiators. In JES2, any member affinity specified through the SYSAFF= parameter as well as the job class and initiator class must match.

Figure 32 illustrates the flow of a job that assigns a scheduling environment previously defined in the WLM service definition. The primary focus in OS/390 Version 2 Release 4 is on JES2; however, other work management environments will use it in the future.

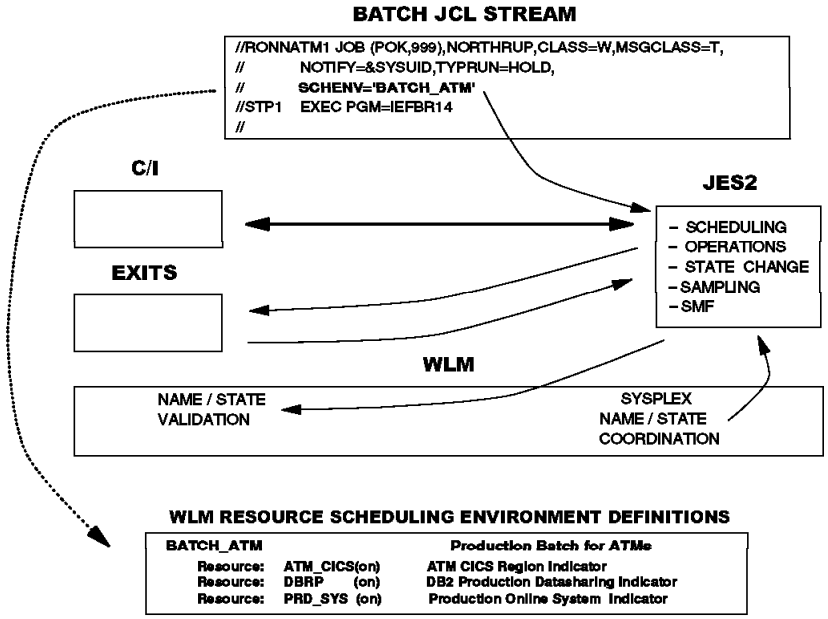

Figure 32. Resource Scheduling Components

JES2 exits can be modified to change the scheduling environment name associated with batch jobs during JCL conversion. These exits can also be used to dynamically generate scheduling environment associations as work is submitted. This would be useful if you where to change from using existing scheduling mechanisms to scheduling environments. Refer to *OS/390 JES2 Installation Exits* for more information.

2.9.2 Scheduling Environment Considerations

The following list contains considerations for implementing scheduling environments:

- Jobs that are part of a critical path, such as those that must complete in the over-night batch window, should remain with JES-managed initiators.
- WLM does not support a secondary JES2 subsystem in the same MAS as the primary. Jobs are not initiated in this configuration.
- It is unnecessary that all systems in the sysplex be in the same MAS complex.
- The WLM policy is sysplex-wide and independent of the JES2 member or MAS complex.
- Jobs with the same classification attributes (job class, priority, for example) are assigned the same service class, irrespective of the MAS they are in.
- Jobs in the same service class have the same goals, irrespective of the MAS they are in.

- All jobs in the sysplex with the same service class should be managed by the same type of initiation. Splitting a service class between WLM-managed and JES-managed initiators causes conflict within WLM′s decision processes. The JOBCLASS MODE= settings should be the same within all MASes in the sysplex.
- JOBCLASS() XEQCOUNT is only MAS-wide, not sysplex-wide.
- Batch delays are factored into the formula for velocity only for work running on WLM-managed initiators.
- Be aware that if you have velocity performance goals set for work running under WLM-managed batch initiators, the batch delays are factored into the velocity formula.
- During normal shutdown, all WLM service queues are de-registered and JES2 disconnects from WLM. JES2 waits for all WLM initiators to terminate before shutting down.
- Priority is not used directly to order jobs in WLM-managed classes, but can be used as a service classification criteria (PRI).
- Once a scheduling environment is assigned to a job, you cannot change it or remove it with operator commands.
- There is no default scheduling environment.

2.10 Displaying and Controlling Scheduling Environments

There are MVS and JES2 commands that you can use to display information about the scheduling environments and associated resource states.

To display information about a scheduling environment on all systems in the sysplex, issue the following command from an MVS console:

```
D  WLM,SCHENV=scheduling_environment
```

To display information about this BATCH_ATM scheduling environment:

```
D WLM,SCHENV=BATCH_ATM

IWM036I  10.27.45  WLM DISPLAY 011
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   NOT AVAILABLE ON ANY SYSTEM
```

This response shows that the scheduling environment is valid but is unavailable as some or all of the associated resource requirements do not have their resource states set correctly on any system. To display more detailed information about a scheduling environment, use the SYSTEM=system_name parameter:

```
D WLM,SCHENV=BATCH_ATM,SYSTEM=SC63

IWM037I  10.37.52  WLM DISPLAY 013
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   SYSTEM:                SC63
   STATUS:                NOT AVAILABLE
                       REQUIRED  CURRENT
    RESOURCE NAME      STATE     STATE
   *ATM_CICS           ON        RESET
   *DBRP               ON        RESET
   *PRD_SYS            ON        RESET

D WLM,SCHENV=BATCH_ATM,SYSTEM=SC64

IWM037I  10.38.15  WLM DISPLAY 015
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   SYSTEM:                SC64
   STATUS:                NOT AVAILABLE
                       REQUIRED  CURRENT
    RESOURCE NAME      STATE     STATE
    ATM_CICS           ON        ON
   *DBRP               ON        RESET
    PRD_SYS            ON        ON
```

Using this command, you can see all of the resource names included in the scheduling environment, along with the required and current states. Resource requirements that are not satisfied are marked with an asterisk. In this situation, SC64 is the system where we want the BATCH _ATM scheduling environment made available. To achieve this, set the resource state for DBRP to ON. See 2.10.1, “Managing Resource States” for details.

2.10.1 Managing Resource States

For every resource name and its corresponding resource state that is defined as part of a scheduling environment, a corresponding resource state must be set on every system in the sysplex. The individual system resource state can be:

ON Satisfies a resource requirement of ON.

OFF Satisfies a resource requirement of OFF.

RESET Specifies that the status of the resource is unknown and consequently not in any of the required states.

When a service definition policy is activated, the corresponding resource state of all resource names is set to RESET on all systems until you explicitly modify that state to either ON or OFF. At IPL, all existing resource states are set to RESET. For as long as the RESET settings remain, resources will not match any resource requirements (whether they be ON or OFF). This means that batch jobs selecting particular scheduling environments associated with resources in this state cannot be scheduled on any system in the sysplex until the resource requirement is changed to the value defined in the service definition.

2.10.2 Resource Name Example

A job requesting resource DB2A in the ON state could be used to guarantee that the job only runs when that copy of DB2 is available. Similarly a job could request DB2A in the OFF state to guarantee that the job only runs when that copy of DB2 is active but is not being used for production purposes.

Resource states can be manipulated in two ways:

1. By using the MODIFY WLM operator command:

```
F  WLM,RESOURCE=resource_name,setting
```

   Where the setting can be ON, OFF, or RESET.

   This is how you would set the resource name DBRP to conform to the value (ON) defined in the service definition on system SC64, and the response you would receive:

```
ROUTE SC64 F WLM,RESOURCE=DBRP,ON

IWM039I RESOURCE DBRP IS NOW IN THE ON STATE


D WLM,SCHENV=BATCH_ATM,SYSTEM=SC64

IWM037I  10.40.31  WLM DISPLAY 018
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   SYSTEM:                SC64
   STATUS:                AVAILABLE
                      REQUIRED  CURRENT
    RESOURCE NAME     STATE     STATE
    ATM_CICS          ON        ON
    DBRP              ON        ON
    PRD_SYS           ON        ON
```

2. By using the WLM application programming interface IWMSESET.

   You can refer to *OS/390 MVS Programming: Workload Management Services*, GC28-1773 for more information on using this interface.

As your installation's usage of scheduling environments grows, you will find it more practical to let automation control the setting of the various resource states across the sysplex. Automation can manage resource settings:

- By listening for messages from a subsystem that indicate that the initialization process has completed and that the subsystem is ready to accept work. The automation script can issue the appropriate F WLM,RESOURCE=subsystem,ON command on that system. When messages are issued indicating that the subsystem has terminated, the script can issue the appropriate F WLM,RESOURCE=subsystem,OFF command on that system.
- For time-related resource settings, a simple script can turn settings on and off at certain times of the day. For example, the PRIMESHIFT resource name could be used to disallow the scheduling of certain resource-intensive batch jobs on production online systems during normal business hours.

More information relating to automation, WLM services, and the coordination between other job scheduling programs can be found in *OS/390 MVS Programming: Workload Management Services*

Note: When you modify the resource state settings on a given system, you do so on that system only.

When we modified the DBRP resource state on the SC64 system, it had no affect on the DBRP setting on SC63. If the intent was to modify the settings on both systems, it would have been necessary to explicitly direct the commands to each individual system.

When all of the resource state settings on a particular system match the resource requirements defined in the scheduling environment, only then is the system eligible to process batch jobs associated with that scheduling environment.

2.10.3 System Command Changes

New information has been added to the following system commands in order to accommodate the functional enhancements to Workload Manager Release 4. They are:

DISPLAY WLM,SCHENV=sname	Display the status of scheduling environments.
DISPLAY WLM,RESOURCE=rname	Display the status of resources.
MODIFY WLM,RESOURCE=rname	Alter resource state settings (either to ON, OFF or RESET).
VARY WLM,POLICY=pname,REFRESH	Discard historical data for the named policy.

For ease of reference, each of the commands is summarized here. Refer to *OS/390 Version 2 Release 4 System Commands*, for more comprehensive information.

Note: The F WLM command *cannot* be specified in the COMMNDxx parmlib member.

2.10.3.1 DISPLAY Command

The following changes to the D WLM command are highlighted as shown.

```
D WLM,SYSTEM=sysname|,SYSTEMS
     ,APPLENV=applenvname|*
     ,SCHENV=schenvname,SYSTEM=sysname|,SYSTEMS
     ,RESOURCE=resourcename,SYSTEM=sysname|,SYSTEMS
     ,L={a|cc|cca|name|name-a}
```

The command displays workload management service policy, service definition, mode, application environments, resources, and scheduling environments.

WLM	Displays the name of the active service policy in effect on all systems in the sysplex, and the time and date that the service policy was activated. It also provides information about the installed service definition.
SYSTEM=sysname	Displays the status of the named system.
SYSTEMS	Displays the status of all systems in the sysplex.

APPLENV=applenvname|* Displays the status of the named application environment. Use * to display all application environments.

SCHENV=schenvname Displays status information for the named scheduling environment. You can display multiple scheduling environments by using wildcard characters (* and ?).

RESOURCE=resourcename Displays status information for the named resource. You can display multiple resources by using wildcard characters (* and ?).

2.10.3.2 MODIFY Command

This command switches modes or resource states, or changes the resource state only on the system where you issue the command. On an OS/390 Version 2 Release 4 system, it has an additional new function: to switch the settings of a resource state for scheduling environment resources.

Note: Before switching to goal mode, your service definition must be installed in the WLM couple data set and a service policy activated. Entering the F WLM,MODE=GOAL command completes the mode switch. Use the F WLM command to:

- Switch the workload management mode in effect on a system
- Switch the settings of a resource state
- Change the state of a resource

Note: The F WLM command cannot be specified in the COMMNDxx parmlib member.

The highlighted syntax is new for the F WLM command as shown:

```
F WLM,MODE={GOAL|COMPAT}
     ,RESOURCE=resourcename,{ON|OFF|RESET}
```

RESOURCE=resourcename Changes the state of resourcename.

- **ON** Specifies that if the required state of the resource in a scheduling environment is ON, then that requirement is satisfied on the target system.
- **OFF** Specifies that if the required state of the resource in a scheduling environment is OFF, then that requirement is satisfied on the target system.
- **RESET** Specifies that this resource setting does not satisfy an ON or an OFF resource requirement. Therefore, if a scheduling environment includes resourcename in its list of resources (whether ON or OFF), then that scheduling environment is not available on the target system.

2.10.3.3 VARY Command

The VARY command is used to activate a workload management service policy for a sysplex.

```
V WLM,POLICY=policyname{,REFRESH}
```

WLM,POLICY=policyname Specifies the 1- to 8-character name of the service policy to be activated.

The VARY command is used to control an application environment. This command has sysplex scope.

```
V WLM,APPLENV=applenvname,{REFRESH}
                          {,QUIESCE|Q}
                          {,RESUME}
```

WLM,APPLENV=applenvname Specifies the 1- to 32-character name of the application environment.

REFRESH Terminates the server address spaces for an application environment and starts new ones in their place.

QUIESCE|Q Terminates the server address spaces for an application environment. Work requests for the application environment are queued but not selected.

RESUME Restarts the server address spaces for an application environment. Work requests that are queued are eligible for selection.

2.11 Migration Considerations

In order to use the WLM functions in OS/390 Version 2 Release 4, you must allocate a new WLM couple data set from an OS/390 Version 2 Release 4 system. This is done using the IXCL1DSU utility. You have the choice of setting the job up yourself or using the ISPF WLM administration application to build it for you. WLM administration activities relating to the reading and writing of service definitions to and from the WLM couple data set require an OS/390 Version 2 Release 4 allocated WLM couple data set. This is true regardless of whether you are using the new WLM functions or not.

One method of migrating to OS/390 Version 2 Release 4, where you have a mixed-release sysplex, is to temporarily continue using the pre-OS/390 Version 2 Release 4 WLM administrative application. This allows you to defer the reallocation of the WLM couple data set. However, be aware that you will not be able to install a service definition or activate a service policy from an OS/390 Version 2 Release 4 system. You need to use the same level of the WLM

administrative application to perform these tasks as was used to create or modify the service definition. When you migrate the last system in your sysplex to OS/390 Version 2 Release 4, allocate a new WLM couple data set at the same time.

Because most of the OS/390 Version 2 Release 4 enhancements to WLM are applicable only to JES2 at this stage, there are special considerations and dependencies, including the following common points:

- OS/390 Version 2 Release 4 must be installed on each participating system.
- The OS/390 Version 2 Release 4 level of JES2 must be installed on each participating system.

  Once satisfied with the stability of the new levels of MVS and JES2, you must enable use of the new functions of JES2 Release 4 by turning on a switch via the $ACTIVATE JES2 operator command. This ensures that no JES2 members are permitted in the MAS that are not at the OS/390 Version 2 Release 4 level. This command requires confirmation, because once processed, lower-level JES2 members attempting to join the MAS will be prevented from doing so. If at a later stage you decide to go back to a mixed level MAS, a cold start will be necessary.

  The results of entering the $ACTIVATE command are as follows:

```
$HASP895 OS/390 RELEASE 4 FUNCTION IS NOW ACTIVE
$HASP895 JES2 CHECKPOINT LEVEL IS NOW OS/390 RELEASE 4
```

- The WLM couple data set must be allocated at the OS/390 Version 2 Release 4 level to accommodate the addition of scheduling environments in the service definition.

2.11.1 Migrating to Batch Initiator Management

In addition to the common points just mentioned, the following function-specific points must be noted:

- Batch Initiator Management:
  - The OS/390 Version 2 Release 4 level of SDSF must be installed on each participating system.
  - You must identify sets of JES2 job classes that will be migrated as a whole. For example, JES2 initiators with job classes W, L, and M in their selection list could indicate that W, L, and M will migrate as a set.
  - You need to be aware of the effect that batch delays, for job classes in WLM-managed mode, will have on your velocity values and adjust your goals accordingly.
  - You may need to change existing automation scripts, such as those that currently examine the JES2 job queue and start initiators.
  - You may need to develop operational procedures for dynamically switching job classes into and out of WLM-managed mode.
- Scheduling Environments:
  - Determine what scheduling environments and resources you need to set up and define them using the WLM ISPF application.

- You may need to develop code to bridge between your current scheduling mechanisms and scheduling environments. This may include name mapping (in exits) or changes to existing automation scripts.
- You may also need to revise dependent software products to coordinate resource state settings.
- Finally, you need to educate your end-users and application developers on how to exploit scheduling environments. You must show the which specific names are in use, and how to associate their incoming work with those scheduling environments.
- You may want to consider implementing a JES2 conversion exit to enforce the association of specific scheduling environments. This would be particularly useful to divert development work away from production images.

- Use of JES2 job priority:
  - Educate your scheduling team and end users regarding use of the PRI classification rule filter. Discuss how best to use it within your installation.

2.11.2 Migration With A New WLM Couple Data Set

The steps shown in the following sections show a migration to OS/390 Version 2 Release 4 workload management from a system already running in goal mode. This method could be used where you have chosen to convert the WLM couple data set to the OS/390 Release 4 format by allocating a new one. This allows you to use the Release 4 WLM administrative application to edit, install, and activate a service definition, and to also make use of the new functions.

If you have systems at lower levels in the sysplex, we recommend that you do not use any new functions until all the systems have been upgraded to OS/390 Version 2 Release 4. Systems at lower levels can install and activate to the newly formatted WLM couple data set provided none of the new functions have been used.

- **Step 1**: Extract the live service definition from the active WLM couple data set, using the pre-OS/390 Version 2 Release 4 WLM administrative application, and save it to a partitioned data set.
- **Step 2**: Make any necessary adjustments to your service definition.
- **Step 3**: Determine the size of the existing WLM couple data set.

2.11.3 Sizing a WLM Couple Data Set

To determine the size of a WLM couple data set, it is necessary that you know how much space to allocate for the new one. The most straightforward method, as shown in the following figure, is to run a job that dumps the WLM couple data set to a flat file (ADRDSSU) and then use the couple data set size utility (IWMARSZ) to give you the details you require.

Example of sizing a WLM couple data set

```
//RONNWMSZ JOB (999,POK),NORTHRUP,CLASS=A,MSGCLASS=T,
//             NOTIFY=&SYSUID
//STP1    EXEC PGM=ADRDSSU,REGION=6M
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//WLMVOL   DD  DISP=SHR,UNIT=3390,VOL=SER=TOTDS0
//WLMDMP   DD  DSN=&&WLM,DISP=(NEW,PASS),UNIT=3390,
//             SPACE=(CYL,(20,10),RLSE),STORCLAS=SCCOMP,
//             VOL=SER=TOTS1
//SYSIN    DD  DATA,DLM='||'
  PRINT        DATASET(SYS1.XCF.WLM00)                        -
               INDDNAME(WLMVOL)                               -
               OUTDDNAME(WLMDMP)                              -
               TOLERATE(ENQFAILURE)                           -
               SHR
||
//STP2    EXEC PGM=IKJEFT01,COND=(4,LT,STP1)
//SYSTSPRT DD  SYSOUT=*
//SYSEXEC  DD  DSN=SYS1.SBLSCLI0,DISP=SHR
//WLMDD    DD  DSN=&&WLM,DISP=(SHR,PASS)
//SYSTSIN  DD  DATA,DLM='||'
  EXEC 'SYS1.SBLSCLI0(IWMARSZ)' 'WLMDD JCL'
||
//

The output from the second step produces the following values:

Number of Policies.................: 10    1
Number of Workloads................: 35      2
Number of Service Classes..........: 100       3
```

Note: If the new WLM couple data set is smaller than the current one, you will not be able to make the new one available to the sysplex. The Cross-System Coupling Facility (XCF) will issue message IXC250I as follows:

```
SETXCF COUPLE,TYPE=WLM,ACOUPLE=(SYS1.XCF.WLM0,TOTDS0)
IXC255I UNABLE TO USE DATA SET
SYS1.XCF.WLM00
AS THE ALTERNATE FOR WLM:
ALLOWABLE NUMBER OF SYSTEM RECORDS IS LESS THAN CURRENT PRIMARY
```

2.11.4 Allocate a WLM Couple Data Set

- Allocate the new WLM couple data set.

The most straightforward way to do this is to setup a batch job that executes the XCF utility, IXCL1DSU. However, this can also be done using the facility provided in the WLM administrative application from an OS/390 Version 2 Release 4 system. In both instances, use the information obtained from the sizing activity to fill in the requested fields which determine the size of the service definition, namely:

- Number of policies 1

- Number of workloads 2
- Number of service classes 3

The following example illustrates WLM couple data set allocation using a batch job. If you would prefer to use the WLM administrative application, refer to Appendix A, "WLM CDS Allocation using Panels" on page 229.

Example of job to allocate R4 WLM couple data sets

```
1. Set up a job that will execute IXCL1DSU and create the new
   WLM couple data sets (primary and alternate) on the DASD
   of your choice:

//RONNWCDS JOB (POK,999),NORTHRUP,CLASS=A,MSGCLASS=T,NOTIFY=&SYSUID
//STEP1    EXEC PGM=IXCL1DSU
//STEPLIB  DD   DSN=SYS1.MIGLIB,DISP=SHR,UNIT=3390,VOL=SER=034RSA
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   *
   DEFINEDS SYSPLEX(WTSCPLX1)
            DSN(SYS1.WLMR4.CDS01)
            MAXSYSTEM(32)
            VOLSER(TOTDS0) NOCATALOG
         DATA TYPE(WLM)
              ITEM NAME(POLICY) NUMBER(10)
              ITEM NAME(WORKLOAD) NUMBER(35)
              ITEM NAME(SRVCLASS) NUMBER(100)
   DEFINEDS SYSPLEX(WTSCPLX1)
            DSN(SYS1.WLMR4.CDS02)
            MAXSYSTEM(32)
            VOLSER(TOTDS1) NOCATALOG
         DATA TYPE(WLM)
              ITEM NAME(POLICY) NUMBER(10)
              ITEM NAME(WORKLOAD) NUMBER(35)
              ITEM NAME(SRVCLASS) NUMBER(100)
/*
//

2. The STEPLIB is only required if this job
   runs on a pre-OS/390 Version 2 Release 4 system to ensure that
   Release 4 couple data sets are allocated.

3. The STEPLIB is only required if this job
   runs on a pre-OS/390 Version 2 Release 4 system.
```

2.11.5 Activate the WLM Couple Data Set

- **Step 1**: Make the Release 4 WLM couple data set available to the sysplex.

To achieve this, you use a series of SETXCF commands to switch from the currently active primary and alternate couple data sets to the new Release 4 couple data sets. All systems in the sysplex then operate with the newly allocated data set.

With OS/390 Version 2 Release 4 formatted WLM couple data sets available to the sysplex, you can continue to use a pre-OS/390 Version 2 Release 4 WLM

administrative application to modify, install and activate your service definition, or you can begin using the OS/390 Release 4 WLM application.

Cycle New WLM Couple Data Set Into Sysplex

```
1. Make SYS1.WLMR4.CDS01 the alternate using the command:

      SETXCF  COUPLE,TYPE=WLM,ACOUPLE=(SYS1.WLMR4.CDS01,SBOX02)

   As part of this processing, SETXCF copies the
   contents of the current primary WLM couple data set
   to SYS1.WLMR4.CDS01, which now is the new alternate.

2. Switch SYS1.WLMR4.CDS01 to primary by using the command:

      SETXCF  COUPLE,TYPE=WLM,PSWITCH

3. Now make SYS1.WLMR4.CDS02 the new alternate by using the command:

      SETXCF  COUPLE,TYPE=WLM,ACOUPLE=(SYS1.WLMR4.CDS02,SBOX02)

   As in Step 1, this causes the contents of the
   new primary WLM couple data set SYS1.WLMR4.CDS01 to be
   copied to the new alternate, SYS1.WLMR4.CDS02.
```

- **Step 2**: To start the WLM application, use a level of the administrative application that is compatible with the functionality level you desire in your service definition. If you wish to use OS/390 Release 4 function, then you must use the OS/390 Release 4 application. If you have systems at pre-OS/390 Release 4 levels in the sysplex, it is recommended that you do not use the new functions until all those systems are upgraded to OS/390 Release 4.
- **Step 3**: Make your changes to the existing service definition by editing it, through one of the following choices:
  - Extract the existing service definition from the existing WLM couple data set.
  - Read the service definition that you previously saved.
- **Step 4**: Install a service definition on the WLM couple data set. Use the same level of the administrative application to install the service definition as you used to edit in Step 3.
- **Step 5**: Activate a service policy.

  Once you have installed a service definition, you can activate a service policy. Use the same level of the administrative application to activate the service definition as you used to to edit and install the service definition. You can activate a policy either from the administrative application, or with the VARY operator command.

  To activate a service policy from the application, choose the Utilities option from the menu bar on the definition menu.

  To activate a service policy with the VARY command, specify VARY WLM,POLICY=xxxx, where xxxx is the name of a policy defined in the installed service definition. Once you issue the command, there is an active policy for the sysplex. Systems in compatibility mode are still processing with the

existing IEAIPSxx and IEAICSxx parmlib members. Systems in goal mode start managing system resources to meet the goals defined in the service policy.

For more information about the VARY command, see *OS/390 MVS System Commands*.

2.11.6 Compatibility Considerations

The following points summarize the steps you will need to follow if you intend to continue using your existing WLM couple data set:

- Check your service definition.

  One change you may need to make to your existing service definition relates to the number of service classes you have defined. If you have more than 100 defined, you will need to reduce the number to 100 or less. With this release, a maximum of 100 service classes are supported. OS/390 Version 2 Release 4 will not activate a service policy which is associated with more than 100 service classes.

- Install required service to downlevel systems.

  If you are running with mixed OS/390 releases in a sysplex, install the compatibility PTFs for APAR OW20913 to all pre-OS/390 Version 2 Release 4 systems. This permits different levels of WLM to coexist until you can upgrade the entire sysplex to OS/390 Version 2 Release 4.

2.11.7 Service Definition Functionality Levels

Each service definition has a functionality level associated with it. When you install a service definition, the system you are installing from determines the functions you have used and sets the level accordingly. If you do not use any of the new functions for a release, then the functionality level does not change, even if you are using the service definition on a new release.

As an example, if you created your service definition on MVS/ESA SP 5.1, then its functionality level is LEVEL001. If you then upgraded to MVS/ESA SP 5.2 and installed this same service definition from the new system but did not use any of the new functions, then the functionality level would remain LEVEL001.

Figure 33 illustrates the service definition functionality levels relating to respective MVS releases.

Release	Service Definition Function Level
MVS/ESA SP 5.1	LEVEL001
MVS/ESA SP 5.2	LEVEL002
OS/390 R3	LEVEL003
OS/390 R4	LEVEL004

Figure 33. Service Definition Functionality Levels

Use of the following OS/390 Version 2 Release 4 functions changes the service definition level to LEVEL004:

- Defining scheduling environments
- WLM managed batch initiators
- Use of priority classification qualifier (PRI)
- Use of SYSTEM and SYSSTC service class names within installation-defined work classification rules

Figure 34 illustrates the functions that can be performed from a system depending on the functionality level of the service definition.

Service Definition Functionality Level	WLM Function Performed	WLM Rel Lvl MVS/SP 5.1	WLM Rel Lvl MVS/SP 5.2	WLM Rel Lvl OS/390 R3	WLM Rel Lvl OS/390 V2.4
LEVEL001	Open/Save PDS	YES	YES	YES	YES
	Extract from CDS	YES	YES	YES	YES
	Install/Activate Pre R3 CDS Format	YES	YES	NO	NO
	Install/Activate R3 Format CDS	YES	YES	YES	NO
	Install/Activate R4 Format CDS	YES	YES	YES	YES
LEVEL002	Open/Save PDS	NO	YES	YES	YES
	Extract from CDS	NO	YES	YES	YES
	Install/Activate Pre R3 CDS Format	NO	YES	NO	NO
	Install/Activate R3 Format CDS	NO	YES	YES	NO
	Install/Activate R4 Format CDS	NO	YES	YES	YES
LEVEL003	Open/Save PDS	NO	NO	YES	YES
	Extract from CDS	NO	NO	YES	YES
	Install/Activate Pre R3 CDS Format	NO	NO	NO	NO
	Install/Activate R3 Format CDS	NO	NO	YES	NO
	Install/Activate R4 Format CDS	NO	NO	YES	YES
LEVEL004	Open/Save PDS	NO	NO	NO	YES
	Extract from CDS	NO	NO	NO	YES
	Install/Activate Pre R3 CDS Format	NO	NO	NO	NO
	Install/Activate R3 Format CDS	NO	NO	NO	NO
	Install/Activate R4 Format CDS	NO	NO	NO	YES

Figure 34. Functionality Level to WLM Release Cross-Reference

Generally, the WLM administrative application must be at the same level as the MVS system it runs on if you plan to use functions that access the WLM couple data set. Referring to Figure 34, note the following points:

- The WLM administrative application can process any service definition at a functionality level equal to or lower than the one that corresponds to the WLM application release level.
- To install a service definition and activate a service policy with the OS/390 Version 2 Release 4 WLM application, you must use a Release 4 allocated WLM couple data set.

2.12 Workload Manager Base Enhancements

The remainder of this chapter describes the other new capabilities of Workload Manager (WLM) in Release 4 of OS/390.

OS/390 Release 3 and OS/390 Release 4 added important new functions to Workload Manager. The new functions can be grouped in three categories:

- Improvements to batch processing, as described previously
- New workloads managed by WLM
  - OpenEdition fork and spawn performance
  - Domain Name System Server supports TCP/IP session placement decisions similar to those of VTAM generic resources.
  - Control of system services classes for started tasks
- WLM internal enhancements
- Buffer pool management

2.13 OpenEdition Fork and Spawn Function Calls

The initial implementation of UNIX fork and spawn function calls within MVS OpenEdition made use of APPC/MVS services as a vehicle for routing requests into MVS address spaces. The objective of this support is to integrate the OpenEdition fork and spawn functions into the MVS workload manager address space management mechanisms, replacing the APPC stack and APPC scheduler paths.

Fork A program running in an address space, called the *parent process*, issues a fork() causing the system to copy one process into a new process, called the *child process*. It places the child process into a new address space, called the *forked address space*

Spawn The spawn-callable service creates a child process to run a specified hierarchical file system (HFS) executable file. The spawn callable service combines the semantics of the fork and exec callable services to provide this function. If the calling parent task is in a WLM enclave, the child is joined to the same WLM enclave. This allows WLM to manage the parent and child as one *business unit of work* entity for system accounting and management purposes.

This should result in better performance for all applications that make use of OpenEdition fork and spawn. This WLM support includes:

- Using existing WLM application environment administrative definitions to describe OpenEdition execution address space environments.
- New WLM interfaces for accepting fork and spawn OpenEdition work requests.
- Using goal-oriented address space management functions to control the creation, pool management, and destruction of server address spaces serving fork and spawn requests.

2.13.1 Open Edition Fork and Spawn Performance

Before Release 4, OpenEdition used APPC/MVS services to create an address space for fork and spawn. This was an asynchronous process. The requesting task had to wait for the creation of the address space. APPC initiator number limits could considerably increase this wait time.

By taking advantage of WLM Server Address Space Management, the cost in terms of CPU instructions of creating a child task via fork or spawn UNIX instruction is considerably decreased.

2.13.1.1 Server Address Space Management

OS/390 Release 3 introduced server address space management, which introduced into WLM controls on the number of address spaces that serve queues from work managers such as DB2 Stored SQL Procedures and DSOM. In OS/390 Release 4, this support is extended to batch initiators and Open Edition fork and spawn address spaces.

An application environment is a way to group similar server programs together and have workload management dynamically create and delete server address spaces as needed to handle the work. Each application environment typically represents a named group of server functions that require access to the same application libraries. Depending on the subsystem′s implementation of application environments, the scope of server address space management is either confined to a single system or is sysplex-wide.

The design of server address space management addresses the following questions:

- If the objectives are not met, can an additional server improve the performance index?
- If there is a resource constraint (CPU or storage) in the system, how do you reduce the activity of server address spaces?
- When should the number of server address spaces be decreased?
- Will the creation of a new server address space adversely impact the performance of other, more important, goals?
- How do you report resource consumption and response time of units of work that span more than one business of work?

The work managers DCE, DSOM, and DB2 stored SQL procedures use WLM services to permit workflow from their network attachment address spaces, through WLM, and into server address spaces for execution. The three components that participate in meeting the objectives, shown inFigure 35 on page 63, are:

Work manager A subsystem or work manager routes transactions to WLM, identifies the server JCL to WLM, and provides shell services for applications.

WLM Creates or deletes server address spaces.

Directs work into the server address spaces.

Decides when a new server address space has to be created.

Reports on goal achievement.

SRM Monitors the performance of *transaction environments*, gathers performance information, and takes decisions for SRM or WLM actions.

The server address spaces managed by WLM, as shown in Figure 35, are created as started tasks to serve work associated with an application environment. The workload requirements for a given application may determine that multiple server address spaces be activated. Figure 35 shows five server address spaces to support the workload. The address spaces making up an application environment are identical in terms of capabilities, as they are all created using the same JCL procedure, and they are interchangeable. An application request can be processed by any of the address spaces comprising the application environment, with identical results.

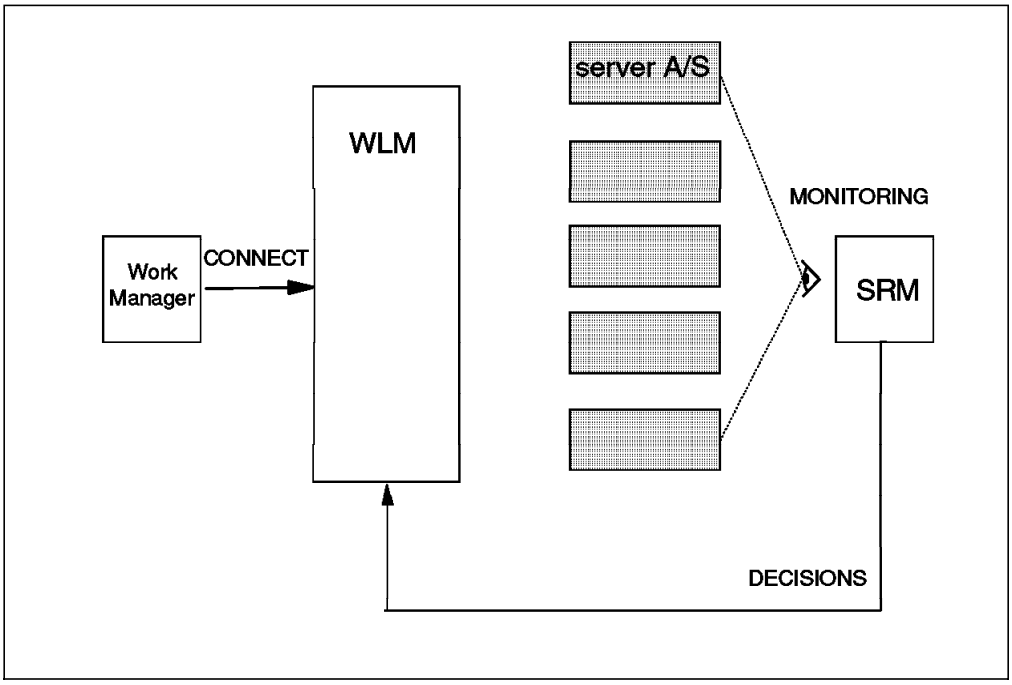

Figure 35. Server Address Space Management Components

The functions of the server address spaces are determined by the set of resources allocated during address space creation, as defined in the startup JCL. This can include definitions of:

- The accessible application program
- Compiler run-time libraries
- Connections to file systems
- The initial program to be given control

Within each address space, an environmental *shell* is provided by the associated work manager. This shell may be either the job step program (as with DB2 stored SQL procedures), or may be run-time services invoked by the executing application programs.

To use server address space management, the following events should be considered:

- The WLM administrator must *define* the application environment to WLM.
- At subsystem initialization, the work manager must *connect* to WLM.
- The work manager, WLM, and the server address space must cooperate to start a new transaction.
- SRM and WLM monitor the execution of the transaction.
- WLM generates information for RMF.
- The operator may control the execution of the application environment.

For a more comprehensive description, see *OS/390 Release 3 Implementation*, SG24-2067 and *OS/390 MVS Planning: Workload Management*, GC28-1761.

In OS/390 Release 4, WLM maintains a free pool of server address spaces to be managed on behalf of OpenEdition fork and spawn requests. A request for a server space is synchronous if the free pool of address spaces is not empty, as shown in Figure 36. When a request arrives, WLM resumes the SRB of the initiator task if the free pool of address spaces When awakened, the initiator task correlates itself to the pending fork/spawn request. When the child task ends, its initiator SRB is suspended, and the address space returns to the free pool.

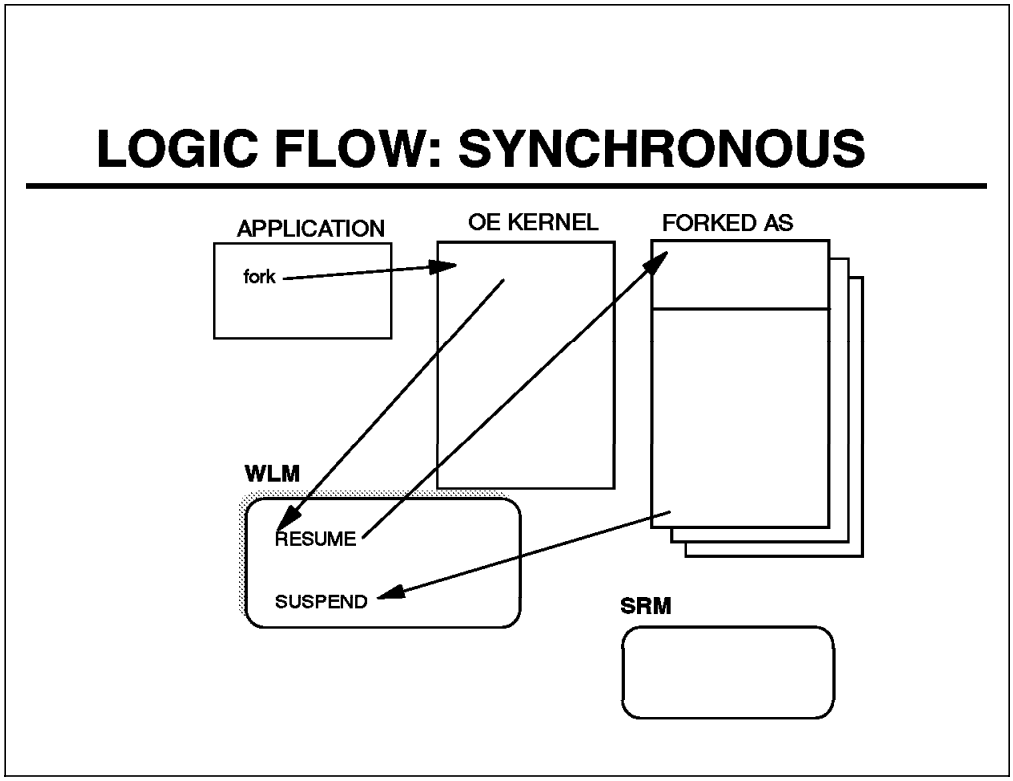

Figure 36. WLM Fork and Spawn Synchronous Processing

If the server address space pool is empty, WLM returns a non-zero return code to OpenEdition, creates an ECB, and the application goes into a wait state, as shown in Figure 37 on page 65. A request for an additional server space is queued to the Environment Manager subcomponent of WLM and then control is returned to the requesting process. The OE kernel then places the requestor into an ECB wait, pending availability of a server space. WLM creates a new server address space, adds it to the pool and posts the ECB.

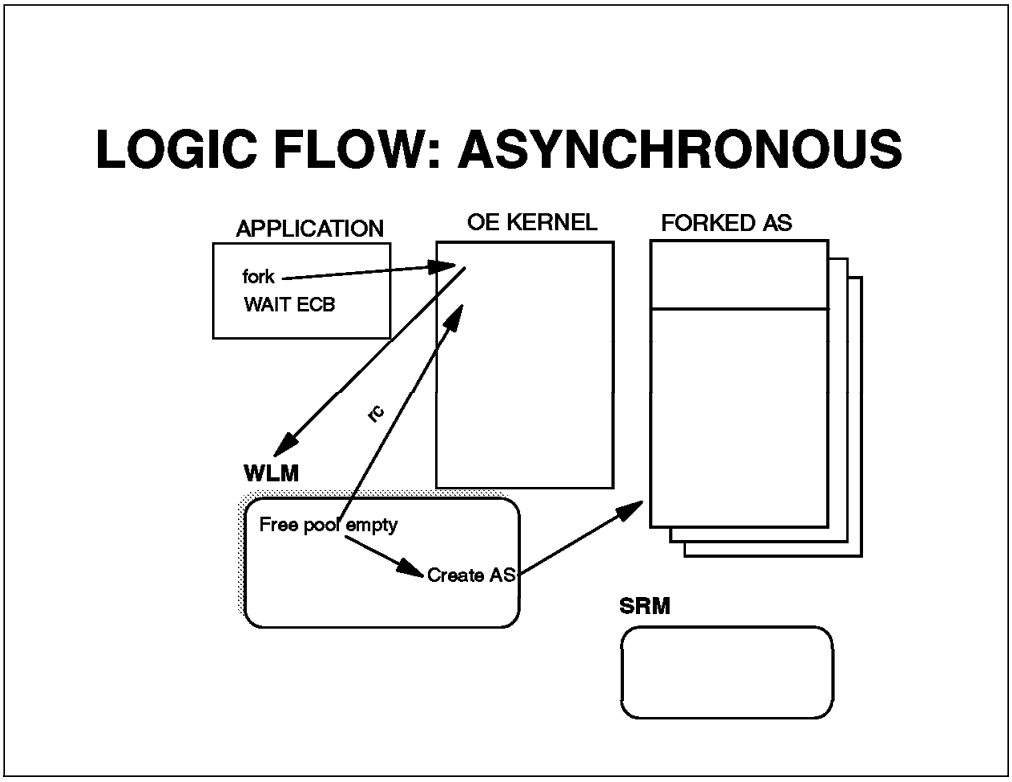

Figure 37. WLM Fork and Spawn Asynchronous Process

This enhancement is operational in either goal or compatibility mode.

2.14 Domain Name Server

For SNA networks, support already has been provided for the intelligent distribution of SNA network connections across multiple systems in a sysplex configuration. This support was provided in MVS/ESA SP 5.2.0 and OS/390 Release 3 for intelligent distribution of client connection to SNA-based applications through cooperation with generic resources support delivered by VTAM. This SNA solution is based upon APPN architecture, that has as its foundation well-formed relationships between application servers (destination LU names) and network sessions. Equivalent capabilities are not available for client connections using TCP/IP network transport protocols.

Domain name server (DNS) support provides a general purpose sysplex connection distribution mechanism for TCP/IP-based workloads. This support provides an integration between existing OS/390 Parallel Sysplex characteristics and the emerging TCP/IP-based client server application environments. WLM services are provided to:

- Register the availability of work managers supporting TCP/IP-based applications with WLM
- Maintain sysplex topology information about the available work management environments
- Dynamically resolve server location requests for an MVS-based domain name server
- Be used by exploiting workloads through direct calls from the C programming language

In a TCP/IP environment, users connect to an application within a *host*. In this much different environment, a clients direct connection requests to a host anticipates that the desired application server exists on that host. The DNS component of the TCP/IP product provides support similar to that provided for SNA applications and now provides a new name resolution mechanism based upon WLM awareness of application server topology and WLM knowledge of available capacity within a sysplex.

2.14.1 Client Connections

A host is identified by a composite name such as wtscplx1.itso.ibm.com. The request first goes to the corporate node IBM.COM, which directs it to zone itso.ibm.com, where wtscplx1x.itso.ibm.com can be identified. These servers may run on any platform (not only OS/390).

Starting with the TCP/IP product that comes with OS/390 Release 4, you can define a server for the sysplex environment.

You define a domain name server *authoritative* for a set of domains which represent families of servers having the same capabilities. To do so, you must include a record in the server boot file identifying the domain owned by the server, as shown in Table 2.

Type	Zone name	File name	Action
Primary	wtscplx1.itso.ibm.com	ns.data	Cluster

Table 2. Server Boot File

Primary A type of primary indicates that this copy of DNS provides authoritative answers for names ending with the zone name. This implies that other instances of DNS within the network should only resolve these names as directed by this instance of DNS.

Zone name The name assigned to a host.

File name The file name is the name of a file containing the address records for each zone. For WLM-managed zones the file name is optional, but it should be used by installations to enable the use of existing network administration tools

Action The keyword, cluster, is appended to the zone definition to trigger DNS use of WLM for resolving names within this domain. The cluster attribute on a DNS zone definition informs DNS of the following facts:

- Names represented by the domain name are not described within DNS configuration files.

- DNS must dynamically retrieve the list of supported domain names from WLM.
- DNS is to apply intelligent routing logic in the resolution of referenced names.

This new boot file record is unique to the MVS platform.

2.14.2 DNS and WLM Example

Figure 38 on page 68 shows the connection of a user to a sysplex. In this example, SC47 is the *primary* host, which means that it is authoritative for names ending with wtscplx1.itso.ibm.com. The other systems in the sysplex only resolve these names as directed by the primary.

Two WLM services are used by DNS to dynamically maintain address resolution information. The WLM services are:

IWMDNLOC This service is used by DNS to retrieve the list of server groups known to WLM. A server group WTSCPLX1 exists. The list instructs DNS to internally make a domain name eligible for resolution: wtscplx1.itso.ibm.com, defined with the cluster attribute. These names can be logically thought of as being sub-domain names (TCP/IP) or server group names (MVS).

Using the list of eligible server group names, DNS retrieves routing information required to intelligently resolve requests for the domain name.

IWMDNSRV This service, when used, supplies a server group name and in response receives a table of information related to that server group. In the example two WLM calls would result, one for each user for wtscplx1.

In the example in Figure 38 on page 68, a request for wtscplx1 results in a table having three rows, with entries for SC50, SC51, and SC52.

The information returned within each row includes a list of IP addresses supported by each server instance and WLM weight information used for DNS decision-making.

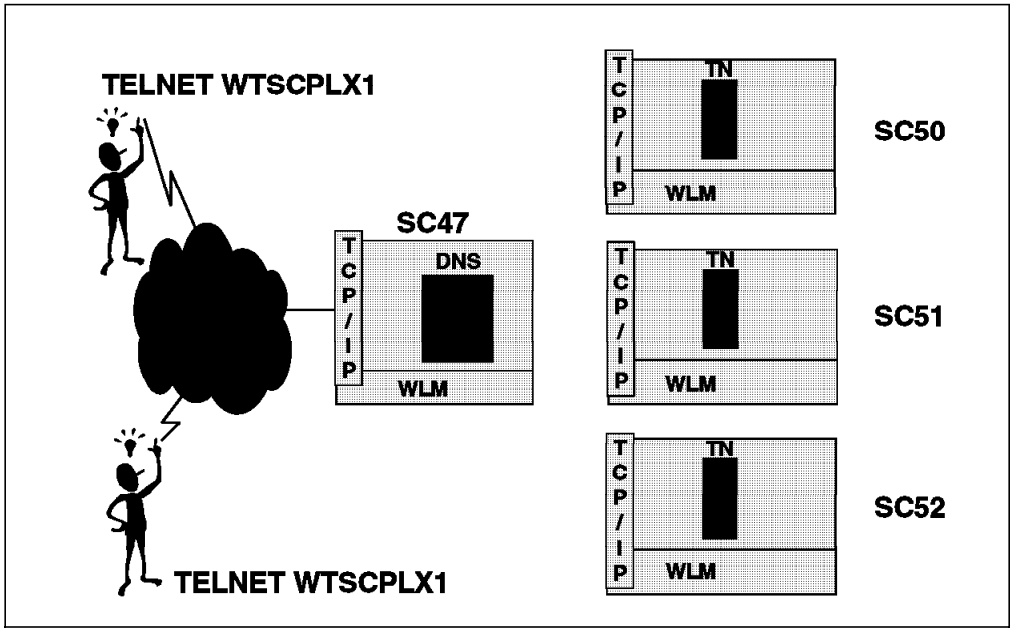

Figure 38. Connecting to a TCP/IP Application

In the figure, two users are connecting to applications TELNET in the sysplex. In both cases, to be able to connect, the user must know the name of the application (TELNET) and the name of the node (wtscplx1.itso.ibm.com).

The way DNS places the sessions depends on the contents of the action entry on the server boot file (see Table 2 on page 66):

- If action does not specify *cluster*, sessions are placed in round robin on the different hosts.
- If action specifies cluster, then the WLM algorithm governs session placement.

For the WLM algorithm, WLMs in other systems in the sysplex periodically send a *weight* to the WLM in the primary. Weight is a number between 0 and 64 which indicates the capacity available on the sending system. If SC51 sends a weight of 5 and SC52 send a weight of 20, the first 10 sessions are placed round robin on SC51 and SC52. The next 15 sessions go to SC52. After 25 sessions, the cycle restarts. A weight of 0 indicates that a system suffers a resource shortage and should receive no sessions.

2.15 WLM Internal Enhancements

The following are changes are made to WLM internal logic for OS/390 Release 4.

- Local performance index evaluation
- Less disruptive policy activation
- Service definition comments

  WLM ISPF panels are modified to include comments to entities such as work classification groups and work classification rules.

2.15.1 Local Performance Index Evaluation

Before OS/390 Release 4, SRM decisions were mainly influenced by the sysplex-wide accomplishment of service class period objectives. This could lead to situations where, even though goals were globally achieved on the sysplex, one particular system could have very poor performance. In OS/390 Release 4, the sysplex-wide performance continues to be predominant but more attention is paid to individual systems. This eliminates the chance of a few users having a very bad response time while the overall average is meeting its goals.

WLM and SRM logic that identifies service classes requiring attention is driven by the performance index (PI) for each service class, periodically calculated to reflect how well each class has been performing with respect to the customer-defined goal for the class. Multiple indices are maintained for each class: a local PI is maintained for each individual system, and a sysplex PI reflects the aggregate for all systems in the configuration. The initial implementation of logic using the indices, which runs on each system in the sysplex, focuses first on the sysplex PI; classes meeting their goals from a sysplex perspective are deemed to be acceptable. The local PIs on each system are only examined if all classes of work are meeting their goals based upon the sysplex PIs, or if no action can be taken to help classes currently missing their sysplex goals; additional logic is then performed to adjust the highest importance service class not meeting goals on the other systems.

2.15.1.1 Sysplex PI Evaluation

Figure 39 on page 70 shows a high-level overview of the SRM policy adjustment loop.

Select receiver With each SRM policy adjustment interval, the current search algorithm first determines if an action is required to assist a service class that is not achieving its goal from a sysplex perspective.

Each service class period is examined, ordered by sysplex performance index value within service class importance, starting with the most important period that is performing the poorest and ending with the least important period that is performing the best. Periods are evaluated until an adjustment is possible, at which point policy adjustment actions are complete for the interval.

Find bottleneck The evaluation of each service class period consists of identifying resource bottlenecks that SRM understands and performing the resource algorithms associated with those bottlenecks.

Resource algorithms The resource algorithms determine if there is an acceptable donor for the resources needed to address those bottlenecks, and they assess the value of adjustments to decide if the adjustments are worth making.

Fix routines If actions are deemed to have value, the appropriate fix routines are invoked to carry out those actions on behalf of the service class period in question. Having made one adjustment, processing for the interval is complete.

Actions on behalf of a service class period are not performed for a large number of reasons, most likely because the period is not experiencing delays due to

resources that SRM manages or because there is no projected value to making any of the possible adjustments. If no actions are taken, the loop continues with the next service class period, which has a lower sysplex PI or has a lower importance level.

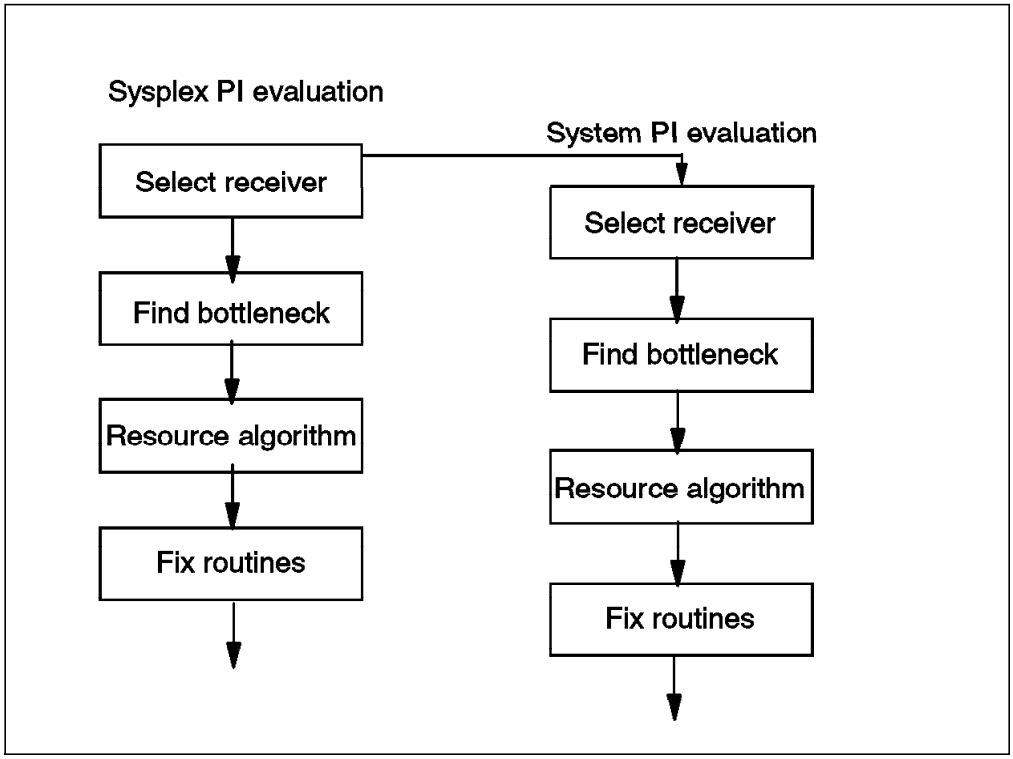

Figure 39. SRM Policy Adjustment Loop

Policy adjustment processing is performed on each system within the sysplex, with each system contributing actions toward successful achievement of the defined service class period goals.

2.15.1.2 OS/390 Release 4 PI Evaluation

With the previously described sysplex PI evaluation loop shown in 2.15.1.1, “Sysplex PI Evaluation” on page 69, each system determines if it can make an adjustment that will have a positive impact on the sysplex PI of a period. Taking no action to help a sysplex PI does not, however, imply that actions cannot be taken to improve system performance.

With the current policy adjustment implementation, the previously described loop is performed a second time, this time focusing on the service class period PI as seen by the local system. This second loop, shown in Figure 39, refines the current situation, given that it is only performed if all service class periods are either already meeting their goals or cannot be helped. An adjustment is made, if possible, to help a service class period from the local system perspective. This compensates for overachievement on one system cancelling out the effects of poor performance on other systems processing the same service class periods.

The policy adjustment loop was modified to interleave local system PI evaluations into the sysplex PI evaluation loop. Sysplex PI evaluation remains the top priority; however, the local PI evaluation is no longer deferred until all sysplex PI processing is complete. Upon completion of sysplex PI evaluation for

a given importance level, the local PIs are examined for the previous higher importance level periods. The sequence will be: importance level 1 sysplex PIs, importance level 2 sysplex PIs, importance level 1 local PIs, importance 3 sysplex PIs, and so forth.

2.15.2 Less Disruptive Policy Activation

In previous releases, changing a WLM policy caused the reset of history data for jobs. In OS/390 Release 4, the history is kept whenever possible so the system adapts more quickly to the new policy.

WLM decisions are based on the historical behavior of the entities under its control. If the policy is changed, WLM needs to start accumulating new historical data from scratch, and it takes some time before the system steadies.

The change described in this paragraph is based on the idea that there are two types of policy activations:

- Structural changes, when the administrator decides to change the definition of entities
- Routine changes, done by a human or automated operator, usually at shift transition

In the second case, some of the historical information remains valid. To be more precise, the following three cases allow you to preserve at least part of the historical data:

- Changing the service class periods and service class period goals
- Altering the relationship between resources and service classes
- Altering the maximum capacity of resources

If for any reason you want to force an erasure of historical data, you may do so with the following command: `VARY WLM,POLICY=`*`name`*`,REFRESH`

Chapter 3. JES2 Enhancements is OS/390 Release 4

OS/390 Release 2 Version 4 JES2 supports the new enhancements for the workload manager:

- Batch initiator management
- Affinity resource scheduling

This support for the workload manager has initiated many other changes to JES2 as follows:

- Operator commands
- JES2 checkpoint
- Initialization statements

This chapter describes the changes made to JES2.

Note: See OS/390 *JES2 Migration Notebook*, GC28-1797 for complete details about this release of JES2.

3.1 Batch Initiator Management

During initialization processing, JES2 connects to WLM as a resource manager. If this connection to WLM fails, JES2 initialization terminates.

JES2 uses the WLM-provided IWMCONN service during initialization to connect to WLM. This initial contact with WLM during initialization specifies the following options:

WORK_MANAGER(YES) This parameter identifies the role that the connecting address space performs.

SUBSYS(JES) This parameter identifies the set of WLM classification rules to be used.

SUBSYSNM(name) This parameter, when combined with the SUBSYS-TYPE value, uniquely identifies the subsystem instance.

This information physically binds the JES2 address space to WLM, enabling use of WLM work classification services by JES2 and establishing WLM address space termination recovery.

JES2 initialization also synchronizes the JES2 job queue with the WLM policy in effect. This can occur when a new WLM policy is activated, when JES2 may have terminated when a new policy was activated, or when new initialization parameters may be in effect governing the JES2 restart. JES2 initialization determines if job queue reconstruction is necessary by comparing the level of WLM service definition that was used to structure the JES2 job queue with the currently active service definition level in effect on the active system. If different, all queued jobs must be reclassified and requeued in a manner consistent with the currently active WLM service definition. This queue consistency checking and reconstruction is only necessary if jobs are resident in the queue associated with WLM service class queues.

3.1.1 Classifying Batch Work

Figure 40 illustrates the flow of a job from the reader through the execution phases. The job processing flow for a normal batch job includes a reader phase, a conversion phase, and then an execution scheduling phase. The reader phase gathers relevant information about the submitted job and writes the data to JES2 spool records. The conversion phase passes the input job stream to the MVS Converter/Interpreter component, for syntactical validity checking, structuring the information for use during execution, and extraction of job-related information needed by the JES2.

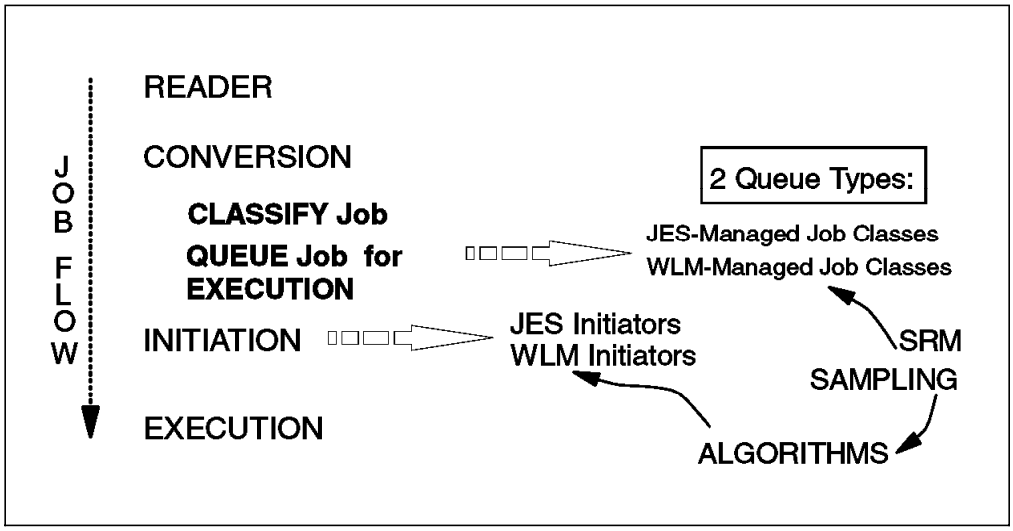

Figure 40. Batch Job Flow Phases

New logic is injected at the completion of conversion processing to perform WLM work classification, passing many of the job attributes extracted by the conversion phase. Prior to OS/390 Release 4, the classification of jobs for WLM purposes was performed during job initiation.

Classify Job: Following conversion, JES2 calls the WLM IWMCLSFY service to classify the job. Passed to WLM work classification are many of the job attributes extracted by the conversion phase, as shown in Figure 41.

```
IWMCLSFY USERID=$DWORK2,  Owning userid
      TRXCLASS=JQXJCLAS,  Job class
      TRXNAME=JQEJNAME,   Job name
      ACCTINFO=(R6),      Accounting string
      ACCTINFL=$DWORK+4,  Address of acctg length
      PERFORM=(R5),       Address of Perf group
      PRIORITY=$DWORK,    Priority
      CONNTKN=(R4),       WLM connect token
      SERVCLS=JQASTOK,    Service class token (rtn)
      SRVCLSNM=JQAWSCN,   Service class name  (rtn)
      MF=(E,$CLSFY)       Address of list form
```

Figure 41. JES2 Job Classification

On return from WLM classification, JES2 receives:

- A four-byte classification token
- A service class name

This early classification enables introduction of goal-oriented logic within JES, driven by the WLM view of job goals and importance levels.

The classification token is kept by JES2 until execution complete. It contains three key pieces of WLM information:

- An index to the WLM service class name
- An index to the report class name
- An indicator that a policy activation has not occurred since classification was done

The service class is used by JES2 to queue jobs for execution scheduling. The service class assigned to a job may be changed by a JES2 operator command, activation of a WLM policy, or through an MVS address space-related operator command.

Note: WLM classification is not performed for started tasks and TSO logons.

3.1.1.1 Using Priority During Classification

Classification of batch jobs is performed independent of the WLM mode in effect once the proper level of JES2 is installed. The intention is that JES2 not be sensitive to or aware of WLM modes and mode transitions.

New WLM support is provided to permit the use of the JES2 job priority within the WLM work classification rules. The priority used in the classification rules must be within the numeric range of 0 to 15, which corresponds to the range of job scheduling priorities currently supported by JES2.

3.1.2 Job Execution Queueing

For jobs to be selected by either JES2- or WLM-started initiators, JES2 now maintains two different queue organizations for all jobs awaiting execution, as shown in Figure 42 on page 76:

1. All jobs are queued by job class, priority, and the order in which they finished conversion. This is the queue from which JES2-managed initiators select jobs for execution.
2. Jobs awaiting execution in WLM-managed job classes are also queued by their WLM assigned service class in the order they were made available for execution.

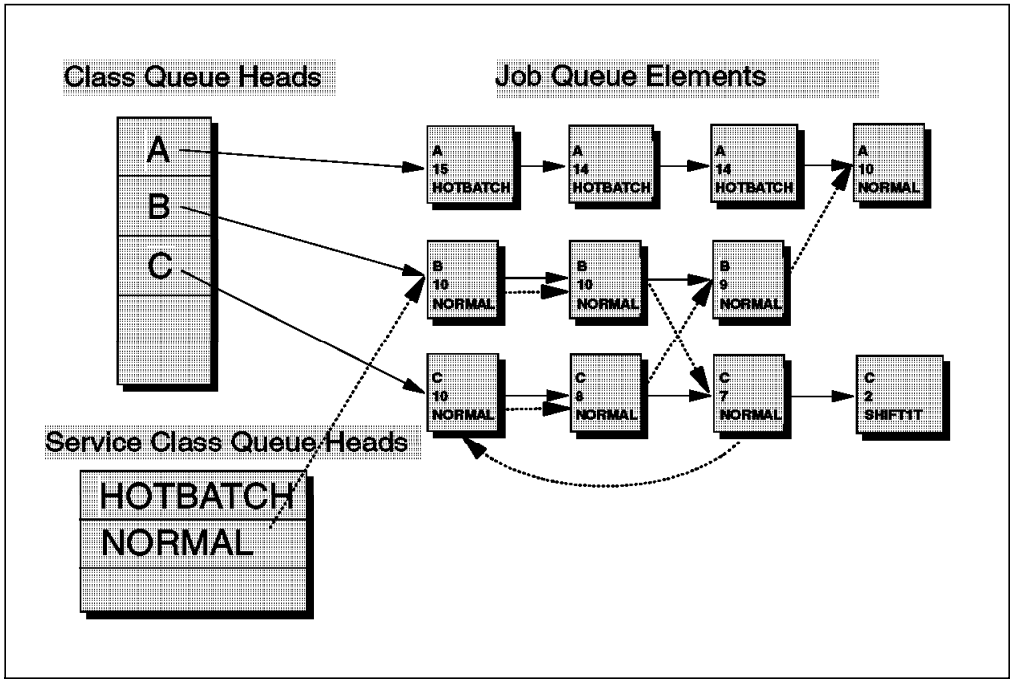

Figure 42. JES2 and Service Class Queues

If a job awaiting execution is in a job class that is managed by WLM, then that job will be on both a job class queue and a service class queue. The JES2 job class queue is ordered the same as the old job class queues (the order a job is added to the queue within priority). The new WLM service class queue is for jobs awaiting execution with a particular service class. Since priority is an input to classification, then priority can be used to dictate which service class queue a job is placed on. In Figure 42, the higher priority jobs are classified into service class HOTBATCH.

Therefore, the service class queues are not in priority order. The queues are ordered by a new time stamp, which is set when a job first becomes eligible to execute.

The JES2 JQE class queue heads are stored in a new JOBCLASS checkpoint section. The new service class queues are also stored in a service class section of the checkpoint. Service class queues are ordered by element age on queue and are not ordered by priority, as shown by the dotted arrow lines in Figure 42.

3.1.3 Job Execution

JES2-managed initiators continue to exist. The class list can include jobs that are in WLM mode. These initiators never select jobs from classes that are in WLM mode. You can change job classes from JES2 mode to WLM mode and back. All initiators, JES2-managed or WLM-managed, honor the existing selection criteria for a job such as, job hold, system affinity values, duplicate job name, spool volume not available, job class held, all job classes held, $PXEQ, job busy (not executing), job class limits, and scheduling environment requirements.

Initiators in WLM mode select only jobs on a service class queue. Processing for these initiators is the same as for normal initiators, except that the service class queue name is passed to JES2 on the job select SSI.

WLM initiators pass a new ECB address to be waited on when there is no work for the initiator. The initiator WAITs on a list of 3 ECBs, the JES2 ECB, the P INIT ECB, and the new WLM ECB. If WLM wants to stop an initiator, it posts its ECB and JES2 returns from the job select SSI.

As jobs are removed from WLM service class queues, the queues eventually empty. Any service class queue that is inactive for more than 1 hour is deregistered on all members and its queue head is freed.

3.1.3.1 Job Select User Exits

JES2 exit 14 is not called for WLM initiators. A new exit 49 is invoked to allow job selection to be stopped for the current job.

3.2 JES2 Checkpoint

A number of new sections are added to the JES2 checkpoint for the WLM support. These new data areas expand the JQE control block with new information for the WLM support.

These new sections are automatically added once all the members are at the Release 4 level, and assuming that the checkpoint data sets are large enough, by doing either:

- A JES2 cold start
- Issuing a new JES2 command, $ACTIVATE

The sections must be present to activate the new Release 4 functions. If the sections cannot be added and an attempt to activate the new functions is made, the activate fails. You should then use the checkpoint reconfiguration process to move to a larger checkpoint data set.

3.2.1 JES2 Job Queue Element

Since it is not possible to increase the size of the JQE without a cold start, a new checkpoint section is added that contains a JQE extension. This fixed size extension contains new information that exists for the life of the job.

A new service is created to access the new JQE. This service also allows installation extensions to be added to the JQE.

With JES2 OS/390 Release 4, the job queue element (JQE) control block is rarely referenced. The management of JQEs in now done using what is called an *artificial* JQE, as shown in Figure 43 on page 78. The artificial JQE that is constructed is called the JQA, and consists of the following parts:

- The JQE which is in the JQE checkpoint table entry (CTENT)
- The spools used by this job mask (SUMSK)
- The JQE extension (JQX) which is in its own CTENT
- Information which is stored in the block extension reuse table (BERT) CTENT

When a JQA is built, almost all of the JQE fields are moved except the offset to the next JQE and the offset of the first work JOE.

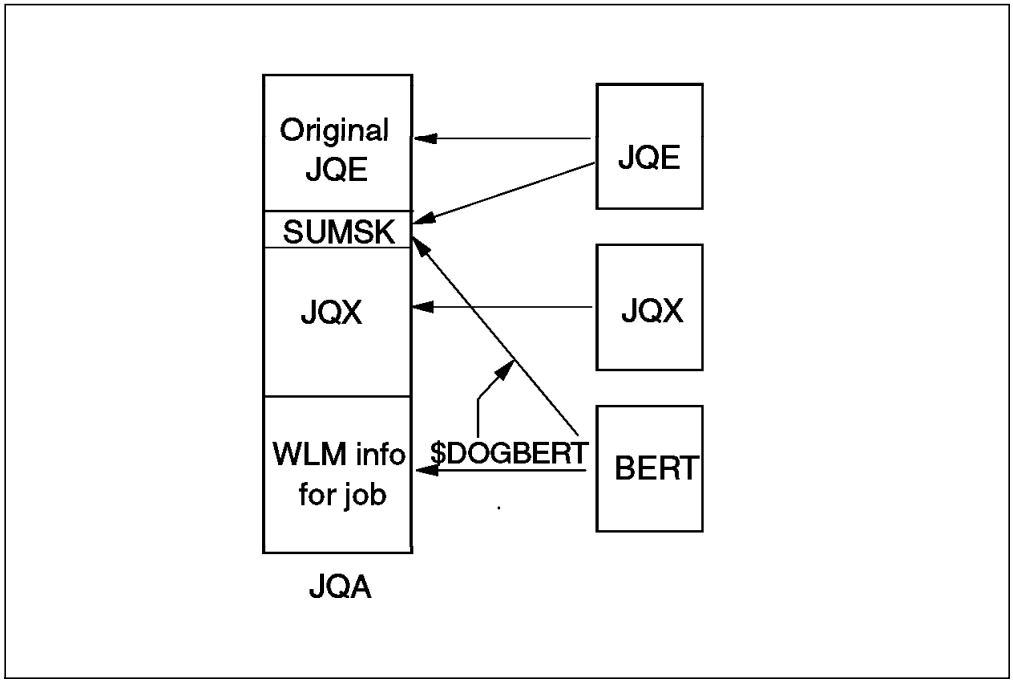

Figure 43. JES2 JQE Control Block

3.2.2 JES2 64K Job Numbers

JES2 Release 4 allows up to 65534 jobs and 161314 output elements. This release of JES2 also allows the number of jobs, output elements, and track groups to be increased via an operator command.

Note: This support can be added to JES2 Release 3 by installing APAR OW20844.

See 3.3.4, “JOBDEF Statement” on page 80 for initialization changes in support of the 64K job numbers.

3.3 JES2 Initialization Changes

The following summarizes the changes to the JES2 initialization statements:

- JOBCLASS
- CKPTSPACE
- Lnnnn.JTn
- JOBDEF

3.3.1 JOBCLASS Statement Changes

The JES2 JOBCLASS statement is expanded and given a MAS scope in this release. A new parameter (MODE=) is added to define whether jobs in this job class are JES2- or WLM-managed. This parameter has a MAS scope. It is modifiable via the operator command ($T JOBCLASS) and retains its value across all JES2 warm starts

Note: All existing parameters on the JOBCLASS initialization statement now have a scope of MAS. All parameters are modifiable via operator command. The initialization statement values are only honored on a cold start. The default is MODE=JES2.

The following changes are made to the JOBCLASS statement in support of the workload manager (WLM) batch initiator. The new keywords are:

MODE MODE=JES|WLM specifies whether jobs are to be run under JES-controlled or WLM-controlled initiators.

XEQCOUNT XEQCOUNT= specifies the maximum number of jobs that can execute concurrently in a class. XEQCOUNT=MAX=nnnnnnnnn limits the number of jobs that can be active in a given class in the MAS. This parameter is modifiable via operator command and does not retain its setting across a JES2 restart (including hot starts). A limit of 0 indicates no jobs will be selected from the job class. A limit of * indicates that there is no limit. The default is LIMIT=*.

Note: JES2 no longer considers the subscript for JOBCLASS(subscript) to be a single character. Therefore, you cannot specify a subscript range such as JOBCLASS(A-Z) as you might have in prior releases.

For example, if you have JOBCLASS(A-9), change it to:

```
JOBCLASS(*)
```

Note: This specification includes JOBCLASS(STC) and JOBCLASS(TSU).

Specifying the following affects only single-character job classes.

```
JOBCLASS(?)
```

If you have JOBCLASS(A-F), specify:

```
JOBCLASS(A,B,C,D,E,F)
```

3.3.1.1 JOBCLASS Migration Considerations

With JES2 Release 4 in Release 4 mode, the JES2 JOBCLASS initialization statements now affect every member of the MAS. The initial values are set from the system that cold starts or issues the $ACTIVATE command. Once set, the JOBCLASS parameters can only be changed via operator command or on a cold start. Once JES2 is running in Release 4 mode, values specified in the initialization deck are ignored on all subsequent warm starts.

3.3.2 Lnnnn.JTn Statement

Lnnnn.JTn is a new statement and it defines a line job transmitter. You can use multiple job transmitters together with work selection parameters WS=(LIMIT= to control the number of jobs transmitted.

You can specify the number of transmitters and receivers on NJE lines globally through the following NJEDEF parameters, which apply to each line generated by the NJEDEF LINENUM parameter:

- JRNUM - Number of Job Receivers for each NJE line
- **JTNUM - Number of Job Transmitters for each NJE line**
- SRNUM - Number of SYSOUT Receivers for each NJE line
- STNUM - Number of SYSOUT Transmitters for each NJE line.

Specifying work selection values for NJE Transmitters allows an NJE transmitter to select files based on their size. This allows you to optimize NJE transmissions. For example, by dedicating some job transmitters to large files and others to small files, you can ensure that small files are not always forced to wait for completion of large file transmissions. You can set the limits based on

the size in cards, lines or pages, depending on whether the file is cards, or page-mode or line-mode data sets. If you specify both page mode and line mode, a job must fit within both limits before JES2 will select it.

The following example shows how to set up two job transmitters, the first one for large jobs, and the second for small jobs:

```
L23.JT1  WS=(LIM/),LIMIT=0-*     /* all files can use  */
L23.JT2  WS=(LIM/),LIMIT=0-100   /* reserved for small */
```

3.3.3 CKPTSPACE Statement

CKPTSPACE is a new statement and it defines how much space JES2 is to allocate to the checkpoint for block extension reuse tables. This checkpoint area is used to store and extend the following data:

- JQEs
- WLM queues
- Other block extension workspace

There are two keywords:

BERTNUM BERTNUM=nnnnnn|0 specifies how much space JES2 allocates to the checkpont for block extension reuse tables (BERTs).

BERTWARN BERTWARN=nnn|80 specifies the percentage of use of BERTs at which the operator is alerted by message $HASP050.

Note: The $T CKPTSPACE command can be used to modify the keyword parameters.

3.3.4 JOBDEF Statement

Changes to the JOBDEF statement are only necessary when you require the use of more than 32767 job numbers. The JOBDEF statement changes are as follows:

JOBDEF

RASSIGN=YES|NO Specifies whether job numbers outside of the RANGE definition can be assigned to jobs received via NJE or spool reload.

YES Jobs with original job numbers outside the JOBDEF RANGE retain their original job number, if possible.

NO Jobs with original job numbers outside the JOBDEF RANGE are always assigned a new job number within the range.

Note: This keyword is new and has single member scope.

RANGE=nnnnn-mmmmm Valid values for nnnnn: 1-65524,

Valid values for mmmmm: n+10-65534

The RANGE parameter can be changed by an operator command as follows:

```
$T JOBDEF,RANGE=nnnnn-mmmmm

   Valid values for nnnnn: 1-65524,
   Valid values for mmmmm: n+10-65534
```

3.3.5 Other Statements for Checkpoint Expansion

This release of JES2 support dynamic expansion of the checkpoint. The three initialization statements that specify the keywords that can now be dynamically expanded are:

JOBDEF

JOBNUM The minimum action to modify this keyword is an operator command to increase the number, and a JES2 cold start to decrease it.

OUTDEF

JOENUM The minimum action to modify this keyword is an operator command to increase the number, and a JES2 cold start to decrease it.

SPOOLDEF

TGSPACE=MAX The minimum action to modify this keyword is an operator command to increase the number, and a JES2 cold start to decrease it.

3.3.6 Operator Commands for Checkpoint Expansion

The following operator commands are used to dynamically expand the JES2 checkpoint by increasing the values of the JQEs, JOEs, and track groups in the checkpoint:

```
$T JOBDEF,JOBNUM=nnnnn

Where:
    nnnnn specifies the maximum number (10-65534) of jobs that
    can be in the JES2 job queue at any given time.  This
    value includes all TSU and STC jobs, as well as batch
    jobs.  This number can only be increased by operator
    command.


$T OUTDEF,JQENUM=nnnnnnn

Where:
   nnnnnnn specifies the maximum number (10-161314) of output
   elements (JOEs) to exist in the JES2 system.


$T SPOOLDEF,TGSPACE=(MAX=nnnnnn)

Where:
   nnnnnnn specifies the maximum number (1-9935680) of track groups
   allowed.


$T CKPTSPACE,BERTNUM=nnnnnnn

Where:
```

nnnnnnn specifies the maximum number (1-9935680) BERTS entries in the checkpoint.

Note: These numbers can only be increased by operator command.

3.3.7 JES2 Initialization Range Changes

The following parameters are changed to support a job number limit of 64K:

JOBDEF	JOBNUM= range is changed from 10-32767 to 10-65534
	RANGE= range is changed from 1-32767 to 1-65534
OFF(n).JR	RANGE= range is changed from 1-32767 to 1-2147483647 and the default is changed from J1-32767 to J1-2147483647
OFF(n).JT	RANGE= range is changed from 1-32767 to 1-65534 and the default is changed from J1-32767 to J1-65534
OFF(n).SR	RANGE= range is changed from 1-32767 to 1-2147483647 and the default is changed from J1-32767 to J1-2147483647
OFF(n).ST	RANGE= range is changed from 1-32767 to 1-65534
OUTDEF	JOENUM= range is changed from 1-120000 to 10-161314
PRT(nnnn)	RANGE= range is changed from 1-32767 to 1-65534 and the default is changed from J1-32767 to J1-65534
PUN(nn)	RANGE= range is changed from 1-32767 to 1-65534
R(nnnn).PR(m)	RANGE= range is changed from 1-32767 to 1-65534 and the default is changed from J1-32767 to J1-65534
R(nnnn)PU(m)	RANGE= range is changed from 1-32767 to 1-65534

3.3.8 JES2 Initialization Statement Processing

On a JES2 warm start, if JES2 detects that certain initialization parameters do not match the settings saved in the JES2 checkpoint, message $HASP442 is issued as follows:

```
  IXZ0001I CONNECTION TO JESXCF COMPONENT ESTABLISHED,
           GROUP WSCJES2   MEMBER WSCJES2$SYSB
  IXL014I IXLCONN REQUEST FOR STRUCTURE J2CKPT1 WAS SUCCESSFUL.
  JOBNAME: JES2 ASID: 001A CONNECTOR NAME: JES2_SYSB
  CFNAME: CF1LPAR
  IEF196I IEF237I 01A8 ALLOCATED TO SYS00001
  $HASP478 INITIAL CHECKPOINT READ IS FROM CKPT1
           (STRNAME J2CKPT1)
           LAST WRITTEN THURSDAY, 13 NOV 1997 AT 09:18:27 (LOCAL TIME)
 *$HASP493 JES2 MEMBER-SYSB QUICK START IS IN PROGRESS
> $HASP442 INITIALIZATION STATEMENTS CONFLICTING WITH SAVED VALUES FOLLOW:
> $HASP496 SPOOLDEF TGSPACE=MAX=16288 SAVED VALUE OF 32576 WILL BE USED
> $HASP496 CKPTSPAC BERTNUM=4100 SAVED VALUE OF 4101 WILL BE USED
  $HASP537 THE CURRENT CHECKPOINT USES 142 4K RECORDS
  IEF196I IEF237I 01A9 ALLOCATED TO $PTS1A9
  IEF196I IEF237I 01AA ALLOCATED TO $PTS1AA
  $HASP492 JES2 MEMBER-SYSB QUICK START HAS COMPLETED
```

Figure 44. Initialization Message $HASP442

This could result from an operator command that updated the parameters while the system was running, from an incorrect attempt that was made to update the initialization statement.

Following $HASP442, the $HASP496 shown in Figure 44 on page 82 displays each incorrect initialization parameter, the value specified on the initialization statement, and the value saved in the checkpoint.

3.4 Operator Commands

The job list commands previously processed by HASPCOMM are now processed by $SCAN. Also, initiator, spool, and the remaining device commands are changed to also use $SCAN. This allows for an easy mechanism for adding additional job set and display keywords to the JES2 commands. The changes to the JES2 command structure provide the following improvements:

- The ability to limit the size of a display by displaying only those subparameters of interest.
  - $DJOBQ,PRIORITY

    This command displays all jobs and lists only the priority.

    ```
    $HASP890 JOB(OSASF)      PRIORITY=15
    ```

- The ability to limit the scope of a command by providing a filter.

  This command displays all jobs with an affinity for system SC47.

  - $DJOBQ,SYSAFF=SC47

    ```
    $HASP890 JOB(ADBDBC1)
    $HASP890 JOB(ADBDBC1)   STATUS=(AWAITING HARDCOPY),
    $HASP890                CLASS=A,PRIORITY=1,SYSAFF=(ANY),
    ```

- The ability to provide multiple subscripts or multiple subscript ranges.
  - $DJ(10,101-102,ROGERSH2)

    This command displays all jobs 10, 101, 102, and a job named ROGERSH2.

- The ability to specify filters on set commands. The same keyword may be used as both a filter and a modify parameter. To use the keyword as a filter, precede the keyword on the command with a slash.
  - $TJOBQ,/P=10,P=15

    This command changes all jobs with a priority of 10 to a priority of 15.

- The ability to enclose symbolic subscripts in single quotes rather than parenthesis.
  - $DJ'D96CLW1A'

    Using single quotes, the characters "*" and "?" are treated as exact character matches, rather than generics.

- The ability to specify filter operations other than equal ("="). Most filterable keywords may now specify the following operations:
  - Equal "="
  - Greater than ">"
  - Less than "<"
  - Greater than or equal ">="

- Less than or equal "<="
- Not equal "<>"
- $DJOBQ,P>=10

  This command displays all jobs with a priority greater than or equal to 10.

  Note: All of these filter operations may be preceded by a "¬ " to form the opposite operation, such as not equal "¬ = . "

- For both job and OUTPUT commands, existing job and OUTPUT level keywords can be used as single keyword displays. Most will also be available as filters. In addition to the existing parameters, several new parameters are available as job display parameters.
- For remote operators, the new filters added to many commands makes it easier to display output jobs for the remote terminal as follows:

  $do jobq,ready,r=rmt2,outgrp

  ```
  JOB03305  $HASP686 OUTPUT(IEBGENER) OUTGRP=1.1.1
  JOB03305  $HASP686 OUTPUT(IEBGENER) OUTGRP=2.1.1
  JOB03305  $HASP686 OUTPUT(IEBGENER) OUTGRP=3.1.1
  JOB03776  $HASP686 OUTPUT(REMOTE02) OUTGRP=1.1.1
  JOB03776  $HASP686 OUTPUT(REMOTE02) OUTGRP=2.1.1
  ```

  Where:

  READY Indicates that only READY output groups are to be displayed.

  OUTGRP Used as a filter to specify to only display the output groups.

3.4.1 Initiator Commands

The JES2 initiator $DI command is modified to display only current, non-held JES2-mode job classes in the class list. WLM-mode job classes and held job classes are displayed on a new INELIGIBLE_CLASS= keyword.

For example, after a $TI,C=ABCDE command, the initiator display would show CLASS=ACE,INELIGIBLE_CLASS=(B-WLM,D-HELD) if class B specified MODE=WLM and class D was held.

Note: The $DI command does not apply to WLM-managed initiators.

To halt and restart job initiation on a particular JES2 member, the new commands $PXEQ and $SXEQ are created. When a $PXEQ is issued, JES2 stops assigning new work to all initiators (JES2- as well as WLM-managed). JES2 also deregisters all service class queues on this member. A $SXEQ causes initiators to be assigned new work, and the service class queues to be reregistered.

3.4.2 Initiate a Batch Job Immediately

The operator can force a job into execution with a new JES2 command $SJnnnnn. For jobs started via a $SJ command, the only difference is that JES2 notifies WLM that a forced execution job exists when the command is entered, and WLM will then have an initiator select jobs with a new force executed indicator. WLM initiators with the forced execution indicator will only be assigned jobs that are forced into execution by the $SJ command. Jobs that are forced into execution can, however, be selected by normal WLM initiators. The $SJ command only applies to jobs in WLM-managed job classes.

3.4.3 JES2 Commands for Scheduling Environments

JES2 can display the scheduling environment for a particular job and list those jobs waiting for a specific scheduling environment. Additional information is displayed when the following commands are processed:

$DJnnnnn,LONG Display the entire environment for a specific job.

```
$DJ1049,LONG

$HASP890 JOB(RONNATM1)
$HASP890 JOB(RONNATM1)  STATUS=(AWAITING EXECUTION),
$HASP890                CLASS=W,PRIORITY=5,SYSAFF=(ANY),
$HASP890                HOLD=(NONE),CMDAUTH=(LOCAL),
                           ...
$HASP890                SRVCLASS=STC_HI,
$HASP890                SCHENV=BATCH_ATM,SCHENV_AFF=(),....
```

In this situation, the job cannot be scheduled because of a problem with the scheduling environment, BATCH_ATM. Some of the associated resources are not in the required state as defined in the WLM service definition. For this reason, the SCHENV_AFF field does not show any system name.

$DJnnnnn,DELAY,SCHENV Display delays caused by any scheduling environment for a specific job.

This is what you would see:

```
$DJ1049,DELAY,SCHENV

$HASP890 JOB(RONNATM1)  DELAY=(SCHENV),SCHENV=BATCH_ATM
```

$DJQ,SCHENV=,SCHENV,SCHENV_AFF Display jobs waiting for a specific scheduling environment.

This is what you would see:

```
$DJQ,SCHENV=BATCH_ATM,SCHENV,SCHENV_AFF

$HASP890 JOB(RONNATM1)  SCHENV=BATCH_ATM,SCHENV_AFF=()
```

The output from this command shows that no scheduling affinity has been assigned. Although the scheduling environment is valid, a system cannot be selected to process the job because one or more resources, associated with the scheduling environment, are not in the required state.

Note: Once a scheduling environment is assigned to a job, you cannot change it or remove it with operator commands.

Using SDSF, you are able to display the job scheduling details for specific jobs. From the Input Queue Display panel, key in and enter the letter I next to the job you want to check. A Job Information pop-up window is displayed with the relevant information. You can also display all of the scheduling environments defined in the WLM service definition, and on which members in the MAS each is available, by using the SDSF **SE** option. If you are an authorized SDSF user, you are able to display and set the state of the resources on each member by

using the **RES** option. You may find this far more practical than entering numerous JES2 commands from a console.

Refer to Chapter 4, “OS/390 Release 4 SDSF Enhancements” on page 93 for further details on the OS/390 Version 2 Release 4 of SDSF.

3.5 JES2 User Exits

Installation exits should be examined to see if they are still required and if so, whether they affected by the job queue changes. In addition, any JES2 exit that processes job information, job class information, the JES2 pre-execution job queue, or the JQE needs to be examined for possible changes as a result of the significant changes made in this release. Consider the following exits:

Exit 5 JES2 command exit. If exit 5 is obsoleted by this support, you may want to examine the sample exit 5 that translates old JES2 command syntax to new syntax for use with translating old local command syntax to new JES2 command syntax.

Exit 14 This exit is not be called for WLM initiator work select.

Exit 49 This is a new exit that gets control after a $QGET selects a job. In this exit, you can reject the element selected. Rejecting the job select can cause a job to be held or selection to continue with the next job.

Exit 44 This exit gets control in the JES2 main task after conversion has completed, but before a job has been classified. This exit could be used to set the scheduling environment for a job.

For more information on the type of control you have over batch jobs entering the system, see *JES2 Installation Exits*, SC28-1793.

3.5.1 JES2 Exit 5

JES2 Release 4 enhanced the job list commands and other commands to include additional filtering capability. In doing so, some commands were changed incompatibly. APAR OW27715 provides a sample exit 5 that is intended for use as a migration aid to convert commands entered in the old format to equivalent JES2 Release 4 commands.

JES2 initialization is updated to automatically include the following statements in the initialization stream:

```
LOAD(HASX05C)
EXIT5  ROUTINES=(HASX5CTR)
```

This automatically enables the command translation. If additional EXIT(5) statements are found in the initialization stream, they override this default. To include the translation function in this case, HASX5CTR should be added to the list of routines on the EXIT(5) statement.

A new message $HASP006 is issued for each command that is translated. This message includes both the untranslated and translated versions of the command:

```
$HASP006 cmd1 TRANSLATED BT HASX05C TO cmd2
```

To disable the command translation, either:

- Remove the HASX5CTR routine from the list of routines on any EXIT(5) statements in your own initialization stream,

  or

- Add an EXIT(5) DISABLE statement to your initialization stream (or issue $TEXIT(5),DISABLE command)

Note: For JES2 Release 4, a warm start or IPL is required only if the new exit HASX05C is to be enabled.

Table 3 on page 88 shows a comparison of the old and new JES2 commands:

Table 3. Comparison of Old and New JES2 Commands

Pre-Release 4 Command	**New Command Translated by Exit 5**
$D′name′,...	$DJOBQ′name′,CMDAUTH=*,...
$T′name′,...	$TJOBQ′name′,...
Note: Similar for $A, $C, $E, $H, $L, $O, $P, $T, $TO	
$DJ1,2,...	$DJ(1,2),...
Note: J can be J, JOB, S, STC, T, TSU.	
Note: Similar for $A, $C, $E, $H, $L, $O, $P, $TO	
$DJ1-2,J3-4,...	$DJ(1-2,3-4)...
Note: Similar for $A, $C, $E, $H, $L, $O, $P, $TO	
$LJnnn,ALL	$DOJnnn
$LJnnn,H	$DOJnnn,HELD
$LJnnn,READY	$DOJnnn,READY
$LJnnn,OUTGRP=xxx	$DOJnnn,OUTGRP=xxx
$CJnnn,OUTGRP=xxx	$COJnnn,OUTGRP=xxx
$PJnnn,OUTGRP=xxx	$POJnnn,OUTGRP=xxx
$PJnnn,Q=x	$POJnnn,Q=x
$vJnnnn,A=|DAYS=|Hours=	$vJnnnn,A>|Days>|Hours>
$TJnnnn,S=sid1,sid2,...	$TJnnnn,S=(sid1,sid2,...)
$DSPL,JOBS=nn	$DJOBQ,SPOOL=(PERCENT>=nn)
$DSPL,V=xxxxxx,JOBS=nn	$DJOBQ,SPOOL=(PERCENT>=nn,
	VOLUME=xxxxxx)
$SSPL,V=xxxxxx,...	$SSPL(xxxxxx),...
$vIxx	$vI(xx)
$TIxx,class-list	$TI(xx),C=class-list
$HQ,ALL	$TJOBCLASS(*),QHELD=Y
$HQ,C=xyz	$TJOBCLASS(x,y,z),QHELD=Y
$AQ,ALL	$TJOBCLASS(*),QHELD=N
$AQ,C=xyz	$TJOBCLASS(x,y,z),QHELD=N
$PQ,ALL,...	$POJOBQ,READY,...
$PQ,Q=xyz,...	$POJOBQ,READY,Q=XYZ,...
$OQ,ALL,...	$OJOBQ,/R=LOCAL.*,...
$OQ,Q=xyz,...	$OJOBQ,/R=LOCAL.*,/Q=xyz,.
$TALL,sid1,sid2,...	$TJOBQ(*),/S=(sid1),S=(sid2,..)
$LSYS	$DMEMBER
$ESYS,sid	$EMEMBER(sid)
$ESYS,RESET=sid	$ECKPTLOCK,HELDBY=sid
$TSYS,IND=Y/N	$TMEMBER(local),IND=Y/N

3.6 Miscellaneous Changes

JES2 extracts and stores the maximum completion code for any locally run job. This completion code is available in the extended status SSI, the $DJ command, and is added to the HASP165 message (JCL notify message). In addition, it is added to the NJE header so that it is available on other nodes. The value stored is the highest step completion code (excluding ABEND codes).

The maximum number of JES2 NJE lines, logons, nodes, printers, and remotes is increased. All (except logons) are increased from 9999 to 32767. The new logon limit is now 999.

Note: If a device is defined above the old limit of 9999, the size of the device name may increase beyond 8 characters.

3.6.1 LOGON(nnn) Statement

The LOGON(nnn) statement identifies JES2 as an application program to VTAM. SNA NJE protocols in JES2 prohibit multiple SNA sessions between two VTAM APPLIDs. With this new support, using multiple APPLIDs per node, each with a unique logon identifier, you can start multiple SNA NJE sessions between two adjacent nodes.

You can define up to 999 unique APPLIDs and LOGONs. At two adjacent nodes, using two or more unique LOGON(n)s and two or more APPLID statements, you can start 2 or more sessions in parallel on two or more different SNA lines.

3.7 Migrating to JES2 Version 2 Release 4

See informational APAR II10760, which documents the current required maintenance when installing JES2 OS/390 Release 4, and before putting JES2 Release 4 into production.

Note: The PSP bucket should be referred to for ongoing maintenance recommendations.

The following APARs are required to avoid a cold start when migrating to this new release of JES2.

- APAR OW26844 (which IF-REQs BCP APAR OW26914)
- APAR OW27317 (which IF-REQs BCP APAR OW27577)
- APAR OW27661

The new JQE Extension (JQX) for an artificial JQE (JQA), which is a new JES2 spool control block, has been expanded by APAR OW27661. If this APAR is applied after migration to JES2 and a $ACTIVATE command has been issued, then a cold start will be required.

If a cold start has already been performed on JES2, then another cold start is required to apply OW27661. Apply the three APARs to JES2 before migrating to the new JES2 to avoid this.

3.7.1 Virtual Storage Considerations

Support for the dynamic expansion on the JES2 checkpoint requires an increase in the amount of virtual storage obtained to represent both the I/O copy and the in-storage copy of the JES2 checkpoint to their maximum size. As a result, it may be necessary to increase the REGION size associated with JES2. A minimum of 100M is recommended; however, more storage may be required in your configuration.

Note: If you specify a REGION size of 0, no action is required.

3.7.2 Checkpoint Storage Considerations

When issuing the $ACTIVATE command or using operator commands to increase the number of JOEs or JQEs, the checkpoint data set size must be large enough to hold the new data areas. You should then use the checkpoint reconfiguration dialog process to move to a larger checkpoint data set.

3.7.3 Understanding JES2 Release 4 Mode

When migrating to JES2 Release 4, you need to be aware of the mode that JES2 is operating in. Once all members of a MAS have migrated to JES2 Release 4, issue the $ACTIVATE command on one of the members to activate the Release 4 functions.

JES2 Release 4 has two modes of operation, as follows:

- There is a compatibility mode or pre-Release 4 mode that allows pre-Release 4 levels of JES2 (back to SP 5.1.0) to coexist in a MAS with Release 4.

  If you install JES2 Release 4 with a warm start, then you are in pre-Release 4 mode.

  In compatibility mode, certain data areas in the checkpoint data set do not exist. This implies that certain functions are not enabled. Functions that are not enabled in pre-Release 4 mode include:

  - WLM-managed job classes
  - Single system image JOBCLASS support
  - Job completion code tracking
  - Greater than 9999 remotes
  - Scheduling environments
  - Job transmitter limits
  - RMF reporting of pre-execution job delays
  - Start job command ($S Job)
  - Age filtering on job level commands (for example, $D J, $T J)

  When Release 4 is warm-started with a checkpoint data set that was used by a pre-Release 4 JES2, it starts in pre-Release 4 mode or compatibility mode.

- The other mode is Release 4 mode.

  If you install JES2 Release 4 with a cold start or issue a $ACTIVATE command when all systems are at Release 4, then you are in Release 4 mode.

  Once the JES2 Release 4 functions have been activated by either a $ACTIVATE command or cold start, there is no way to deactivate them. If you need to fall back to an earlier JES2 release on any MAS member, a cold start is required. If you cold-start a MAS containing both Release 4 and

pre-Release 4 members, you *must* perform the cold start from a pre-Release 4 member.

As mentioned earlier in this chapter, with JES2 Release 4 in Release 4 mode, the JES2 JOBCLASS initialization statements now affect every member of the MAS. The initial values are set from the system that cold-starts or issues the $ACTIVATE command. Once set, the JOBCLASS parameters can only be changed via operator command or on a cold start. Once JES2 is running in Release 4 mode, values specified in the initialization deck are ignored on all subsequent warm starts.

Chapter 4. OS/390 Release 4 SDSF Enhancements

SDSF is an efficient way to control JES2 job processing and devices. It allows you to monitor jobs while they are running, and browse output without printing it. You can also browse the system logs, including the sysplex-wide operations log. SDSF does this with an interactive panel interface that provides immediate, up-to-date information about jobs, output, printers, initiators and other system resources.

OS/390 Release 4 SDSF includes all of the function present in OS/390 release 3 SDSF. OS/390 Release 4 SDSF provides enhancements in the following areas:

- **Scheduling Environment Display**

  OS/390 Release 4 Workload Manager (WLM) allows you to create groups of resources and give them a name. This *scheduling environment* is associated with a job through a new JCL statement.

  SDSF provides a Scheduling Environment panel which simplifies the WLM support for scheduling environments. The Scheduling Environment panel gives a MAS-wide (or optionally, sysplex-wide) display that shows all the scheduling environments that are defined, a description of each, and the systems on which the scheduling environments are available.

- **Resource Display**

  Resources are defined using the Workload Manager (WLM) dialog. A scheduling environment consists of one or more resources and a desired state for each resource (ON or OFF). SDSF allows you to display the resources defined to the system via the Resource panel.

  The Resource panel can be accessed by typing an action character on the Scheduling Environment display. It shows the resources defined through WLM and the state of each (ON, OFF or RESET) for the selected scheduling environment. The Resource display can also be accessed as a primary display, with a command or from a pull-down, in which case all the resources in the MAS (or optionally, the sysplex) are displayed. You can modify the state of each resource on the system.

- **Job Class Display**

  OS/390 release 4 SDSF allows you to manage JES2 and WLM classes using the new Job Class panel. The display shows the classes defined in the MAS, the status of the class, and the number of jobs waiting and active in each class. Action characters and overtypeable fields can be used to hold or release a class or modify characteristics of a class.

- **WLM-managed Batch Support**

  With OS/390 release 4, there are WLM initiators as well as JES2 initiators.

  OS/390 release 4 JES2 has new commands to control WLM managed initiators. SDSF has corresponding action characters to:

  – Force a WLM-managed job to start

  – To start job execution on a MAS member

  – To stop job execution on a MAS member.

  New columns have been added to existing SDSF panels to provide job specific information on WLM-managed jobs.

- **Job Information Pop-up**

  The new Job Information pop-up display shows information that might be responsible for delaying the processing of a job.

  The Job Information pop-up displays information such as:

  – Job name and ID
  – Scheduling environment name
  – Job class limit indicator
  – Job class held indicator
  – Estimated (average) time until execution
  – System or systems which satisfy the resource requirements

- **Maximum Return Code**

  OS/390 release 4 SDSF has a new column on the H, O, and ST displays which shows the maximum return code for each job.

4.1 SDSF Primary Option Menu

The OS/390 Release 4 SDSF option menu has been updated; see Figure 45 for an example of the new primary option menu.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
HQX1800 -----------------  SDSF PRIMARY OPTION MENU  ------------------
COMMAND INPUT ===>                                          SCROLL ===> PAGE

    LOG       - Display the system log
    DA        - Display active users in the sysplex
    I         - Display jobs in the JES2 input queue
    O         - Display jobs in the JES2 output queue
    H         - Display jobs in the JES2 held output queue
    ST        - Display status of jobs in the JES2 queues
    PR        - Display JES2 printers on this system
    INIT      - Display JES2 initiators on this system
    MAS       - Display JES2 members in the MAS
    LINE      - Display JES2 lines on this system
    NODE      - Display JES2 nodes on this system
    SO        - Display JES2 spool offload for this system
    PUN       - Display JES2 punches on this system
    RDR       - Display JES2 readers on this system
    JC        - Display job classes in the MAS
    SE        - Display scheduling environments in the MAS or sysplex
    RES       - Display WLM resources in the MAS or sysplex
    ULOG      - Display user session log
    END       - Exit SDSF
```

Figure 45. Primary Option Menu

The primary option menu has been modified to display the SDSF release identifier (FMID) in the upper left hand corner. The menu also contains new options:

```
JC        - Display job classes in the MAS
SE        - Display scheduling environments in the MAS or sysplex
RES       - Display WLM resources in the MAS or sysplex
```

The new options are described in greater detail in the relevant sections.

4.2 SDSF Display Pull-Down

SDSF gives you a choice of how you enter commands. You can enter the command directly on the command input line, or you can use the action bar and pull-down menus:

```
_Display  Filter  View  Print  Options  Help
-------------------------------------------
```

One of the pull-down menus from the action bar is DISPLAY. It allows you to select an SDSF panel; see Figure 46.

```
                                                  ----------------------------
 __ 1. Status of jobs on any queue (ST)           U  ---------------------------
    2. Active users (DA)                                 SCROLL ===> PAGE
    3. Input queue (I)
    4. Output queue (O)
    5. Held output queue (H)                      x
    6. Printers on this system (PR)               ue
    7. Initiators on this system (INIT)           eue
    8. Members in the MAS (MAS)                   ut queue
    9. User session log (U)                        queues
   10. System log (LOG)                           em
   11. Lines on this system (LI)                  stem
   12. Nodes on this system (NO)
   13. Spool offload for this system (SO)
   14. Punches on this system (PUN)
   15. Readers on this system (RDR)               s system
   16. Job classes in the MAS (JC)                m
   17. Scheduling environments in the MAS (SE)    m
   18. WLM resources in the MAS (RES)
   19. Exit SDSF                                   the MAS or sysplex
                                                  r sysplex
   ULOG      - Display user session log
   END       - Exit SDSF




 F1=HELP      F2=SPLIT     F3=END       F4=RETURN    F5=IFIND     F6=BOOK
 F7=UP        F8=DOWN      F9=SWAP     F10=LEFT     F11=RIGHT    F12=RETRIEVE
```

Figure 46. DISPLAY Pull-down

The DISPLAY pull-down has been modified to contain the new panels:

16. Job classes in the MAS (JC)
17. Scheduling environments in the MAS (SE)
18. WLM resources in the MAS (RES)

The new panels are described in greater detail in the relevant sections.

4.3 Scheduling Environment Display

OS/390 Release 4 Workload Manager (WLM) provides enhanced batch processing capabilities. For a detailed explanation of the new WLM enhancements, see Chapter 2, “Workload Manager (WLM)” on page 13. One of the batch processing enhancements available with OS/390 release 4 WLM is *Resource Affinity Scheduling*. Resource affinity scheduling allows you a flexible method of controlling job scheduling.

A scheduling environment is made up of a group of resources and a desired state for each resource (ON or OFF). A job is linked to a particular scheduling environment through a new JCL JOB statement parameter (SCHENV=*se_name*).

If a job is associated with a scheduling environment, it is not selected for execution by JES2 until the scheduling environment is available. A scheduling environment is available when all the resources that make up the environment are in the desired state (ON or OFF). The state of a resource can only be altered by an operator command:

```
F WLM,RESOURCE=resource_name,ON|OFF|RESET
```

After an IPL, all resources are in RESET state; therefore, the state has to be changed by operator (or automated operator), command.

A simple example of this is a job that requires the production CICS systems (CICSP01 and CICSP02) to be inactive when it executes. You would define a scheduling environment using the WLM ISPF dialog. In this example we have named the scheduling environment NO_CICS_PROD. NO_CICS_PROD contains two resources that we have called CICSP01 and CICSP02, and the desired state for both resources is OFF.

When the CICS systems are started, an operator (or automated operator) would issue the commands:

```
F WLM,RESOURCE=CICSP01,ON
F WLM,RESOURCE=CICSP02,ON
```

At this point, the scheduling environment NO_CICS_PROD is unavailable because the state of resources CICSP01 and CICSP02 is ON and the desired state is OFF. If a job is submitted that is associated to scheduling environment NO_CICS_PROD (see Figure 47), it would not be selected by JES2 for execution.

```
//NORMANDA JOB (POK,999),DALE,CLASS=A,MSGCLASS=T,NOTIFY=&SYSUID,
//             SCHENV='NO_CICS_PROD'
//*
//* N.B. This job requires CICSP01 and CICSP02 to be inactive
//*
```

Figure 47. JCL JOB Statement SCHENV Parameter

If CICSP01 is subsequently inactivated, the operator (or automated operator) would set the state of resource CICSP01 to OFF by issuing the command:

```
F WLM,RESOURCE=CICSP01,OFF
```

The scheduling environment NO_CICS_PROD would still be unavailable because it requires that both resources (CICSP01 and CICSP02) be in OFF state. At this point resource CICSP01 has a state of OFF but resource CICSP02 has a state of ON.

When CICSP02 is inactivated the operator, (or automated operator) would set the state of resource CICSP01 to OFF by issuing the command:

```
F WLM,RESOURCE=CICSP02,OFF
```

The scheduling environment NO_CICS_PROD is now available (resources CICSP01 and CICSP02 are both in OFF state), and the job is selected by JES2 for execution.

In order to support the OS/390 Release 4 Workload Manager (WLM) scheduling environments, SDSF provides the Scheduling Environment panel. The scheduling environment display is a MAS-wide (or optionally, sysplex-wide) display showing all the scheduling environments that are defined, a description of each, and the systems on which they are available.

4.3.1 Invoking the Scheduling Environment Panel

The Scheduling Environment panel can be invoked by either the SE command or from the Display pull-down.

Note: The Scheduling Environment display requires OS/390 R4 JES2 to be active in each member of the MAS.

An example of a command to display the Scheduling Environment panel, showing all scheduling environments in the MAS, is:

```
COMMAND INPUT ===> SE
```

The SE command displays the scheduling environment panel, as shown in Figure 48.

```
  Display  Filter  View  Print  Options  Help
------------------------------------------------------------------------
SDSF SCHEDULING ENVIRONMENT DISPLAY MAS SYSTEMS        LINE 1-9 (9)
COMMAND INPUT ===>                                       SCROLL ==
PREFIX=*  DEST=(ALL)  OWNER=*
NP   SCHEDULING-ENV   DESCRIPTION                     SYSTEMS
     BATCH_DFLT       Default Batch                   SC64
     BATCH_PROD       Production Batch                SC63
     CICS_DEV         Development CICS active         SC64
     CICS_PROD        Production CICS active          SC63
     NO_CICS_DEV      Development CICS inactive       SC63
     NO_CICS_PROD     Production CICS inactive
     OVERNIGHT        Over Night (1800 - 0700)
     PRIMESHIFT       Prime Shift (0800 - 1700)       SC63,SC64
     SANDBOX          Test Sysplex: SC63 and SC64     SC63,SC64
```

Figure 48. Scheduling Environment Panel

An example of a command to display the Scheduling Environment panel, showing all scheduling environments in the sysplex, is:

```
COMMAND INPUT ===> SE ALL
```

The SE command displays the scheduling environment panel, as shown in Figure 49 on page 98.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF SCHEDULING ENVIRONMENT DISPLAY ALL SYSTEMS        LINE 1-9 (9)
COMMAND INPUT ===>                                           SCROLL ==
PREFIX=*  DEST=(ALL)  OWNER=*
NP   SCHEDULING-ENV   DESCRIPTION                     SYSTEMS
     BATCH_DFLT       Default Batch                   SC64
     BATCH_PROD       Production Batch                SC63
     CICS_DEV         Development CICS active         SC64
     CICS_PROD        Production CICS active          SC63
     NO_CICS_DEV      Development CICS inactive       SC63,SC65
     NO_CICS_PROD     Production CICS inactive
     OVERNIGHT        Over Night (1800 - 0700)
     PRIMESHIFT       Prime Shift (0800 - 1700)       SC63,SC64,SC65
     SANDBOX          Test Sysplex: SC63 and SC64     SC63,SC64
```

Figure 49. Sysplex-wide Scheduling Environment Panel

The Scheduling Environment display consists of a row for each scheduling environment. There is a column for the name of the scheduling environment, the description of the scheduling environment, and the the systems on which the scheduling environment is available.

4.3.2 Action Characters on the Scheduling Environment Display

Table 4 shows the action characters available on the scheduling environment panel:

Table 4. Action Characters Available on the Scheduling Environment Display

Action Character	Description
=	Repeat the last command
//	Repeat a block of commands
+	Temporarily extend the NP column
R	Display resources for the scheduling environment
ST	Display all jobs requiring the schedule environment

4.3.2.1 The Status (ST) Action Character

From the scheduling environment display, you can issue the st action character to display all jobs that require a particular scheduling environment.

An example of the st action character being issued against a scheduling environment is shown in Figure 50 on page 99.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF SCHEDULING ENVIRONMENT DISPLAY MAS SYSTEMS        LINE 1-9 (9)
COMMAND INPUT ===>                                       SCROLL ==
PREFIX=*  DEST=(ALL)  OWNER=*
NP   SCHEDULING-ENV   DESCRIPTION                     SYSTEMS
     BATCH_DFLT       Default Batch                   SC64
     BATCH_PROD       Production Batch                SC63
     CICS_DEV         Development CICS active         SC64
     CICS_PROD        Production CICS active          SC63
     NO_CICS_DEV      Development CICS inactive       SC63
st   NO_CICS_PROD     Production CICS inactive
     OVERNIGHT        Over Night (1800 - 0700)
     PRIMESHIFT       Prime Shift (0800 - 1700)       SC63,SC64
     SANDBOX          Test Sysplex: SC63 and SC64     SC63,SC64
```

Figure 50. ST Action Character Issued on the Schedule Environment Panel

In Figure 50, you can see that the st action character is being issued against scheduling environment NO_CICS_PROD. This displays the Status panel for all jobs that require scheduling environment NO_CICS_PROD. See example Figure 51.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF STATUS DISPLAY ALL CLASSES                       LINE 1-3 (3)
COMMAND INPUT ===>                                       SCROLL ==
PREFIX=*  DEST=(ALL)  OWNER=*
NP   JOBNAME  JOBID    OWNER    PRTY QUEUE      C  POS  SAFF  ASYS STAT
i    NORMAND1 JOB03638 NORMAND    10 EXECUTION  9     1
     HOUSEKP1 JOB03656 NORMAND    10 EXECUTION  9     2
     FERGMIC1 JOB03692 FERGMIC    10 EXECUTION  9     3
```

Figure 51. ST Panel Showing Jobs Requiring NO_CICS_PROD

In Figure 51, you can see the jobs that require scheduling environment NO_CICS_PROD. If you require further information for a particular job, you can issue the I action character against the job to display the job information pop-up. In Figure 51, you can see the I action character being issued against a job. The resulting job information pop-up can be seen in Figure 52 on page 100.

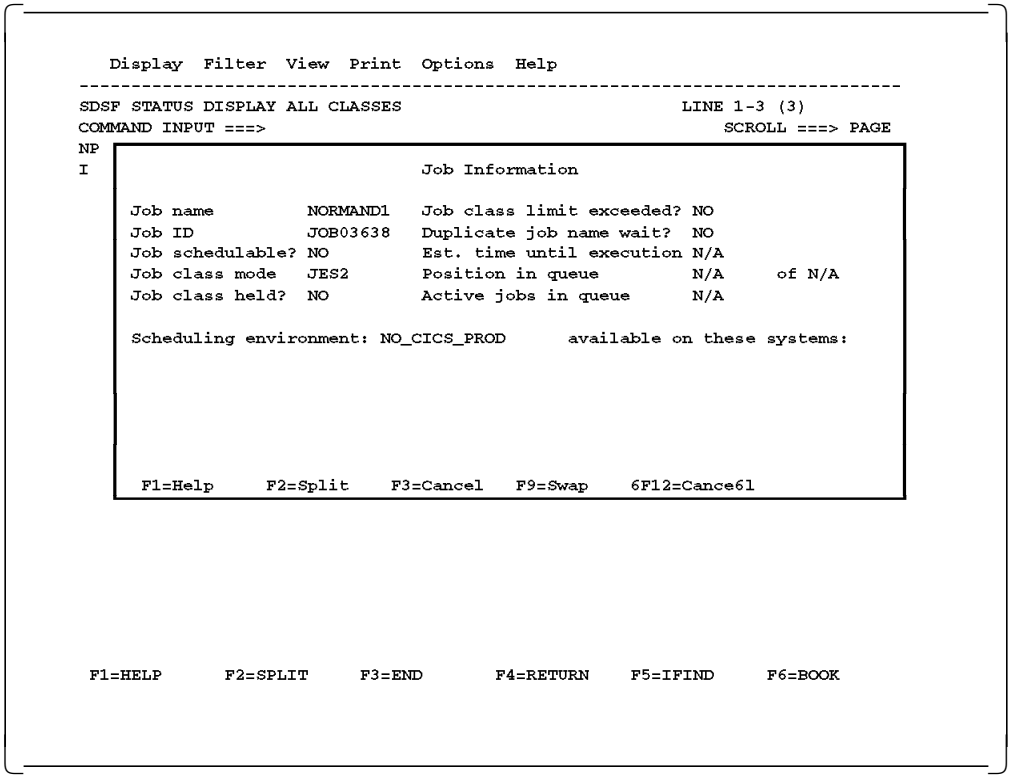

Figure 52. Job Information Pop-up

Example Figure 52 shows that the job requires scheduling environment NO_CICS_PROD. It also shows that NO_CICS_PROD is not available on any system.

4.3.2.2 The Display Resources (R) Action Character

From the scheduling environment display, you can issue the R action character to display the resources for a particular scheduling environment.

An example of the r action character being issued against a scheduling environment is shown in Figure 53.

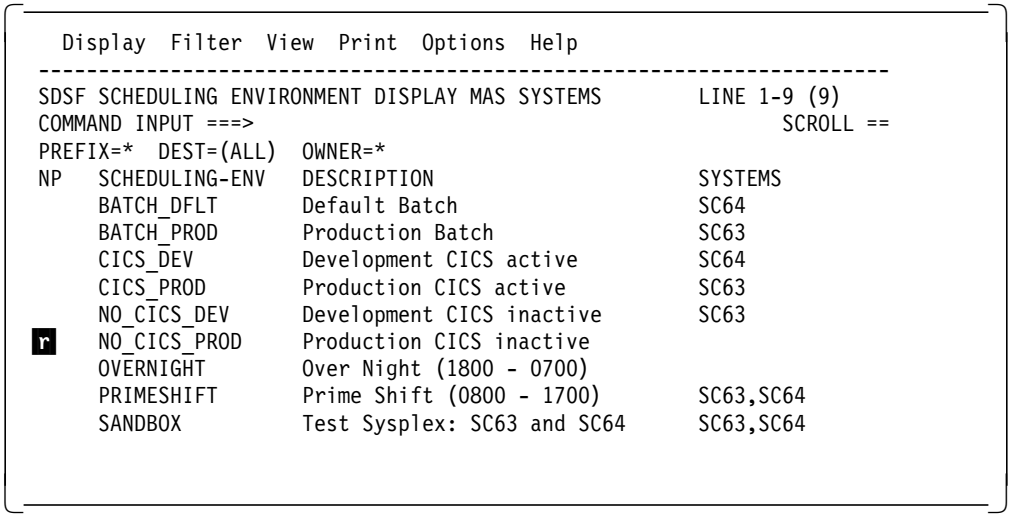

Figure 53. R Action Character Issued on the Scheduling Environment Panel

In Figure 53, you can see that the r action character is being issued against scheduling environment NO_CICS_PROD. This displays the resources defined for scheduling environment NO_CICS_PROD. See example Figure 54 on page 101.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS NO_CICS_PROD         LINE 1-3 (3)
COMMAND INPUT ===>                                          SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         REQSTATE SC63     SC64
     CICSPO1          OFF      ON       RESET
     CICSPO2          OFF      ON       RESET
```

Figure 54. Resource Panel

For a more detailed explanation of the resource display, see 4.4, “Resource Display.”

4.4 Resource Display

Resources are defined using the Workload Manager (WLM) dialog. A scheduling environment consists of one or more resources and a desired state for each resource, (ON or OFF). SDSF allows you to display the resources defined to the system via the Resource panel. The resource display shows the resources that have been defined through WLM and the state of each (ON, OFF or RESET). The Resource display can be accessed as a primary display, with a command or from a pull-down, in which case all the resources in the MAS (or optionally, the sysplex) are displayed. If the resource display is accessed using the r action character on the scheduling environment display, only the resources required by the scheduling environment are shown. You can modify the state of each resource on the system.

The resource display consists of a row for each resource. There is a column for the resource name and a column which shows the state of the resource for every system in the MAS (or optionally, sysplex). For an example of the resource display accessed using the RES command, see Figure 55 on page 102.

When the resource display is accessed via the r action character, against a particular scheduling environment, there is an extra column displayed which shows the required state of each resource for that scheduling environment. Only resources for that particular scheduling environment are displayed. For an example of the resource display accessed using the RES command, see Figure 56 on page 102.

Overtypes can be used to alter the state of the resources directly from the resource display.

4.4.1 Invoking the Resource Panel

The resource display can be accessed using the RES command or from the DISPLAY pull-down. It can also be invoked from the scheduling environment display using the R action character.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS                      LINE 1-11 (11)
COMMAND INPUT ===>                                        SCROLL ==
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         SC63     SC64
     BATCH_DFLT       OFF      ON
     BATCH_PROD       ON       OFF
     CICSDO1          OFF      ON
     CICSDO2          OFF      ON
     CICSDO3          OFF      ON
     CICSPO1          ON       RESET
     CICSPO2          ON       RESET
     OVERNIGHT        RESET    RESET
     PRIMESHIFT       ON       ON
     SC63             ON       ON
     SC64             ON       ON
```

Figure 55. Resource Panel, Accessed With the RES Command

Figure 55 shows the resource display accessed via the RES command. It shows all the resources that have been defined. There are two systems in the MAS (SC63 and SC64), and you can see the status of the resources in each system.

When the resource display is accessed with the r action character, against a particular scheduling environment, another column is also shown on the display. The *REQSTATE* column displays the status of each resource that the scheduling environment requires in order to be available.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS NO_CICS_PROD        LINE 1-3 (3)
COMMAND INPUT ===>                                       SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         REQSTATE SC63     SC64
     CICSPO1          OFF      ON       RESET
     CICSPO2          OFF      ON       RESET
```

Figure 56. Resource Panel, Accessed With the R Action Character

Figure 56 shows the resource display accessed using the r action character. In this example the r action character has been issued against scheduling environment NO_CICS_PROD. You can see the resources that are required by scheduling environment NO_CICS_PROD, the required state for each resource, and the actual state of each resource for each system in the MAS.

4.4.2 Overtypable Fields on Resource Panel

After an IPL of a system, all resources in the system are in RESET status. An operator command,

F WLM,RESOURCE=*resource_name***,ON|OFF|RESET**

must be used to set the resource to the desired state.

On the resource display, the system name field for each resource is overtypable. This allows the status of a resource on a particular system to be altered directly from SDSF, without having to use the MVS modify command.

Figure 57 shows a scheduling environment display, and you can see that none of the scheduling environments is available on any system. The r action character is being issued against scheduling environment CICS_DEV. This invokes the resource display, as shown in Figure 58.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF SCHEDULING ENVIRONMENT DISPLAY MAS SYSTEMS        LINE 1-9 (9)
COMMAND INPUT ===>                                       SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   SCHEDULING-ENV   DESCRIPTION                     SYSTEMS
     BATCH_DFLT       Default Batch
     BATCH_PROD       Production Batch
 r   CICS_DEV         Development CICS active
     CICS_PROD        Production CICS active
     NO_CICS_DEV      Development CICS inactive
     NO_CICS_PROD     Production CICS inactive
     OVERNIGHT        Over Night (1800 - 0700)
     PRIMESHIFT       Prime Shift (0800 - 1700)
     SANDBOX          Test Sysplex: SC63 and SC64
```

Figure 57. Scheduling Environment Panel

In Figure 58, you can see the resources that make up scheduling environment CICS_DEV: they are CICSD01, CICSD02, and CICSD03.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS CICS_DEV             LINE 1-3 (3)
COMMAND INPUT ===>                                       SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         REQSTATE SC63     SC64
     CICSD01          ON       RESET    RESET
     CICSD02          ON       RESET    RESET
     CICSD03          ON       RESET    RESET
```

Figure 58. Resource Panel Output

The required state for each resource is ON; therefore, to make scheduling environment CICS_DEV available on system SC63, you need to set the resources

to ON for that system. You can do this by simply overtyping the current state, as shown in Figure 59 on page 104.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS CICS_DEV            LINE 1-3 (3)
COMMAND INPUT ===>                                          SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         REQSTATE SC63     SC64
     CICSD01          ON       on ET    RESET
     CICSD02          ON       on ET    RESET
     CICSD03          ON       on ET    RESET
```

Figure 59. Resource Panel

Once you have pressed Enter, you see that the state of each of the resources that you overtyped has changed to ON; see Figure 60. The state of the resources on system SC63 now match the required state for scheduling environment CICS_DEV.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS CICS_DEV            LINE 1-3 (3)
COMMAND INPUT ===>                                          SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         REQSTATE SC63     SC64
     CICSD01          ON       ON       RESET
     CICSD02          ON       ON       RESET
     CICSD03          ON       ON       RESET
```

Figure 60. Resource Panel, Showing the Status of CICS_DEV Resources

If you return to the scheduling environment display, you see that CICS_DEV is now available on system SC63; see Figure 61.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF SCHEDULING ENVIRONMENT DISPLAY MAS SYSTEMS        LINE 1-9 (9)
COMMAND INPUT ===>                                          SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   SCHEDULING-ENV   DESCRIPTION                     SYSTEMS
     BATCH_DFLT       Default Batch
     BATCH_PROD       Production Batch
     CICS_DEV         Development CICS active         SC63
     CICS_PROD        Production CICS active
     NO_CICS_DEV      Development CICS inactive
     NO_CICS_PROD     Production CICS inactive
     OVERNIGHT        Over Night (1800 - 0700)
     PRIMESHIFT       Prime Shift (0800 - 1700)
     SANDBOX          Test Sysplex: SC63 and SC64
```

Figure 61. Scheduling Environment Panel. CICS_DEV is Available on System SC63.

4.4.3 Action Characters on the Resource Display

Table 5 shows the action characters available on the resource panel:

Table 5. Action Characters Available on the Resource Display

Action Character	Description
=	Repeat the last command
//	Repeat a block of commands
+	Temporarily extend the NP column
D	Display resources in the logs and ULOG (issues D command)

4.4.3.1 The Display (D) Action Character

From the resource display, you can issue the D action character to display a resource. The D action character issues the

D WLM,RESOURCE=*resource_name*

command, and the output is displayed in the logs and ULOG.

Figure 62 shows the d action character being issued on the resource display, against resource BATCH_PROD.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF RESOURCE DISPLAY MAS SYSTEMS                      LINE 1-12 (12)
COMMAND INPUT ===>                                        SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   RESOURCE         SC63     SC64
     BATCH_DFLT       RESET    RESET
 d   BATCH_PROD       RESET    RESET
     CICSDO1          OFF      ON
     CICSDO2          OFF      ON
     CICSDO3          OFF      ON
     CICSPO1          ON       RESET
     CICSPO2          ON       RESET
     CICSPO3          ON       RESET
     OVERNIGHT        RESET    RESET
     PRIMESHIFT       ON       ON
     SC63             ON       RESET
     SC64             RESET    ON
```

Figure 62. D Action Character Issued on the Resource Panel

Figure 63 on page 106 shows the output from the D action character.

```
SC63   97148  11:57:43.99          -D WLM,RESOURCE=BATCH_PROD
                                   IWM038I  11.57.43  WLM DISPLAY 489
                                     RESOURCE:     BATCH_PROD
                                      DESCRIPTION: Production Batch
                                      SYSTEM   STATE
                                      SC63     RESET
```

Figure 63. Output from the D Action Character Issued on the Resource Panel

4.5 Job Class Display

With OS/390 Release 4, you can have WLM-managed job classes as well as JES2 job classes.

The JES2 JOBCLASS statement now has a MODE keyword, which has a value of either JES or WLM. If you specify MODE(JES), the class will be managed by JES (JES will manage the batch initiators the same way it does in prior releases). If you specify MODE(WLM), the class is managed by WLM. Classes can be switched between JES-managed and WLM-managed modes.

SDSF has a Job Class panel that can be used to display and control both WLM and JES2 job classes. The display contains information on each class, such as counts of jobs waiting and in hold status. Overtypes and action characters can be used to control classes.

4.5.1 Invoking the Job Class Panel

Note: The Job Class display requires OS/390 R4 JES2 to be active in each member of the MAS.

The Job Class panel can be invoked by either the JC command or Display pull-down.

An example of a command to display the Job Class panel, showing all classes, is:

```
COMMAND INPUT ===> JC
```

This command displays the job class panel, displaying all classes, as shown in Figure 64 on page 107.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF JOB CLASS DISPLAY ALL CLASSES                    LINE 1-24 (32)
COMMAND INPUT ===>                                           SCROLL ===
NP   C STATUS   MODE WAIT-CNT XEQ-CNT  HOLD-CNT ODISP        QHOLD HOLD
     A NOTHELD  WLM
     B NOTHELD  WLM
     C HELD     JES         6                  (KEEP,KEEP    YES   NO
     D NOTHELD  JES         2       1         2 (KEEP,KEEP)  NO    NO
     E NOTHELD  JES                             (KEEP,KEEP)  NO    NO
     F NOTHELD  JES                 1           (PURGE,KEEP) NO    NO
     G NOTHELD  JES                           7 (KEEP,KEEP)  NO    YES
     H NOTHELD  JES                             (KEEP,KEEP)  NO    NO
     I HELD     JES         2                   (KEEP,KEEP)  YES   NO
     J NOTHELD  JES                             (PURGE,KEEP) NO    NO
     K NOTHELD  JES         4       1           (KEEP,KEEP)  NO    NO
     L NOTHELD  WLM
```

Figure 64. Job Class Panel

An example of a command to display the Job Class panel, showing only job classes A, B and C, is:

```
COMMAND INPUT ===> JCabc
```

This command displays the job class panel, displaying only classes A, B and C, as shown in Figure 65. Up to six classes can be entered on the JC command and only those classes are displayed.

Note: There is no blank between JC and *classes*.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF JOB CLASS DISPLAY CLASS ABC                      LINE 1-3 (3)
COMMAND INPUT ===>                                           SCROLL ===
NP   C STATUS   MODE WAIT-CNT XEQ-CNT  HOLD-CNT ODISP        QHOLD HOLD
     A NOTHELD  WLM
     B NOTHELD  WLM
     C HELD     JES         6                  (KEEP,KEEP)  YES   NO
```

Figure 65. Job Class Panel, Limited to Classes A, B and C

4.5.2 Action Characters on the Job Class Display

Table 6 shows the action characters available on the job class panel:

Table 6 (Page 1 of 2). Action Characters Available on the Job Class Display

Action Character	Description
=	Repeat the last command
//	Repeat a block of commands
+	Temporarily extend the NP column
A	Release a job class (issues $A JES2 command)

Table 6 (Page 2 of 2). Action Characters Available on the Job Class Display	
Action Character	**Description**
D	Display information about a job class (issues $D JES2 command)
H	Hold a job class (issues $H JES2 command)
ST	Display the ST panel for all jobs in the class

4.5.2.1 The Status (ST) Action Character

From the job class display you can issue the ST action character to display all jobs that are assigned to a particular job class.

An example of the ST action character being issued against a job class is shown in Figure 66.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF JOB CLASS DISPLAY CLASS ABC                        LINE 1-3 (3)
COMMAND INPUT ===>                                        SCROLL ==
NP   C STATUS   MODE WAIT-CNT XEQ-CNT  HOLD-CNT ODISP         QHOLD
     A NOTHELD  WLM
     B NOTHELD  WLM
st   C HELD     JES         2                  (KEEP,KEEP)     YES
```

Figure 66. ST action character Issued on the JC Panel

Figure 67 shows the status panel. You can see the two jobs that are waiting execution in job class C.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF STATUS DISPLAY CLASS C                          LINE 1-2 (2)
COMMAND INPUT ===>                                     SCROLL ===
NP   JOBNAME  JOBID    OWNER    PRTY QUEUE      C  POS  SAFF  ASYS STATU
     NORMAND1 JOB02467 NORMAND     5 EXECUTION  C     1
     FERGMIC1 JOB02879 FERGMIC     5 EXECUTION  C     2
```

Figure 67. Status Panel, Showing the Jobs Assigned to Job Class C

4.5.3 Display-Only Fields on the Job Class Panel

There are six display-only fields on the job class panel. For a brief description of the display only fields on the job class panel, see Table 7.

Table 7 (Page 1 of 2). Display-only Fields of the Job Class Panel	
Field Title	**Description**
C	Job class
STATUS	Job class status

Table 7 (Page 2 of 2). Display-only Fields of the Job Class Panel	
Field Title	**Description**
MODE	Manager of class (JES2 or WLM)
WAIT-CNT	Number of jobs waiting
XEQ-CNT	Number of jobs executing
HOLD-CNT	Number of jobs on hold

4.5.4 Overtypable Fields on the Job Class Panel

The Job Class panel has various overtypable fields to allow an operator or systems programmer to control the classes easily, without having to remember complex JES2 commands. For a brief description of the overtypable fields available on the job class panel, see Table 8.

Table 8 (Page 1 of 2). Overtypable Fields of the Job Class Panel	
Field Title	**Description**
ODISP	Output disposition for normal and abnormal end of jobs.
QHOLD	Specifies whether the queue (class) is to be held until the class is released from held status ($TJOBCLASS,QHELD=NO or $AQ command).
HOLD	Specifies whether jobs in this class are to be held until they are released from held status.
XBM	Specifies the procedure name JES2 uses as the target of an EXEC statement.
JCLIM	Job class limit for the system.
PGN	Default performance group number for this job class.
AUTH	MVS operator commands that are to be executed.
BLP	Bypass label processing.
COMMAND	Disposition of commands read from the input stream.
LOG	Specifies whether the job log is to be printed for this job class.
MSGLV	Message level value.
OUT	Specifies whether SYSOUT data is to be printed for jobs executed in this class.
PL	Default procedure library number to be used for this job class.
REGION	Default region size assigned to each jobstep executed in this class.
SWA	Placement of eligible schedule work area (SWA) control blocks created for jobs in this class.
MAX-TIME	Default for the maximum time that each job step may run.
ACCT	Specifies whether an account numbar is required on the JCL JOB statement.
CPY	Specifies whether jobs are to be queued for output processing as though TYPRUN=COPY were specified on the JCL JOB statement.
JRNL	Specifies whether job-related information is to be saved in a job journal.
PGNM	Specifies whether a programmer name is required on a JCL JOB statement.

Table 8 (Page 2 of 2). Overtypable Fields of the Job Class Panel	
Field Title	**Description**
RST	Specifies whether JES2 is to requeue for execution any job in this class that had been executing before an IPL of the system and a JES2 warmstart.
SCN	Specifies whether jobs in this job class are to be queued for output immediately after JCL conversion, as though TYPRUN=SCAN were specified on the JCL JOB statement.
UJP	Specifies whether the IEFUJP exit is to be taken when job is purged.
USO	Specifies whether the IEFUSO exit is to be taken when the SYSOUT limit is reached for a job in this job class.
TP6	Specifies whether JES2 is to produce type 6 SMF records for this job class.
TP26	Specifies whether JES2 is to produce type 26 SMF records for jobs in this job class.
CPRG	Specifies whether system data sets (such as JESMSG and SYSMSG) in this time-sharing-user class (JOBCLASS(TSU)) are conditionally purged.
MC	Message class for all time-sharing sessions. This class specification is the default logon message class for all TSO/E logons.

4.5.5 Using Overtypable Fields

The overtypable fields on the job class panel make it easier for an operator or systems programmer to manage job classes without having to remember complex JES2 commands.

Figure 68 shows a job class display. Job class C is in HELD status and has six jobs waiting to execute.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF JOB CLASS DISPLAY CLASS ABC                       LINE 1-3 (3)
COMMAND INPUT ===>                                        SCROLL ===
NP   C STATUS   MODE WAIT-CNT XEQ-CNT  HOLD-CNT ODISP         QHOLD HOLD
     A NOTHELD  WLM
     B NOTHELD  WLM
     C HELD     JES         6                  (KEEP,KEEP)   YES   NO
```

Figure 68. Job Class Panel, Showing JES2 Job Class C in HELD Status

To display the jobs that are assigned to job class C, you can use the ST action character. You issue the ST action character against the job class that you are interested in. Figure 69 on page 111 shows ST being issued against job class C.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF JOB CLASS DISPLAY CLASS ABC                        LINE 1-3 (3)
COMMAND INPUT ===>                                             SCROLL ===
NP   C STATUS   MODE WAIT-CNT XEQ-CNT  HOLD-CNT ODISP         QHOLD
     A NOTHELD  WLM
     B NOTHELD  WLM
st   C HELD     JES         6                  (KEEP,KEEP)     YES
```

Figure 69. Job Class Panel, Showing ST Issued against Job Class C

Figure 70 shows the status panel. You can see the six jobs that are waiting for execution in job class C. To display job specific information, you can use the I line command, which causes the job information pop-up to be displayed.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF STATUS DISPLAY CLASS C                             LINE 1-6 (6)
COMMAND INPUT ===>                                             SCROLL ===
NP   JOBNAME  JOBID    OWNER    PRTY QUEUE      C  POS  SAFF  ASYS STATU
i    NORMAND1 JOB03487 NORMAND     5 EXECUTION  C     1
     EVEBYE1  JOB03498 EVEBYE      5 EXECUTION  C     2
     FERGMIC1 JOB03512 FERGMIC     5 EXECUTION  C     3
     DUMP001  JOB03523 NORMAND     5 EXECUTION  C     4
     NORMAND2 JOB03536 NORMAND     5 EXECUTION  C     5
     NORMAND3 JOB03539 NORMAND     5 EXECUTION  C     6
```

Figure 70. Status Panel, Showing Six Jobs Waiting to Execute in Job Class C

Figure 71 on page 112 shows the job information pop-up. The job information pop-up confirms that the job class is held.

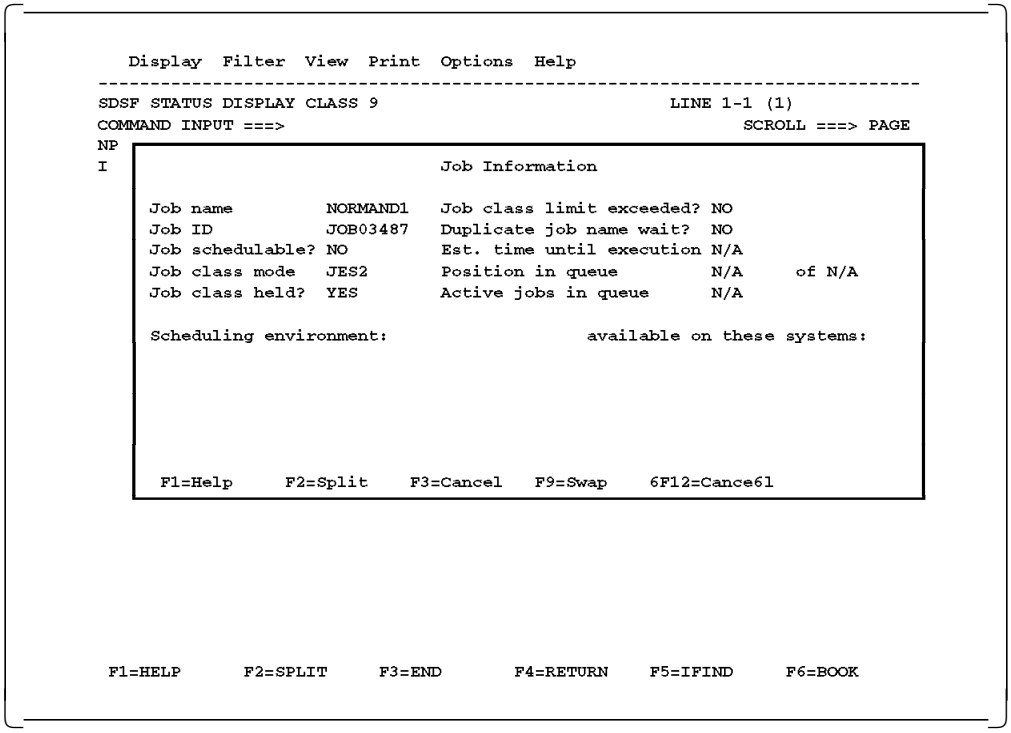

Figure 71. Job Information Pop-Up

If you return to the job class display, you can overtype the QHOLD field for job class C, as shown in Figure 72. In this example, the QHOLD field for JES2 job class C is being altered from YES (job class held) to NO (job class not held).

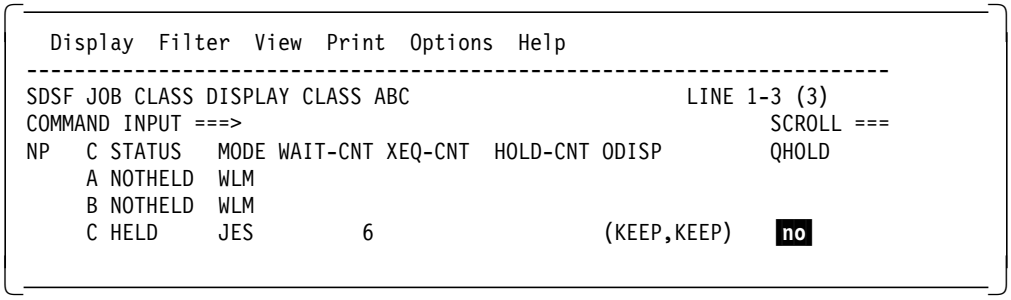

Figure 72. Job Class Panel, Showing the QHOLD Field Being Overtyped

In Figure 73 on page 113, you can see that the staus of job class C has changed from HELD to NOTHELD. The display shows that one job is now executing and five are waiting for execution.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF JOB CLASS DISPLAY CLASS ABC                        LINE 1-3 (3)
COMMAND INPUT ===>                                             SCROLL ===
NP   C STATUS   MODE WAIT-CNT XEQ-CNT  HOLD-CNT ODISP         QHOLD
     A NOTHELD  WLM
     B NOTHELD  WLM
     C NOTHELD  JES         5       1           (KEEP,KEEP)   NO
```

Figure 73. Job Class Panel, Showing the Job Class C Status of NOTHELD

4.6 WLM-Managed Batch Support

With OS/390 release 4, WLM can manage batch initiators that serve WLM-managed (that is, MODE(WLM)) classes.

WLM can dynamically manage the number of initiator address spaces, starting and stopping initiators as required, to meet performance goals.

WLM initiator address spaces are not shown on the INIT display; they are shown on the DA display.

For further information on WLM-managed batch, see 2.4, “Batch Initiator Management” on page 21.

Operators have less control of WLM-managed initiators than they have of JES-managed initiators; WLM stops and starts initiator address spaces as required. OS/390 JES2 provides three new commands to provide the operator with the ability to stop and start the scheduling of jobs on a member of a MAS, and to force WLM to provide an initiator address space to start a job. SDSF provides action characters to allow operators to issue the new JES2 commands directly from SDSF.

4.6.1 New Action Characters for WLM-Managed Batch

There are two new action characters for use on the MAS display.

4.6.1.1 The SX Action Character

The SX action character is issued on the MAS display against a member of the MAS. It generates the $SXEQ JES2 command. The SX action character starts the scheduling of jobs on the member it is issued against.

4.6.1.2 The PX Action Character

The PX action character is issued on the MAS display against a member of the MAS. It generates the $PXEQ JES2 command. The PX action character stops the scheduling of jobs on the member it is issued against, and both JES2 and WLM initiators are drained and terminated.

4.6.1.3 The J Action Character

The J action character is issued on the I and ST displays against a job. It generates the $SJ JES2 command. In response to the J action character, WLM immediately supplies an initiator address space to execute the job. The example shown in Figure 74 shows the syslog entries a job that has been released via the J action character. The $SJ JES2 command is issued and WLM starts an initiator address space for the job.

```
$SJ3995
$HASP890 JOB(NORMAND5) 905
$HASP890 JOB(NORMAND5)  STATUS=(AWAITING EXECUTION),
$HASP890                CLASS=W,PRIORITY=5,SYSAFF=(ANY),
$HASP890                HOLD=(NONE)
IWM034I PROCEDURE INIT STARTED FOR SUBSYSTEM JES2 906
APPLICATION ENVIRONMENT SYSBATCH
PARAMETERS SUB=MSTR
IEF196I         1 //INIT     JOB MSGLEVEL=1
IEF196I         2 //STARTING EXEC INIT
.
.
ICH70001I NORMAND  LAST ACCESS AT 17:54:58 ON WEDNESDAY, MAY 28, 1997
$HASP373 NORMAND5 STARTED - WLM INIT  - SRVCLASS          - SYS SC63
```

Figure 74. A $SJ Command Generated by the SDSF J Action Character

In Figure 74, you can see the $SJ command being issued. WLM immediately starts an initiator address space and the job starts.

4.6.2 New Fields for WLM Managed Batch

OS/390 Release 4 JES2 provides a new SSI for job information. Fields that are new with OS/390 Release 4 JES2 are only available via the new SSI. SDSF exploits the new SSI to display the new fields.

Note: This support requires OS/390 Release JES2 to be active.

SDSF provides job-level information about jobs in WLM-managed classes by adding to existing displays:

- SRVCLASS, which identifies the service class of a job - I and ST panels
- DLY, which indicates that a job is delayed - I and ST panels
- EST-WAIT, which indicates WLM′s estimated time until processing of the job begins - I and ST panels.
- MODE, which indicates whether a job is managed by JES2 or WLM - I and ST panels
- SCHEDULING-ENV, which shows the name of the scheduling environment associated with a job - I and ST panels
- SRVCLASS, which identifies the service class of a job - I and ST panels
- WPOS, which identifies the position of the job in the WLM queue - DA, I, and ST panels.

The SRVCLASS field is overtypable. It generates the JES2 command:

```
$T J,SRVCLASS=
```

4.7 Job Information Pop-up

The job information allows a user to see at a glance various things that might be delaying the processing of a job. The information in the job information pop-up is supplied by JES2. The information displayed on the Job Information pop-up is:

- Job name
- Job ID
- Job schedulable (YES or NO)
- Job class mode (JES or WLM)
- Job class held (YES or NO)
- Job class limit exceeded (YES or NO)
- Duplicate job name (YES or NO)
- Estimated time until execution (hh:mm:ss) (WLM managed classes only)
- Position in queue (WLM managed classes only)
- Active jobs in queue (WLM managed classes only)
- Scheduling environment
- Systems on which sceduling environment is available

4.7.1 Invoking the Job Information Pop-up

The job information pop-up is invoked by the I action character on the I or ST SDSF panels. Figure 75 shows the I action character being issued against a job on the ST panel.

```
  Display  Filter  View  Print  Options  Help
------------------------------------------------------------------------
SDSF STATUS DISPLAY ALL CLASSES                       LINE 1-23 (790)
COMMAND INPUT ===>                                       SCROLL ==
PREFIX=*  DEST=(ALL)  OWNER=*
NP   JOBNAME  JOBID    OWNER    PRTY QUEUE      WPOS  SCHEDULING-ENV
     NORMAND5 JOB03995 NORMAND     5 EXECUTION      0
     NORMAND1 JOB04241 NORMAND     5 EXECUTION      0
     NORMAND1 JOB04249 NORMAND     5 EXECUTION      0
i    RONNSCHA JOB03587 RONN       10 EXECUTION      0 NO_CICS_PROD
     NORMAND1 JOB03638 NORMAND    10 EXECUTION      0 OVERNIGHT
     HOUSEKP1 JOB03639 NORMAND    10 EXECUTION      0 OVERNIGHT
```

Figure 75. Status display, Showing the I Action Character Being Issued

The job information pop-up resulting from this action is shown in Figure 76 on page 116.

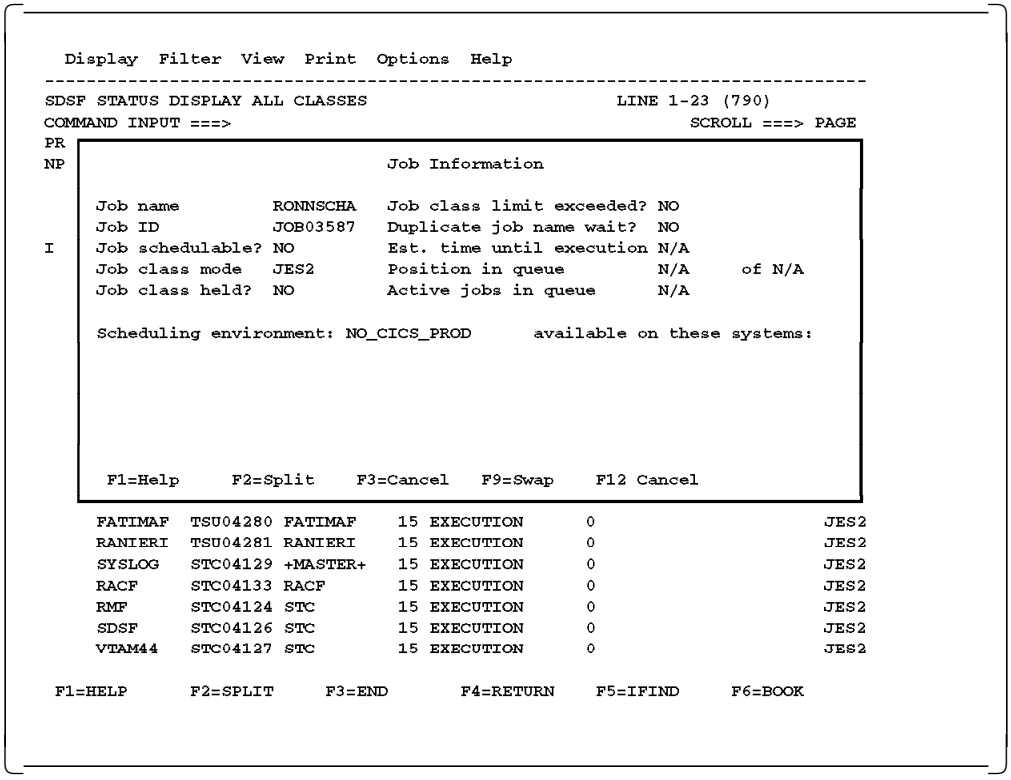

Figure 76. Job Information Pop-Up, Showing a Job in JES2-Managed Class

In Figure 76, you can see that this job is in a JES-managed class. This job is not schedulable, because it requires scheduling environment NO_CICS_PROD, and NO_CICS_PROD is not available on any system. Fields that are not applicable to this job (that is, fields only applicable to WLM-managed classes), are shown as N/A.

An example of a schedulable job in a JES-managed class is shown in Figure 77 on page 117.

```
   Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF STATUS DISPLAY ALL CLASSES                          LINE 1-23 (799)
COMMAND INPUT ===>                                         SCROLL ===> PAGE
PR
NP                              Job Information

     Job name          NORMAND1   Job class limit exceeded? NO
     Job ID            JOB03968   Duplicate job name wait?  NO
     Job schedulable? YES         Est. time until execution N/A
I    Job class mode    JES2       Position in queue         0       of N/A
     Job class held?   NO         Active jobs in queue      N/A

     Scheduling environment:                  available on these systems:




      F1=Help     F2=Split    F3=Cancel   F9=Swap    F12=Cancel
```

Figure 77. Job Information Pop-Up, Showing a Schedulable Job

An example of a job that is in a WLM-managed class and is not schedulable is shown in Figure 78.

```
   Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF STATUS DISPLAY ALL CLASSES                          LINE 1-23 (802)
COMMAND INPUT ===>                                         SCROLL ===> PAGE
PR
NP                              Job Information
I
     Job name          NORMAND2   Job class limit exceeded? NO
     Job ID            JOB04311   Duplicate job name wait?  NO
     Job schedulable? NO          Est. time until execution N/A
     Job class mode    WLM        Position in queue         N/A     of N/A
     Job class held?   NO         Active jobs in queue      N/A

     Scheduling environment:                  available on these systems:




      F1=Help     F2=Split    F3=Cancel   F9=Swap    F12=Cancel
```

Figure 78. Job Information Pop-Up, Showing a Job in WLM-Managed Class

An example of a schedulable job in a WLM-managed class is shown in Figure 79 on page 118

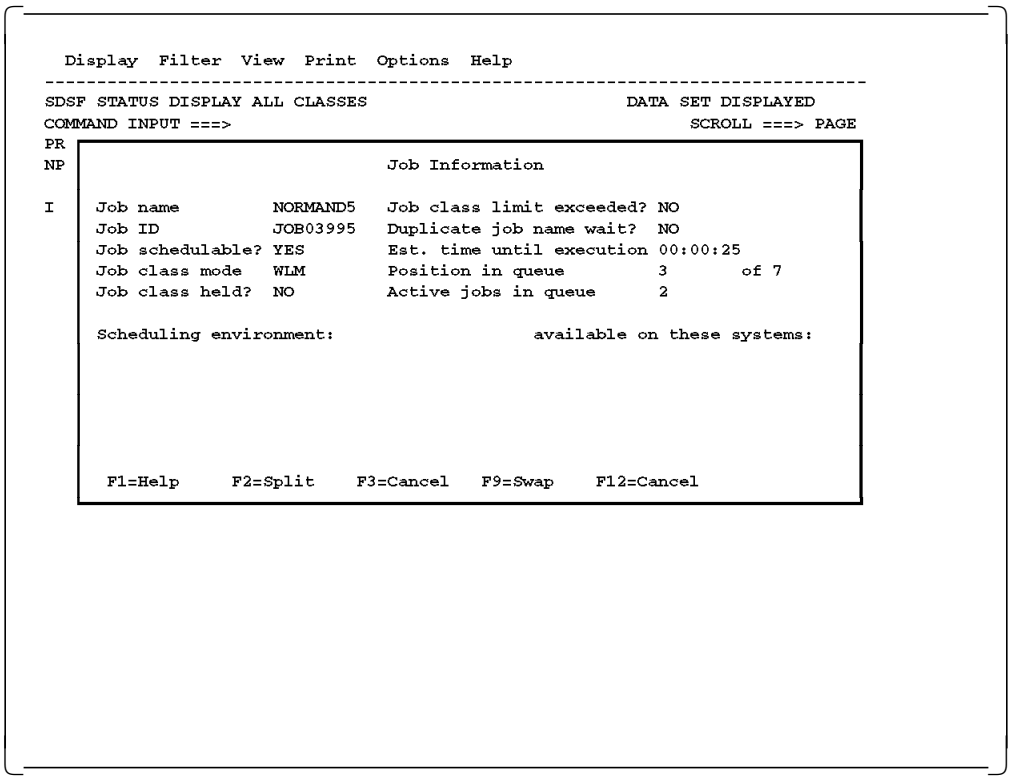

Figure 79. Job Information Pop-Up, Showing a Schedulable Job

4.8 Maximum Return Code

OS/390 release 4 SDSF has been enhanced to take advantage of the new SSI available with OS/390 release 4 JES2. The highest return code returned for a job is provided by a new JES2 SSI. SDSF displays this maximum return code on the H, O, and ST panels under the title *MAX-RC*. If there was an abend, the field provides the appropriate information. The MAX-RC field contains one of the values shown in Table 9.

Table 9. Possible Values of the MAX-RC Field

MAX-RC	**Description**
blank	No completion information
CC xxxx	Job ended normally
CC xxxx	Job ended by CC
JCL ERROR	JCL error
CANCELED	Job canceled
CONV ABEND	Converter abended
SEC ERROR	Security error
ABEND Uxxx	Job abended (user abend)
ABEND Sxxx	Job abended (system abend)

The MAX-RC field has been added to the end of the primary and alternate field lists, but can be arranged forward. An ISFFLD macro, or an FLD statement along with FLDENT statements, defines the fields that are displayed on an SDSF panel. You can define a primary and alternate field list for each SDSF panel. The primary field list contains the fields that are shown upon entry into a panel. The alternate field list contains fields that can be displayed by use of the ? command (or the View pull-down). A user can also use the ARRANGE command (or View pull-down) to change the order of the fields on an SDSF display.

In the example shown in this section, the MAX-RC field has been moved forward in the O display by using the the GROUP, FLD, and FLDENT statements.

Figure 80 shows the GROUP, FLD, and FLDENT statements required to move the MAX-RC forward. The example is for OFIELD2, the alternate O display. Users that have TSO authority of JCL, OPER, and ACCT are able to display the alternate fields using the ? command.

```
   /*****************************************/
   /* GROUP ISFSPROG - System Programmers */
   /*****************************************/

GROUP NAME(ISFSPROG),        /* Group name                          */
TSOAUTH(JCL,OPER,ACCT),      /* User must have JCL, OPER, ACCT      */
CMDLEV=7,
OFIELD2(OFLD2)               /* Alternate field list for O display  */

   /*****************************************/
   /* Alternate field list for OUT display */
   /*****************************************/

FLD NAME(OFLD2) TYPE(OUT)

  FLDENT COLUMN(JOBID),TITLE(JOBID),WIDTH(D)
  FLDENT COLUMN(OWNERID),TITLE(OWNER),WIDTH(D)
  FLDENT COLUMN(RETCODE),TITLE('MAX-RC'),WIDTH(D)
  FLDENT COLUMN(JPRIO),TITLE(JP),WIDTH(D)
  FLDENT COLUMN(DPRIO),TITLE(PRTY),WIDTH(D)
  FLDENT COLUMN(OCLASS),TITLE(C),WIDTH(D)
  FLDENT COLUMN(FORMS),TITLE(FORM),WIDTH(D)
  FLDENT COLUMN(FCBID),TITLE(FCB),WIDTH(D)
  .
  .
  .
```

Figure 80. GROUP, FLD and FLDENT Statements

Figure 81 on page 120 shows examples of the ISFGRP and ISFFLD macros that can be used to move MAX-RC forward. The example is for OFIELD2, the alternate O display. Users that have TSO authority of JCL, OPER, and ACCT are able to display the alternate fields.

```
*********************************************************************
*                                                                   *
*        SAMPLE ISFGRP MACRO FOR SYSTEM PROGRAMMERS.                *
*                                                                   *
*********************************************************************
         SPACE 5
ISFSPROG ISFGRP TSOAUTH=(JCL,OPER,ACCT),                            -
               AUTH=(LOG,I,O,H,DA,DEST,PREF,SYSID,ABEND,ACTION,     -
               INPUT,FINDLIM,ST,INIT,PR,TRACE,ULOG,MAS,SYSNAME,     -
               LI,NO,SO,PUN,RDR,SE,RES,JC),                         -
               OFIELD2=OFLD2

*********************************************************************
*                                                                   *
*        THE ISFFLD MACRO CAN CHANGE THE DEFAULT                    *
*        FIELD LISTS FOR THE SDSF PANELS.  THIS ISFFLD DEFINES THE  *
*        FIELD LIST FOR THE ALTERNATE O PANEL FOR THE               *
*        SYSTEM PROGRAMMERS GROUP ABOVE.                            *
*                                                                   *
*********************************************************************
         SPACE 5
OFLD2    ISFFLD JOBID,JOBID,8,OWNERID,OWNER,8,                      -
               RETCODE,MAX-RC,10 ,JPRIO,JP,2,                       -
               DPRIO,PRTY,4,OCLASS,C,1,                             -
               FORMS,FORMS,8,FCBID,FCB,4,                           -
               .                                                    -
               .                                                    -
               .                                                    -
               TYPE=OUT
```

Figure 81. ISFGRP and ISFFLD Macros

Figure 82 shows the MAX-RC displayed on the Output Queues (O) display.

```
  Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------
SDSF OUTPUT ALL CLASSES ALL FORMS     LINES 1,314      LINE 1-11 (11)
COMMAND INPUT ===>                                          SCROLL ===
PREFIX=*  DEST=(ALL)  OWNER=*
NP   JOBNAME  JOBID    OWNER    MAX-RC     JP PRTY C FORM     FCB  UCS
     NORMAND1 JOB03883 NORMAND  CC 0000     1  144 A STD      **** ****
     NORMAND2 JOB03953 NORMAND  CC 0004     1  144 A STD      **** ****
     FERGMICA JOB03954 FERGMIC  ABEND U999  1  144 A STD      **** ****
     NORMAND2 JOB03602 NORMAND  CANCELED    1  144 H STD      **** ****
     ESILVAB  JOB03603 ESILVA   CC 0000     1  144 H STD      **** ****
     NORMAND1 JOB03523 NORMAND  ABEND S13E  1  144 X STD      **** ****
     EVEBYE1  JOB03562 EVEBYE   ABEND S806  1  144 X STD      **** ****
     NORMAND2 JOB03673 NORMAND  JCL ERROR   1  144 X STD      **** ****
     NORMAND5 JOB03924 NORMAND  SEC ERROR   1  144 X STD      **** ****
     FERGMIC  TSU03815 FERGMIC  CC 0000     1  144 B STD      **** ****
     NORMAND  TSU04012 NORMAND  CC 0000     1  144 B STD      **** ****
```

Figure 82. Output Queues Display, Showing MAX-RC Field

4.9 Securing the New Functions of SDSF Using ISFPARMS

ISFPARMS defines SDSF global options, the functions each user is authorized to perform, and panel formats. ISFPARMS is required, but you can use SAF to control user authorization.

There are two formats of ISFPARMS:

- Assembler macros that you define, assemble, and then link into the SDSF load library.
- Statements, which are simpler to code and can be updated without reassembling and link-editing.

OS/390 Release 4 SDSF adds parameters to ISFPARMS for the new enhancements.

Within ISFPARMS, an ISFGRP macro or GROUP statement defines a group of users and specifies which functions the members of the group may perform.

The AUTH parameter of the ISFGRP macro or GROUP statement identifies the SDSF panels that members of this user group are allowed to display, and the SDSF commands that they are allowed to issue. The SE, RES, and JC commands must be coded on the AUTH parameter to allow the members of the user group to issue the commands and view the panels.

Table 10 describes the authorized commands that have been added for OS/390 Release 4 SDSF.

Table 10. Authorized Commands New to OS/390 Release 4 SDSF

Command	**Description**	**AUTH=**
SE	Display Scheduling Environment panel	SE
RES	Display Resource panel	RES
JC	Display Job Class panel	JC

An example of the ISFGRP macro that gives members of a user group the authority to issue several commands, including SE, RES, and JC is:

```
ISFGRP TSOAUTH(JCL,OPER),       /* User must have JCL, OPER   */
       AUTH(LOG,I,O,H,DA,       /* Authorized functions       */
            SE,RES,JC)          /*                            */
```

In the preceding example, TSOAUTH defines a group of users that have TSO authority of JCL and OPER. The AUTH parameter identifies the commands and panels that members of this group are allowed to issue and display.

An example of the GROUP statement that gives members of a user group the authority to issue several commands, including SE, RES, and JC is:

```
GROUP NAME(ISFSPROG),       /* Group name                      */
  TSOAUTH(JCL,OPER,ACCT),   /* User must have JCL, OPER, ACCT  */
  AUTH(LOG,I,O,H,DA,        /* Authorized functions            */
       INIT,PR,SE,RES,JC)   /*                                 */
```

In the preceding example, TSOAUTH defines a group of users that have TSO authority of JCL, OPER and ACCT. The AUTH parameter identifies the

commands and panels that members of this group are allowed to issue and display.

The CMDLEV parameter of the ISFGRP macro or GROUP statement specifies the command level to which a group member is authorized. The command level determines the action characters that a group member can issue, the fields that a group member can overtype, and the MVS and JES2 commands that a group member can issue from the command line.

The CMDLEV parameter works with the CMDAUTH, ICMD and XCMD parameters. The CMDAUTH, ICMD and XCMD parameters indicate the jobs, TSO user IDs, started tasks, initiators, printers and punches for which members of this group can:

- Issue action characters and overtype fields
- Issue MVS and JES2 commands

For a full explanation of CMDLEV, CMDAUTH, ICMD and XCMD parameters, see *OS/390 SDSF Customization and Security*.

The default command level, if it is not specified on the ISFGRP macro or GROUP statement, is 0. Each command level is inclusive of all those with a lower number. For example, a user with command level of 2 can perform the functions requiring command level 2, 1, and 0.

Chapter 5. Dynamic LPA

OS/390 release 4 introduces the ability to dynamically add or delete modules from the Link Pack Area (LPA) after the IPL. This new facility, the Dynamic Link Pack Area (DLPA), allows optional and new products to be loaded into the system without an IPL. Release 4 also enables you to display modules residing in LPA.

All modules added to LPA are authorized. As part of its LPA search, the system will find modules added dynamically. A module added dynamically will be found before one of the same name added during the IPL. Modules added dynamically to LPA will be loaded into the Common Area (CSA OR ECSA).

OS/390 R4 enables that LPA be buildable from PDSEs. While PDSEs are not available early enough during the IPL process to be processed in the building of LPA, modules in PDSEs can be added to LPA using dynamic LPA services.

5.1 Dynamic LPA Operation Functions

You can perform the following operations through dynamic LPA functions:

- Add a module/alias (or modules/aliases) to LPA.
- Delete dynamically added modules or aliases from LPA.
- Set minimum CSA/ECSA thresholds.

The dynamic LPA functions may be invoked through one of the following functions:

- The PROGXX parmlib member includes the LPA statements, which is used to define modules that can be added to or deleted from LPA after the IPL.
- The SETPROG LPA command, to initiate a change (add,delete) to LPA.
- The Display PROG,LPA command may be used to display information about modules that have been added to LPA.
- The new CSVDYLPA macro allows an authorized program to initiate dynamic LPA services.

5.2 Syntax Format of the LPA Statements

You can use the LPA statement to specify:

- Modules that are to be added to the LPA following IPL.
- Modules that are to be deleted from the LPA following IPL.
- Threshhold values for minimum amounts of CSA storage that must still be available after an ADD operation.

Syntax Format For LPA Statements

Syntax Format For LPA ADD

```
LPA ADD
    MODNAME(modname,...) | mask(mask)
    DSNAME(dsname | LNKLST)
    (FIXED|PAGEABLE)
    (PAGEPROTPAGE)
```

Syntax Format For LPA DELETE

```
LPA DELETE
    MODNAME(modname,...)
    FORCE(YES)
    (CURRENT|OLDEST)
```

Syntax Format For LPA CSAMIN

```
LPA CSAMIN
    (below,above)
```

LPA Statement type indicating that an action may be performed on LPA.

Note: LPA ADD and LPA DELETE cannot be used during IPL. They are for use in PROGxx members pointed to by SET PROG=XX after IPL.

Where:

ADD Specifies that one or more modules is to be added to LPA.

Default Value: *None*

DELETE Specifies that one or more modules is to be delete from LPA. Only modules is added to LPA after an IPL are eligible for dynamic deletion.

Default Value: *None*

CSAMIN Specifies the minimum amount of CSA and ECSA that must remain after a module is added to LPA. If the requested ADD operation would reduce the CSA or ECSA below the defined minimum, the system rejects the operation.

Default Value: *(0,0)*

below The minimum amount of below-16M CSA storage that must remain after the ADD operation, expressed in the format (n|nK|nM), where n is a decimal number, n is n*1024, and nM is n*1024*1024.

above The minimum amount of above-16M CSA storage that must remain after the ADD operation, expressed in the format (n|nK|nM), where n is a decimal number, n is n*1024, and nM is n*1024*1024.

MODNAME MODNAME=(modname,...,modname) refers to the 1-8 character LPA module name or alias. If the last character of the modname is "*"&,. it will be treated as X'C0'. This lets you directly specify the name of a load module that ends with that nonprintable character. Wildcard characters are not supported within modname. A maximum of 128 module names can be provided.

Note: MOD and MODULE can be used as synonyms of MODNAME.

Default Value:: If MODNAME is not specified, MASK must be specified.

MASK MASK(mask) refers to the 1-8 character mask that is to be applied to all the members of the specified data set. It can contain wildcard characters "*" and "?," and all members that match will be processed.

Default Value: If MASK is not specified, MODNAME must be specified.

DSNAME DSNAME(dsname) refers to the 1-44 character data set name that contains the module(s) or alias(es). If MODNAME is specified, you can specify DSNAME(LNKLST) or DSNAME(LINKLIST) if you want the system to use its normal search protocol (search LPA then search the lnklst) instead of a particular data set. The data set must be cataloged.

The attribute of the CSA for each module is assigned as OWNER=SYSTEM. DSN, LIB, and LIBRARY can be used as synonyms of DSNAME.

Default Value: None

FIXED|PAGEABLE This indicates whether the modules are to be placed in fixed or pageable storage.

PAGE can be used as a synonym of PAGEABLE.

Default Value: PAGEABLE

PAGEPROTPAGE This indicates whether or not to page protect the modules entirely. The default is to page protect the entire module. Be aware that when that default is taken, storage utilization for the modules increases, as each module gets allocated a number of whole pages (so that they can be page protected), rather than just the amount of storage that is truly necessary to load the module.

When PAGEPROTPAGE is requested, however, only the whole pages within each load module are page protected, which keeps the storage use to the minimum amount but which makes it possible that a storage overlay of the beginning or end of the load module can occur.

PPPAGE and PPP can be used as synonyms of PAGEPROTPAGE.

Default Value: Page protect entire modules.

FORCE(YES) This confirms that the delete requestor understands the ramifications of deleting a module from LPA, when the system can have no knowledge of whether any code is currently executing within the specified module.

Default Value: None. This is a required parameter.

CURRENT|OLDEST CURRENT specifies that the current copy is to be deleted. OLDEST specifies that the oldest copy other than the current one is to be deleted. CUR can be used as a synonym of CURRENT. OLD can be used as a synonym of OLDEST.

Default Value: CURRENT

5.3 DLPA Module Management

The Dynamic LPA service can be initiated after the IPL by:

- Using the LPA statement in the PROGXX parmlib member

Use the LPA statement in PROGxx to specify:

- Modules that are to be added to the LPA following an IPL
- Modules that are to be deleted from the LPA following an IPL
- ThreshHold values for minimum amounts of CSA storage that must still be available after an ADD operation

5.3.1 Using the LPA Statement in PROGxx Parmlib Member

Following is an example of the ADD LPA parameter required in the PROGXX parmlib member to add a module dynamically to LPA.

LPA ADD PROGxx Definition

```
LPA ADD MODNAME(IEFACTRT) DSNAME(SYS1.SANDBOX.LPALIB) PPALL
```

The output appears in the following format:

```
T PROG=JS
IEE252I MEMBER PROGJS   FOUND IN SYS1.PARMLIB
CSV551I 22.24.06 LPA ADD 792
SUCCESSFUL: 1  UNSUCCESSFUL: 0  NOT PROCESSED: 0
MODULE    RESULT
IEFACTRT  SUCCESSFUL
IEE536I PROG     VALUE JS NOW IN EFFECT
```

Following is an example of the DELETE LPA parameter required in the PROGXX parmlib member to delete a module dynamically added to LPA.

LPA DELETE PROGxx Definition

```
LPA DELETE MODNAME(IEFACTRT) FORCE(YES) CURRENT
```

The output appears in the following format:

```
      T PROG=DJ
      IEE252I MEMBER PROGDJ   FOUND IN SYS1.PARMLIB
      CSV551I 22.24.30 LPA DELETE 796
      SUCCESSFUL: 1  UNSUCCESSFUL: 0  NOT PROCESSED: 0
      MODULE    RESULT
```

Following is an example of the CSAMIN LPA parameter required in the PROGXX parmlib member to set the minimum amounts of CSA storage that must still be available after a ADD to LPA operation.

Set CSAMIN PROGxx Definition

```
LPA CSAMIN(10K,100K)
```

The output appears in the following format:

```
T PROG=JC
IEE252I MEMBER PROGJC   FOUND IN SYS1.PARMLIB
CSV554I LPA CSAMIN HAS BEEN SET TO (10K,100K)
IEE536I PROG     VALUE JC NOW IN EFFECT
```

Once the PROGXX member has been created with either the ADD/DELETE/CSAMIN statement, a SET PROG=XX operator command must be issued to enable the change.

Note: The CSAMIN LPA statement is the only LPA statement allowed at IPL time. If you do not set CSAMIN, the default value of (0,0) will allow all available CSA/ECSA storage to be exhausted by the dynamic addition of modules.

5.4 SETPROG Command Syntax

Use the SETPROG LPA command any time after IPL to specify:

- Modules that are to be added to the LPA
- Modules that are to be deleted from the LPA
- Threshhold values for minimum amounts of CSA storage that still must be available after an ADD operation

SETPROG Command Syntax

```
SETPROG LPA,{ADD,MODNAME=(modname...,modname) | MASK=mask}
           {    ,DSNAME=(dsname | LNKLST)     }
           {    [,FIXED] [,PAGEPROTPAGE]          }
           {                                }
           {DELETE,MODNAME=(modname...,modname)      }
           {    ,FORCE=YES [,CURRENT | ,OLDEST]   }
           {                                }
           {CSAMIN=(below,above)                 }
```

5.4.1 Statements/Parameters for the SETPROG LPA Command

SETPROG LPA. Statement type indicating that an action may be performed on LPA.

Where:

ADD Specifies that one or more modules is to be added to LPA.

Default Value: *None*

DELETE Specifies that one or more modules is to be deleted from LPA. Only modules added to LPA after an IPL are eligible for dynamic deletion.

Default Value: *None*

CSAMIN Specifies the minimum amount of CSA and ECSA that must remain after a module is added to LPA. If the requested ADD operation would reduce the CSA or ECSA below the defined minimum, the system rejects the operation.

Default Value: *(0,0)*

below The minimum amount of below-16M CSA storage that must remain after the ADD operation. The value can be expressed as *n*, *nK*, and *nM*.

above The minimum amount of above-16M CSA storage that must remain after the ADD operation. The value can be expressed as *n*, *nK*, and *nM*.

MODNAME=(modname,...) The term modname refers to the 1-8 character LPA module name or alias. If the last character of the modname is "*," it is treated as X′C0′. This lets you directly specify the name of a load module that ends with that nonprintable character. Wildcard characters are not supported within modname. A maximum of 128 module names can be provided. MOD and MODULE can be used as synonyms of MODNAME.

Default Value: *If MODNAME is not specified, MASK must be specified.*

MASK(mask) The term mask refers to the 1-8 character mask that is to be applied to all the members of the specified data set. It can contain wildcard characters "*" and "?," and all members that match will be processed.

Default Value: *If MASK is not specified, MODNAME must be specified.*

DSNAME(dsname) The term dsname refers to the 1-44 character data set name that contains the module(s) or alias(es). If MODNAME is specified, you can specify DSNAME(LNKLST) or DSNAME(LINKLIST) if you want the system to use its normal search protocol (search LPA, then search the lnklst) instead of a particular data set. The data set must be cataloged.

The attribute of the CSA for each module is assigned as OWNER=SYSTEM. DSN, LIB, and LIBRARY can be used as synonyms of DSNAME.

Default Value: *None*

FIXED This indicates that the modules are to be placed in fixed storage.

Default Value: *If FIXED is not specified, the modules will be placed in pageable storage*

PAGEPROTPAGE This indicates whether or not to page protect the modules entirely. The default is to page protect the entire module. Be aware that when that default is taken, storage utilization for the modules increases because each module gets allocated a number of whole pages (so that they can be page protected), rather than just the amount of storage that is truly necessary to load the module.

When PAGEPROTPAGE is requested, however, only the whole pages within each load module are page protected, which keeps the storage use to the minimum amount but which makes it possible that a storage overlay of the beginning or end of the load module can occur.

PPPAGE and PPP can be used as synonyms of PAGEPROTPAGE.

Default Value: *Page protect entire modules.*

FORCE(YES) This confirms that the delete requestor understands the ramifications of deleting a module from LPA, when the system can have no knowledge of whether any code is currently executing within the specified module.

Default Value: *None. This is a required parameter*

CURRENT|OLDEST CURRENT specifies that the current copy is to be deleted. OLDEST specifies that the oldest copy other than the current one is to be deleted. CUR can be used as a synonym of CURRENT. OLD can be used as a synonym of OLDEST.

Default Value: *CURRENT*

5.4.2 Managing Dynamic LPA Content (SETPROG LPA)

The Dynamic LPA service can be initiated after the IPL by:

- Using the SETPROG LPA system command

Use the SETPROG command to specify:

- Modules that are to be added to the LPA following an IPL
- Modules that are to be deleted from the LPA following an IPL
- ThreshHold values for minimum amounts of CSA storage that must still be available after an ADD operation

Following is an example of the SETPROG LPA,ADD command to add a module dynamically to LPA through console services.

SETPROG LPA,ADD Command

```
SETPROG LPA,ADD,MODNAME(IEFACTRT),DSNAME(SYS1.SANDBOX.LPALIB),PPALL
```

The output appears in the following format:

```
SETPROG LPA,ADD,MODNAME=(IEFACTRT),DSNAME(SYS1.SANDBOX.LPALIB),PPP
IEF196I IEF237I 0CC1 ALLOCATED TO SYS00004
IEF196I IEF285I   SYS1.SANDBOX.LPALIB                          KEPT
IEF196I IEF285I   VOL SER NOS= SBOX01.
CSV551I 17.20.13 LPA ADD 311
SUCCESSFUL: 1  UNSUCCESSFUL: 0  NOT PROCESSED: 0
MODULE    RESULT
IEFACTRT  SUCCESSFUL
```

Follwing is an example of the SETPROG LPA,ADD command using the MASK parameter to dynamically ADD multiple modules to LPA.

SETPROG LPA,ADD,MASK Command

```
SETPROG LPA,ADD,MASK(ICE*),DSNAME(SYS1.SORTLPA),PPALL
```

The output appears in the following format:

```
SETPROG LPA,ADD MASK=ICE*,DSNAME=SYS1.SORTLPA,PPALL
IEF196I IEF237I 2051 ALLOCATED TO SYS00005
IEF196I IEF285I   SYS1.SORTLPA
IEF196I IEF285I   VOL SER NOS= 034RSB.
CSV551I 11.34.55 LPA ADD 622
SUCCESSFUL: 31  UNSUCCESSFUL: 0  NOT PROCESSED: 0
MODULE    RESULT
ICEADEX   SUCCESSFUL
ICEBLDX   SUCCESSFUL
ICECBIF   SUCCESSFUL
ICECOB2   SUCCESSFUL
ICECPUT   SUCCESSFUL
ICEDSVA   SUCCESSFUL
ICEDSVO   SUCCESSFUL
ICEEXIT   SUCCESSFUL
ICEGENER  SUCCESSFUL
ICEIEFS   SUCCESSFUL
ICEIPDT   SUCCESSFUL
ICEIPFT   SUCCESSFUL
ICEIPUT   SUCCESSFUL
ICEIPVT   SUCCESSFUL
ICEMAN    SUCCESSFUL
ICEMAN2   SUCCESSFUL
ICEMCDU   SUCCESSFUL
ICEMESA   SUCCESSFUL
ICEMESO   SUCCESSFUL
ICENLS    SUCCESSFUL
ICEOUTF   SUCCESSFUL
.
.
```

Following is an example of the SETPROG LPA,DELETE command to delete a module that has been added dynamically to LPA with the ADD command.

SETPROG LPA,DELETE Command

```
SETPROG LPA,DELETE,MODNAME(IEFACTRT),FORCE(YES),CURRENT
```

The output appears in the following format:

```
SETPROG LPA,DELETE,MODNAME=(IEFACTRT),FORCE=YES,CURRENT
CSV551I 17.21.41 LPA DELETE 313
SUCCESSFUL: 1  UNSUCCESSFUL: 0  NOT PROCESSED: 0
MODULE    RESULT
IEFACTRT  SUCCESSFUL
```

The system protects modules that have been loaded at IPL. If an attempt is made to delete a module that has not been dynamically loaded, the system will prevent the deletion from being successfully processed.

Following is an example of an attempt to delete a module that is flagged as being loaded at IPL time.

```
D PROG,LPA,MOD=IEFACTRT
CSV550I 17.25.55 LPA DISPLAY 366
FLAGS  MODULE    ENTRY PT  LOAD PT   LENGTH    DIAG
   P   IEFACTRT  823BF000  023BF000  00001238  07372E00

SETPROG LPA,DELETE,MOD=IEFACTRT,FORCE=YES,CURRENT
CSV551I 17.28.09 LPA DELETE 370
SUCCESSFUL: 0  UNSUCCESSFUL: 1  NOT PROCESSED: 0
MODULE    RESULT
IEFACTRT  NOT SUCCESSFUL. NOT IN DYNAMIC LPA
```

The system will display, through console services, that the module was not deleted because it was not found in dynamic LPA.

Following is an example of the SETPROG CSAMIN,LPA command to set the minimum amounts of CSA storage that must still be available after a module is added to LPA.

SETPROG LPA,CSAMIN Command

```
SETPROG LPA,CSAMIN(100K,200K)
```

The output appears in the following format:

```
 SETPROG LPA,CSAMIN=(100K,200K)
 CSV554I LPA CSAMIN HAS BEEN SET TO (100K,200K)
```

Once the PROGXX member has been created with either the ADD/DELETE/CSAMIN statement, a SET PROG=XX operator command must be issued to enable the change.

5.5 DISPLAY Command Syntax for LPA

The complete syntax for the DISPLAY PROG,LPA command is:

Display PROG,LPA Syntax

```
D  PROG,LPA(,MODNAME=modname)
           (,CSAMIN    )

(,L={a|cc|cca|name|name-a})
```

Where:

MODNAME= This displays the entry point and load point/length information fo the LPA module. MOD and MODULE can be used as synonyms of MODNAME .

modname This is the 1-8 character LPA module name. If the last character of the modname is "*," it will be treated as X′C0′

CSAMIN This displays the current CSA and ECSA minimum values.

L = a, cc, cca, name, or name-a

This specifies where the system is to present the display:

The console identifier (cc)
The display area (a)

Both the console identifier and the display area (cca
The console name (name)
Both the console name and the display area (name-a).

For cc, you must specify a decimal number from 1 to 9

Use the DISPLAY PROG,LPA command to display entry point and load point/length information about modules currently in LPA, and to display the minimum amount of CSA and ECSA that must remain after a module is added to LPA.

Following is an example of the DISPLAY PROG,LPA,CSAMIN command to display the current CSA and ECSA minimum values.

```
D PROG,LPA,CSAMIN
CSV557I LPA CSAMIN VALUE IS (0,0)
```

Following is an example of the DISPLAY PROG,LPA,MODNAME command to display the LPA module entry point and load point/length information.

```
D PROG,LPA,MOD=IEFACTRT
CSV550I 22.23.49 LPA DISPLAY 786
FLAGS   MODULE    ENTRY PT  LOAD PT   LENGTH    DIAG
  P     IEFACTRT  822BF000  022BF000  00001238  07272E00
```

Note: This displays the module that has been added at IPL time. The P below FLAGS indicates that the module was added at IPL time.

Following is an example of the DISPLAY PROG,LPA,MODNAME command to display the current CSA and ECSA minimum values.

```
D PROG,LPA,MODNAME=IEFACTRT
CSV550I 18.29.30 LPA DISPLAY 863
FLAGS   MODULE    ENTRY PT  LOAD PT   LENGTH    DIAG
 D P    IEFACTRT  86B5A000  06B5A000  00001238  00F99FB0
```

Note: This displays the module that has been added dynamically. The D P below FLAGS indicates that the module was added at IPL time, and has subsequently been added dynamically to LPA.

5.6 CSVDYLPA Macro

CSVDYLPA allows you to to request dynamic LPA services. Be aware, however, that changes to LPA itself are not actually done. This set of services truly only lets you add modules to, and delete modules from, common storage. When searching by module name, the system will locate the copy of a module added by dynamic LPA services even if it was present in PLPA, MLPA, or FLPA.

With CSVDYLPA, you can request services for the following operations:

- Add one or more modules to common storage (REQUEST=ADD)
- Delete one or more modules that were previously added using dynamic LPA services (REQUEST=DELETE)
- Query information about support for LPA services (REQUEST=QUERYDYN)

5.7 SLIP/PER

SLIP needs to interact with dynamic LPA with respect to its LPAMOD and LPAEP parameters. Currently, the name specified by LPAMOD and LPAEP must exist within LPA.

In order to enable SLIP traps to be set for dynamic LPA modules prior to their being present in LPA, LPAMOD and LPAEP checking is deferred. The processing still checks for their existence, but if the entry is not yet defined, an informational message is displayed and processing continues (as opposed to entering an error state).

Once the specified module is added to LPA, SLIP processing will detect this and update its control structures with the addresses that apply. If this is the currently active SLIP trap, SLIP will also update the control registers with the now-known information about the module's address range.

When a dynamically-added module is deleted, it is also necessary for SLIP to update its control structures. It is possible at that point that there is no remaining LPA module of the right name. It is also possible that an older LPA module of the right name exists and needs to become the one on which the information is based.

5.8 Restrictions

There is no MASK parameter associated with the delete function. The only way to delete modules added with the MASK parameter is to issue the command for each module that has been loaded.

The display function allows you to display individual modules only. You are not able to list all modules that have been dynamically added to LPA since the IPL process.

Chapter 6. System Logger

System logger is an MVS component that allows an application to log data from a sysplex environment. Data can be logged from one system or from multiple systems across the sysplex.

Examples of applications using system logger are:

- CICS log manager.

  You can find more information about this in the following manuals:

  – *CICS Transaction Server for OS/390 Release Guide*
  – *CICS Installation Guide*
  – *CICS System Definition Guide*

- Operations log stream (OPERLOG). You can find more information about this in the following manual:

  – *OS/390 MVS Planning: Operations*

In addition to the existing coupling facility log streams, OS/390 Release gives you the ability to use DASD-only log streams also.

6.1 DASD-Only Log Streams

Prior to OS/390 Release 4, you could only define log streams with an associated coupling facility list structure. Therefore, you had to configure at least one coupling facility (or, from an availability standpoint, it is better to configure two) into a sysplex. You had to define list structures into the coupling facility resource manager policy (CFRM policy) to use log streams.

But there are also a large number of users who do not require a sysplex consisting of multiple physical systems and physical coupling facilities. For those customers, DASD-only log streams are an alternative.

In a DASD-only log stream, the log data is contained in local storage buffers for the system, and then offloaded to DASD log data sets.

Figure 83 on page 138 shows a DASD-only log stream configuration.

A DASD-only log stream has the following scope:

- A single-system scope
- Only one system can connect to it or write it
- Multiple applications from the same system can connect to it and write it

When a system logger application writes a log block to a DASD-only log stream, system logger writes it first to the local storage buffers for the system and duplexes it to a DASD staging data set associated with the system. When the staging data set space allocated for the log stream reaches the installation-defined threshold, system logger offloads the log blocks from local storage buffers to VSAM linear DASD data sets. Both a DASD-only log stream and a coupling facility log stream can have data in multiple DASD log data sets; as a log stream fills log data sets on DASD, system logger automatically allocates new ones for the log stream.

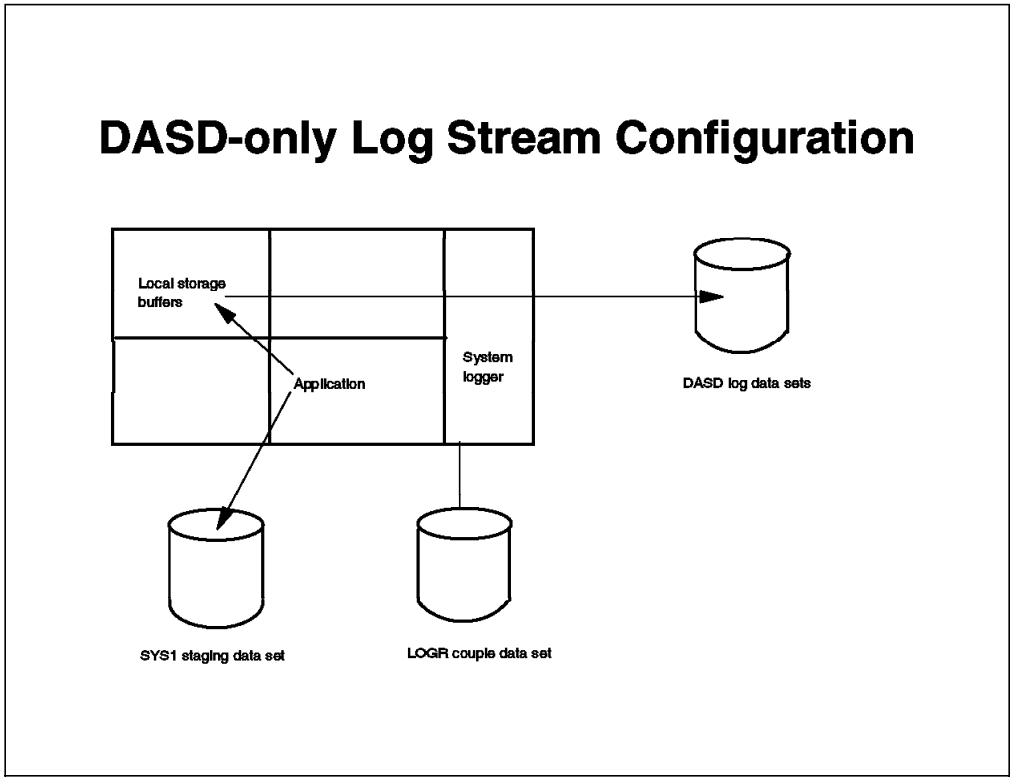

Figure 83. A DASD-Only Log Stream Configuration

6.1.1 Migration to DASD-Only Log Streams

If you plan to migrate to DASD-only log streams, you have to take a closer look at the following aspects:

- LOGR couple data set level

  There is no reformat needed if a LOGR couple data set already exists at a minimum MVS SP Version 5 Release 2.0 level or at an OS/390 level below Release 4.

  The systems where DASD-only log streams are planned to be used *must* be at the OS/390 Release 4 level; only with that release can DASD-only log streams be recognized and processed.

- Definition of DASD-only log streams

  Before you define DASD-only log streams to your installation you should insure that subsystems and applications which make use of the DASD-only log streams are prepared to do so. That means structure names and other related fields are to be set to binary zeroes for certain events.

- Considerations in sysplex configuration

  To use system logger applications in any way (that means coupling facility list structure or DASD-only log streams), your system environment must be in multisystem or monoplex mode.

  You must specify either:

  ```
  PLEXCFG=MONOPLEX
  ```

  or

  ```
  PLEXCFG=MULTISYSTEM
  ```

as a parameter in IEASYSxx member of SYS1.PARMLIB.

MONOPLEX means that the system is a single system sysplex that uses a sysplex couple data set. MULTISYSTEM means that the system is part of a sysplex consisting of one or more MVS systems that reside on one or more processors. The same sysplex couple data sets must be used by all systems.

Other options, as specified in *OS/390 V2R4.0 MVS Setting Up a Sysplex*, are not allowed in this case.

- Coexistence with older OS/390 Releases

  If you are using mixed release system levels within your sysplex. then you must apply a compatibility PTF, OW26722, for those systems which are at a level before OS/390 Release 4. This PTF insures that the down-level systems do not try to handle DASD-only log streams.

  Note: If you do not apply the PTF and your systems are at a mixed level, the MVS system logger may not operate correctly.

6.1.2 Administrative Tasks

If you plan to use DASD-only logs streams within your sysplex environment, you might have to do some administrative tasks. These tasks, which depend on the situations you may have at your installation, and are as follows:

- Defining a DASD-only log stream
- Deleting a DASD-only log stream
- Updating a DASD-only log stream to a structure log stream
- Connecting to a DASD-only log stream
- Writing to a DASD-only log stream
- Browsing from a DASD-only log stream
- Deleting data from a DASD-only log stream

These tasks and the related work you might have to do are described in the following topics.

6.1.2.1 Defining a DASD-only Log Stream

If you define DASD-only log streams, you must use macro IXGINVNT with the following new parameter:

```
IXGINVNT REQUEST=DEFINE,
   DASDONLY=YES,
      :
      :
```

You must also use the utility program IXCMIAPU, which resides in SYS1.MIGLIB. This utility program is used to define, update, list, or delete policy data on the log stream to a sysplex. In Figure 84 on page 140, it is used to define a DASD-only log stream.

```
//DEFINE JOB
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSIN DD *

   DATA TYPE(LOGR) REPORT(YES)
      DEFINE STRUCTURE NAME(LOGSTRUCTURE)
             LOGSNUM(1)
             AVGBUFSIZE(4096)
             MAXBUFSIZE(4096)

      DEFINE LOGSTREAM NAME(PLEXTEST.OPERLOG)
             STRUCTNAME(LOGSTRUCTURE)
             DASDONLY(YES)
             MAXBUFSIZE(40000)
/*
```

Figure 84. Job to Define a DASD-Only Log Stream

In Figure 84 you can see the following new parameters on a DEFINE LOGSTREAM request:

DASDONLY(NO|YES) This is an optional parameter that indicates whether DASD-only log streams are to be used or not.

- **(NO)** Indicates that the log stream will use the coupling facility list structure.
- **(YES)** Indicates that the log stream will not use the coupling facility list structure. Instead, it will instead use the DASD-only log streams.

The default value is NO.

MAXBUFSIZE(maxbufsize) Specifies the size of the largest block that can be written to this (in maxbufsize) DASD-only log stream. The values for MAXBUFSIZE must be between 1 and 65532 bytes. The default is 65532 bytes.

6.1.2.2 Deleting a DASD-only Log Stream

There are no changes to a previous release of OS/390. If you request a deletion, a job to execute looks as shown in Figure 85.

```
//DELETE JOB
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSIN DD *

   DATA TYPE(LOGR) REPORT(YES)
       DELETE LOGSTREAM NAME(PLEXTEST.OPERLOG)
/*
```

Figure 85. Job to Delete a DASD-Only Log Stream

6.1.2.3 Upgrading DASD-Only Log Streams to Structure Log Streams

If you upgrade DASD-only log streams to structure log streams in the coupling facility, you must use the macro IXGINVNT with the following parameter:

```
IXGINVNT REQUEST=UPDATE,
   STRUCTNAME=LOGSTRUCTURE,
      :
      :
```

Then use IXCMIAPU, as shown in Figure 86.

```
//UPDATE JOB
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSIN DD *

   DATA TYPE(LOGR) REPORT(YES)
      UPDATE LOGSTREAM NAME(PLEXTEST.OPERLOG)
             STRUCTNAME(LOGSTRUCTURE)
/*
```

Figure 86. Job to Upgrade a DASD-Only Log Stream to a Structure Log Stream

Note: LOGSTRUCTURE here is used as an example value of parameter STRUCTNAME. As you can see in Figure 86, there is a new parameter now allowed on an UPDATE LOGSTREAM request:

STRUCTNAME(structname) Specifies input of up to 16 characters that contains the name of the structure where all blocks of the log stream are to be written before they are written to DASD. This parameter can only be specified when the existing log stream which is to be modified is a DASD-only log stream.

With the specification of this parameter, the DASD-only log stream can be upgraded to use a coupling facility and become a coupling facility log stream.

If you upgrade a DASD-only log stream to a coupling facility log stream, you have to consider the following points:

- The structure specified on the STRUCTNAME parameter must be defined in a structure definition in the LOGR policy and in the CFRM policy couple data set.
- All connections (active or failed) to the DASD-only log stream must be disconnected.

For further information about set-up steps, see the system logger chapter in *OS/390 MVS Setting Up a Sysplex*.

Note: You cannot change a coupling facility log stream to a DASD-only log stream.

6.1.2.4 Connecting and Disconnecting to a DASD-Only Log Stream

First remember that only one system can be connected to a DASD-only log stream at a time. Multiple applications from the same system can connect to a DASD-only log stream. This single system scope is enforced by a sysplex-wide ENQ. The ENQ is held until a connection from a system to a DASD-only log stream disconnects.

The macro IXGCONN in SYS1.MACLIB must be used to connect or disconnect from a log stream. An application must issue REQUEST=CONNECT before it can read, write, or delete data in a log stream.

If you use macro IXGCONN with coupling facility log stream, the keywords STRUCTNAME, AVGBUFSIZE, and ELEMENTSIZE are filled with alphanumerical characters. When using DASD-only log streams, these keywords keep hexadecimal zeros as return values.

In the system logger processing of disconnection, you can get 3 different types of disconnection. Let us take a short look at these types, because 2 of them handle the processing in a different way when using DASD-only log streams.

- Disconnection for an application:

  An application issues macro IXGCONN with REQUEST=DISCONNECT to disconnect from a log stream. After that, system logger rejects any new requests from that application. There is no difference between coupling facility log streams and DASD-only log streams.

- Last disconnection for log stream on a system:

  System logger rejects any new requests for a specific log stream after processing the last disconnect for that log stream. For DASD-only log streams, the log data in local storage buffers (see also Figure 83 on page 138) is written to DASD log data sets.

- Last disconnection for a system in the sysplex:

  System logger writes all log data from a connection to DASD log stream data sets. For DASD-only log streams, the log data in local storage buffers (see also Figure 83 on page 138) is written to DASD log data sets.

6.1.2.5 Writing to a DASD-Only Log Stream

In the case, there are no changes in execution compared to the previous releases of OS/390. If you request a write, you must execute macro IXGWRITE from SYS1.MACLIB. The IXGWRITE macro allows a log user to write a log block to a log stream.

IXGWRITE first writes the data to local storage buffers for the system associated with the log stream and simultaneously duplexes it to DASD staging data sets. The safe import point for a DASD-only log stream is changed for every successful write request coming in.

For a DASD-only log stream, the safe import point is the highest log block identifier that has been successfully written to the log stream.

With DASD-only log streams, it is not a problem for you to consider data committed after issuing an IXGWRITE service because log data is duplexed automatically to a staging data set. The chance of a failure resulting in data loss is minimal, because there is no space of time between a successful write and the write of the data to DASD. There is also no problem if the coupling facility

fails. An application can use the IXGQUERY service macro to ensure that log data is safe to import from a DASD-only log stream. This macro is used to retrieve information about a log stream.

6.1.2.6 Browsing from a DASD-only Log Stream

There are no changes in executing compared to previous release of OS/390. If you request a write, you must execute macro IXGBRWSE from SYS1.MACLIB. The IXGBRWSE macro browses log blocks in a log stream.

6.1.2.7 Deleting Data from a DASD-Only Log Stream

There are no changes in execution compared to the previous releases of OS/390. If you request a write, you must execute macro IXGDELET from SYS1.MACLIB. The IXGDELET macro specifies a range of log blocks that are to be deleted.

For DASD-only log streams, the group of blocks which you specify for deletion can reside on both the local storage buffers and the DASD log data sets.

6.1.3 Staging Data Sets in DASD-Only Log Streams Environments

If you plan to use DASD-only log streams, you should review the following topics regarding the use of staging data sets:

- 6.1.3.1, “Planning for Staging Data Sets.”
- 6.1.3.2, “Staging Data Set Name Change.”

6.1.3.1 Planning for Staging Data Sets

In a DASD-only log stream environment, staging data sets are required if you use system logger. Staging data sets should be SMS-managed. If you do not plan on using SMS for the staging data sets, system logger automatically allocates them, using the maximum VSAM data set size. You can specify your own size for the data set using the STG_SIZE parameter in the log stream definition in the LOGR policy.

6.1.3.2 Staging Data Set Name Change

With the introduction of DASD-only log streams, naming conventions for staging data sets have changed.

To form the staging data set name during the first connection from a system to a DASD-only log stream, system logger uses the sysplex name as the last level qualifier. The sysplex name is the SYSPLEX parameter in member COUPLExx of SYS1.PARMLIB. In general this looks like:

```
hlq.logstreamname.sysplex
```

For example, a staging data set on system SC52 in sysplex WTSCPLX1 with high level qualifier IXGLOGR (which is the default) and associated with log stream LOG1.OPERLOG.RECOVERY.WTSCPLX1 has the following name:

```
IXGLOGR.LOG1.OPERLOG.RECOVERY.WTSCPLX1
```

In the past, system logger used the system name specified in IEASYSxx member of SYS1.PARMLIB as last level qualifier of a staging data set.

6.1.4 Application Failure and Recovery

An application which is connected to a DASD-only log stream may fail. This may happen independent of the system logger itself.

Using a DASD-only log stream, the processing that occurs after the application fails is the same as the processing which is performed for a failing application using a coupling facility log stream.

Let us assume there are no other connections to the log stream. The following actions happen:

- All log data that was written is offloaded to log stream DASD data sets.
- The application associated staging data set is freed.
- The LOGR inventory is updated to indicate a clean disconnection from the log stream.

6.1.4.1 System Level Recovery

While a system has connections to DASD-only log streams, the MVS system logger address space or the system itself may fail. System level recovery of the DASD-only log streams will not occur after the failing system gets a new IPL or after the system logger address space is restarted, because data has already been stored on DASD staging data sets. An explicit connection is required to the DASD-only log stream for it to be recovered (see also 6.1.4.2, “Log Stream Recovery”). Subsequent connections to the DASD-only log stream may be done from the same system where the failure occurred, or from another system within the sysplex, after the system logger address space is restarted.

6.1.4.2 Log Stream Recovery

Recovery of DASD-only log streams needs to be performed in the following cases:

- If the system logger is unable to successfully offload all of its written log block data to log stream DASD data sets during the last disconnect for the log stream
- If the system logger address space (or the system itself) fails while a system has connections to DASD-only log streams

Recovery of DASD-only log streams is be performed in the following cases:

- If an explicit connection is required to the DASD-only log stream for it to be recovered
- If the new first connection to the DASD-only log stream that had been in a failed state will cause that log stream to be recovered

As part of the connect processing, system logger reads log data from the staging data set into the local storage buffers of the connecting system. That log data is associated with the last system that connected to the log stream.

Note: OS/390 systems which are pre-Release 4 do not perform any recovery for DASD-only log streams.

Chapter 7. OS/390 Version 2 Release 4 JES3

To provide continuous availability, an installation must configure their MVS system such that no planned or unplanned outage will cause all hardware and/or software to fail at any one time. The MVS system and the JES3 component do not yet quite reach this goal. Before OS/390 Version 2 Release 4 JES3, any change to the JES3 operating definitions caused an installation to bring down and restart the JES3 complex (and therefore the sysplex as well). Such changes included adding hardware, installing new software releases and some maintenance, and changing the JES3-managed configuration. Outages caused by these processes are reduced by OS/390 Version 2 Release 4 JES3, if not eliminated.

This chapter describes enhancements provided by OS/390 Version 2 Release 4 JES3. When appropriate JES3 publications are referenced to provide additional information.

7.1 OS/390 Version 2 Release 4 JES3 Enhancements

The following sections list and briefly describe what is new, changed, or deleted in the OS/390 Version 2 Release 4 JES3:

- Hot start with refresh
- Segmented Initialization stream
- Dynamic configuration change support
- Faster JES3 restart
- JES3 64K Job Number Support
- MVS COMMAND JCL statement support
- JES3 statistics
- Deletion of JES3LIB support
- Deletion of initialization stream support from card readers

7.1.1 Hot Start With Refresh

JES3, prior to OS/390 Version 2 Release 4, requires at least a JES3 warm start (which implies a JES3 complex-wide restart with IPL of all processors) to change any keyword or parameter in the JES3 initialization stream. In OS/390 Version 2 Release 4 JES3, a new start type called *hot start with refresh* is added. Hot start with refresh reads the initialization stream and processes certain initialization statements (for example, FSSDEF, RJPWS, non-execution DEVICE).

The hot start with refresh allows JES3, on a hot start, to reread the initialization stream to allow many, but not all, of the parameters to be changed without an IPL. Note that a JES3 hot start only requires JES3 to restart on the global. Jobs currently in execution either on the global or on any local processors in the JES3 complex continue to execute.

Hot start with refresh syntax checks all initialization statements and parameters that are defined in the selected “inish” origin. However, not all initialization statements and parameters on initialization statements are processed during a hot start with refresh. Note that any errors found on the initialization statements

cause the hot start with refresh to fail even if the error is on a statement that is only syntax checked.

If a hot start with refresh fails, any changes that were made to the initialization stream are discarded. That is, a subsequent hot start obtains the configuration specifications from the last successful cold, warm, or hot start with refresh.

See 7.2, “JES3 Hot Start With Refresh” on page 149 for further details.

7.1.2 Segmented Initialization stream

Prior to OS/390 Version 2 Release 4 JES3, the JES3 initialization stream had to be contained as a single member of a partitioned data set, or it had to be specified as a single input stream.

OS/390 Version 2 Release 4 JES3 allows the initialization stream to be segmented. That is, different sections of the initialization stream can be put into different members and included in the primary initialization stream member.

You may use a new INCLUDE initialization statement to specify a member whose contents are to be included in the initialization stream. The included member must be in the data set specified on the first JES3IN DD data set in the JES3 procedure.

See 7.3, “JES3 Segmented Initialization stream” on page 163 for further details on the INCLUDE initialization stream statement.

7.1.3 Dynamic Configuration Change Support

Prior to OS/390 Version 2 Release 4 JES3, only a few dynamic configuration changes were allowed through operator commands. The *MODIFY,NJE,ADD= is an example of such an operator command. OS/390 Version 2 Release 4 JES3 dynamic configuration change support dynamic configuration change support provides for some additional configuration changes through operator commands, as follows.

7.1.3.1 *MODIFY,CONFIG Command

The OS/390 Version 2 Release 4 JES3 *MODIFY,CONFIG command allows you to make dynamic JES3 configuration changes which are equivalent to adding the following initialization statements to the initialization stream and taking a warm start:

- SNA RJP
  - RJPWS - SNA RJP work station characteristics
  - CONSOLE - SNA RJP consoles
  - DEVICE - SNA RJP devices
- Non-channel attached FSS printers
  - FSSDEF - Functional Subsystem Definition
  - DEVICE - I/O Device Definition

In the *MODIFY,NJE,ADD= operator command, all the parameters being processed are contained within the actual command itself. Unlike the NJE command, *MODIFY,CONFIG allows a subset of the JES3 initialization stream that resides in a member of the JES3 initialization data set, to be read in and processed dynamically. JES3 then makes the appropriate changes to its configuration without requiring a hot start with refresh or a warm start.

See 7.4, “Modify JES3 Configuration” on page 164 for further details on the *MODIFY,CONFIG operator command.

7.1.3.2 *MODIFY,T and *MODIFY W Command Improvements

OS/390 Version 2 Release 4 JES3 improves the *MODIFY,T and *MODIFY,W, operator commands by allowing you to modify the following for RJP work stations:

- G=, PL=, and TRACE= options with the *F,T command
- DYNAMIC= (DYN=), DGRPONLY (DGRPY=), PDIR, and XLATE options with the *F,W command.

Any information that is changed with these commands is not preserved over a hot start.

See 7.4.2, “*MODIFY,T and *MODIFY,W Command Enhancements” on page 167 for additional information on these command changes.

7.1.4 Faster JES3 Restart

Faster JES3 restart reduces the amount of JES3 downtime during a JES3 global restart and thus minimizes the delay that jobs in execution experience during a JES3 outage caused by a dynamic system interchange (DSI) or global hot start without an IPL. This performance enhancement is for global hot starts (including DSI) and warm starts.

Faster JES3 restart in OS/390 Version 2 Release 4 is accomplished by initiating I/O in parallel during the phase of JES3 initialization which is responsible for validating job control blocks and reallocating the spool space assigned to jobs.

Depending on the amount and complexity of the job structures in your installation, with this support you might see an improvement in the restart time over previous releases of JES3.

Some installations have observed a nearly 50% restart time reduction during a DSI from OS/390 Version 1 Release 3 JES3 global to a OS/390 Version 2 Release 4 JES3 local when there have been between 5000 to 15000 jobs in the job queue.

See 7.5, “OS/390 Version 2 Release 4 JES3 Global Restart Performance” on page 168 for additional information on this improvement.

7.1.5 JES3 64K Job Number Support

OS/390 Version 2 Release 4 JES3 allows up to 65,534 (64K) job numbers. Changes that result with this support occur in messages, the OPTIONS JOBNO= JES3 initialization statement, the GROUP JSPAN= JES3 initialization statement, various JES3 commands, and so on.

Prior this release of JES3, the maximum number of jobs in the JES3 complex was limited to 32,767 (32K) job numbers.

See 7.8, “OS/390 Version 2 Release 4 JES3 64K Job Numbers” on page 176 for additional information and considerations concerning the expansion of the allowable JES3 job number range.

7.1.6 JES3 MVS JCL COMMAND Statement Support

OS/390 Version 2 Release 4 JES3 allows the commands specified on the MVS JCL COMMAND statement in a job stream to be executed. The MVS COMMAND JCL statement command execution is controlled through new keywords (AUTH= and Command=) to the CIPARM initialization statement specifying the command groups allowed and their disposition.

AUTH= This specifies which command groups are accepted through COMMAND JCL statements in a job stream.

COMMAND= This specifies the disposition of commands entered through JCL COMMAND statements in the job stream.

Note that in prior JES3 releases, JES3 did not support the use of the MVS JCL COMMAND statement. If you attempted to use the statement in a JCL stream, it would be ignored by the system and no error message would be issued.

See 7.6, “OS/390 Version 2 Release 4 JES3 MVS JCL COMMAND Statement” on page 171 for additional information concerning the use of the MVS JCL COMMAND statement.

7.1.7 JES3 Statistics

OS/390 Version 2 Release 4 JES3 collects statistics about JES3 functions such as initialization, connect/restart, and output service. The statistics are collected using a new macro IATXSTAT. Dump Core (DC utility) can be used to dump the statistics for JES3 diagnostic and problem determination purposes.

Dump Core has a new start command option (*S DC OPTION=SDA) to produce a hexadecimal dump of the statistics data area for JES3 diagnostics and problem determination functions.

See 7.7, “OS/390 Version 2 Release 4 JES3 Statistics” on page 174 for further details on the diagnostic facilities that can be used to view and process the collected statistics.

7.1.8 Deletion of JES3LIB Support

OS/390 Version 2 Release 4 JES3 drops support for the JES3LIB initialization stream statement.

In prior releases of JES3, the JES3LIB initialization stream statement could be used to dynamically allocate a private JES3 TASKLIB for the IATNUC task. The TASKLIB data sets are searched for an entry name before the data sets defined either through the STEPLIB or JOBLIB DD statement and system libraries are searched.

A STEPLIB DD statement may be used in the JES3 procedure as the replacement for the removed JES3LIB initialization statement. If JES3 encounters a JES3LIB statement in the initialization stream, it issues a warning message (IAT3263) and allows initialization to continue.

7.1.9 Deletion of Initialization Stream Support from Card Readers

In OS/390 Version 2 Release 4 JES3, the support to allow the initialization stream to be read from a card reader (″U = ″ in message IAT3012) is deleted.

7.2 JES3 Hot Start With Refresh

OS/390 Version 2 Release 4 JES3 introduces a new type of JES3 start called *hot start with refresh*. Hot start with refresh rereads the initialization stream to allow many, but not all, of the parameters to be changed without an IPL. Previously, a JES3 warm start was required to perform this function. A JES3 warm start requires a JES3 complex re-IPL. A JES3 hot start, on the other hand, only requires JES3 to restart on the global. Jobs currently in execution either on the global or on any local processors continue to execute over hot start with refresh.

Hot start with refresh requires a *complete* initialization stream as input. During a hot start with refresh JES3 does not process all initialization statements and parameters on the initialization statements. Those statements and parameters which are not processed are, however, syntax-checked and will fail the initialization attempt if found erroneous. 7.2.1, “JES3 Initialization Statement Summary” on page 150 summarizes the JES3 initialization statements and the actions they can take during a hot start with refresh.

If a hot start with refresh fails, all changes that were made to the initialization stream are discarded. That is, a subsequent hot start obtains the information from the last successful cold, warm, or hot start with refresh.

A hot start with refresh overrides most previously issued *MODIFY commands relating to initialization statements. See Table 12 on page 155 in 7.2.2, “Summary of JES3 Restart Effect on Commands” on page 155 for the list of commands and how they remain in effect over a warm, hot, or hot start with refresh.

The JES3 operator prompt to specify the start type (message IAT3011) has new HR and HAR options that the operator can specify in the reply.

Note: HR stands for hot start with refresh and HAR for hot start with refresh and analysis.

Figure 87 on page 150 depicts an operator dialog for hot start with refresh. The dialog also shows the IAT3012 message with the changed initialization stream origin text (″U = ″ is deleted).

```
*RETURN,PASSWORD
*002 IAT3807 CONFIRM "RETURN" COMMAND FOR **GLOBAL** SC49 (CONTINUE (U) OR CANCEL)
2,U
 IEE600I REPLY TO 002 IS;U

 IAT7380 DJOB35 (JOB15878): IATDJSV - ESTAE WAS ENTERED - NO RETRY
S JES3
 IAT3040 STATUS OF JES3 PROCESSORS IN COMPLEX
 IAT3040 LO6RS007(  ), LO6RS008(  ), SC50    (UP), SC49    <  >
 IAT3040 SC43    (  )
 IAT3011 SPECIFY JES3 START TYPE
*003 IAT3011 (C, L, H, HA, HR, HAR, W, WA, WR, WAR, OR CANCEL)
3,HR
 IEE600I REPLY TO 003 IS;HR
*004 IAT3012 SELECT JES3 INISH ORIGIN (N OR M=), AND OPTIONAL EXIT PARM (,P=) OR CANCEL
4,M=B1
 IEE600I REPLY TO 004 IS;M=B1
 IAT4030 0001 SPOOL DATA SET  IN USE
 IAT4075 MAXIMUM NUMBER OF JOBS IS LIMITED TO 0006686 BY JCT DATA SET SIZE
*IAT3100 JES3 OS240 SYSTEM  HOTSTART ON 1997.308 AS SC49
 IAT7131 NJECONS NOW ACTIVE
 IAT9225 NJERDR IS ACTIVE
 IAT2801  SNARJP IS ACTIVE
*S JSS
 :::::::
****** SC49     CONNECT COMPLETE ******
 :::::::
****** SC50     CONNECT COMPLETE ******
```

Figure 87. JES3 Hot Start with Refresh Operator Dialog

7.2.1 JES3 Initialization Statement Summary

The following is a summary of the initialization statements and whether they can be changed, added, or deleted during a hot start with refresh. Dependencies on some of these statements are listed if there are other initialization statements or keywords that intersect with hot start with refresh processing. See *OS/390 JES3 Initialization and Tuning Reference* for the specific statements and the restrictions in place concerning these dependencies.

When you change the initialization stream, remember that some initialization statements are processed and some are ignored during a hot start with refresh. For example, if you add a new SPART statement to define a spool partition, and add or change a SYSOUT statement to reference that spool partition, JES3 issues message IAT3222 indicating that you specified a non-valid spool partition. This is because during a hot start with refresh, the SPART statements are ignored and the SYSOUT statements are processed.

Table 11 (Page 1 of 2). JES3 Initialization Statement Summary

Statement	Change	Add	Delete	Note(s)
ACCOUNT	Yes	Yes	Yes	
BADTRACK	No	No	No	
BADTRACK	No	No	No	
BUFFER	Yes	Yes	Yes	12
CIPARM	Yes	Yes	Yes	1
CLASS	Yes	Yes	Yes	2
COMMDEFN	Yes	Yes	Yes	

Table 11 (Page 2 of 2). JES3 Initialization Statement Summary				
Statement	**Change**	**Add**	**Delete**	**Note(s)**
COMPACT	Yes	Yes	Yes	
CONSOLE	Yes	Yes	No	
CONSTD	No	No	No	
DEADLINE	Yes	Yes	Yes	9
DEVICE	Yes	Yes	Yes	3
DYNALDSN	No	No	No	
DYNALLOC	Yes	Yes	Yes	13
FORMAT	No	No	No	
FSSDEF	Yes	Yes	Yes	5
GROUP	Yes	Yes	Yes	14
HWSNAME	Yes	Yes	Yes	1
INTDEBUG	Yes	Yes	Yes	
JES3LIB	No	No	No	10
MAINPROC	No	No	No	8
MSGROUTE	Yes	Yes	Yes	6
NJECONS	Yes	Yes	Yes	
NJERMT	Yes	Yes	Yes	
OPTIONS	Yes	Yes	Yes	4
OUTSERV	Yes	Yes	Yes	11
RESCTLBK	Yes	Yes	Yes	
RESDSN	Yes	Yes	Yes	1
RJPLINE	Yes	Yes	Yes	
RJPTERM	Yes	Yes	Yes	
RJPWS	Yes	Yes	Yes	
SELECT	Yes	Yes	Yes	8
SETACC	Yes	Yes	Yes	
SETNAME	No	No	No	
SETPARAM	No	No	No	
SETRES	Yes	Yes	Yes	
SPART	No	No	No	
STANDARDS	Yes	Yes	Yes	7
SYSID	Yes	Yes	Yes	
SYSOUT	Yes	Yes	Yes	
TRACK	No	No	No	

Table 11 on page 150 Notes:

1. If you change the HWSNAME, CIPARM, or RESDSN statements, you must restart any C/I FSSs that are active in order to make sure that JES3 honors any changes. If you perform an orderly shutdown of JES3, you can restart the C/I FSSs by specifying TERM=YES on the C/I FSSs′ FSSDEF statements.
2. The following items refer to the CLASS statement:

- If you delete a class, JES3 issues existing message IAT3414: CLASS class UNDEFINED OVER RESTART.
- If you delete a class that is still referenced by jobs, the jobs will be put into operator hold status at one of the following times:
  - During initialization if the job was active in main service
  - When the job is scheduled for main service

  You can use the *MODIFY,J command to change the job's class and to release it from operator hold.

  Do not delete a job class if is referenced by jobs that are in execution. This will cause JES3 failures.
- The SPART parameter has a dependency on the SPART statement, which is not processed during a hot start with refresh.
- MLIMIT, MDEPTH, SYSTEM parameters have dependencies on the MAINPROC statement which is not processed during a hot start with refresh.

3. The following items refer to the DEVICE statement:

   a. DEVICE statement - main processors
   Do not add or delete a device statement that is associated with a main processor; that is, one which contains the DTYPE=SYSMAIN parameter.
   b. DEVICE statement - I/O Devices

   You cannot add or delete devices that are defined as execution devices (that is, which have XTYPE and XUNIT parameters specified on the DEVICE statement).

   - You cannot add or delete a device number and system name from the XUNIT parameter of an existing DEVICE statement.
   - You cannot switch device numbers in the XUNIT parameter of two or more DEVICE statements.
   - You can, however, change any of the non-execution device-related parameters on the DEVICE statement, such as JNAME, FORMS, FLASH, and WS.

   You can only add or delete non-execution devices (that is, on which the XTYPE and XUNIT parameters are not specified); for example, non-channel-attached printers and NJE lines.

   JES3 is not sensitive to the order of the DEVICE statements. That is, new devices can be inserted anywhere and existing devices can be deleted.

   You cannot add or delete a DEVICE statement for an FSS-managed printer if the FSS is active at the time you perform the hot start with refresh. However, you can do this if the FSS is active on the global and the global was IPLed prior to performing the hot start with refresh.

   You cannot change or delete the device number or the system name from the JUNIT parameter if the FSS is active and a writer is using the device. (This only applies to the device number and system where the FSS is active).

   If you change the DEVICE statement for an FSS-managed printer and the FSS is active, the following checkpointed information overrides the information on the DEVICE statement:

   - MODE=COMP/FSS
   - Message destination (JUNIT)
   - Online/offline status (JUNIT)
   - CKPNT/CKPNTPG/CKPNTSEC
   - NPRO

- SETUPMSG

All other information from the DEVICE statement takes effect. Note that if the FSS was active on the global and the global was IPLed prior to performing the hot start with refresh, the above parameters from the DEVICE statement are used.

4. Only the DUMP, DUMPLINS, JOBNO, INTRDR, MT, and WANTDUMP parameters on the OPTIONS statement can be changed. All other parameters are syntax-checked only.

   You can decrease the low job number or increase the high job number over a JES3 hot start with refresh through the JOBNO parameter.

   You cannot delete job numbers from the job number range specified on the JOBNO parameter during a hot start with refresh. That is, you cannot raise the low job number value or lower the high job number value. It requires a warm start to do this.

5. The following items refer to the MAINPROC and SELECT statements:

   - You cannot delete an FSSDEF statement for an active FSS unless the FSS was active on the global and the global was IPLed prior to performing the hot start with refresh.
   - If an FSS is active, changes to the parameters on the FSSDEF statement are ignored. However, if the FSS was active on the global and the global was IPLed prior to performing the hot start with refresh, then the changes take place.
   - See the DEVICE statement for restrictions on adding, changing, and deleting devices associated with an FSS.
   - The SYSTEM parameter has a dependency on the NAME parameter of the MAINPROC statement. During a hot start with refresh, the MAINPROC statements are not processed; the MAINPROC statements from the last warm or cold start are used. Therefore, if you add a MAINPROC statement during a hot start with refresh, it is ignored. This causes error messages to be issued if the SYSTEM parameter of the FSSDEF statement references the MAINPROC statement that you attempted to add during a hot start with refresh.

6. Changes made to the global's MSGROUTE statement take effect immediately. Changes made to the local's MSGROUTE statement will not take affect until you restart the JES3 local address space.

7. The following items refer to the JES3LIB statement:

   - If you change the third subparameter of the BYTES, CARDS, LINES, and PAGES parameter, or if you change the SYSLINES keyword, the change takes effect immediately on the global processor. The change will not take effect on the local processor until you restart the JES3 local address space.
   - JES3 ignores any changes made to the SETUP and THWSSEP keywords.

8. The following items refer to the OUTSERV statement:

   - A select mode is associated with a main processor as follows:
     - Through the SELECT parameter of the MAINPROC statement
     - Through the *MODIFY,G,main,SELECT,MODE,mode command

   During a hot start with refresh, the main processors are reassigned the select mode that was in effect prior to the restart. JES3 ignores any changes to the SELECT parameter on the MAINPROC statement.

Therefore, do not delete the SELECT statements for any select modes that were in affect prior to the hot start with refresh.
- During a hot start with refresh, only the select mode name that was in effect prior to the hot start with refresh is reassigned to the main processors. The actual select mode values are taken from the parameters on the SELECT statements in the initialization stream.

9. DEADLINE statements can be deleted. But if you delete a DEADLINE statement that is referenced by a job, JES3 issues message IAT7415 when deadline processing begins. Any job that references the deleted DEADLINE statement uses the deadline algorithm information that was last defined for this DEADLINE statement.
10. The JES3LIB statement is deleted in OS/390 Release 4. Use STEPLIB in the JES3 procedure instead.
11. If you change the OUTLIM parameter of the OUTSERV statement, the change takes effect immediately on the global processor. The change will not take effect on the local processor until you restart the JES3 local address space.
12. The following items refer to the BUFFER statement:
    - You can only change the FD, MINBUF, and TRUNC parameters on the BUFFER statement; all other parameters on the BUFFER statement are syntax-checked only.
    - Changes to the FD and MINBUF parameters do not take effect immediately for C/I FSS if it is active, and therefore you must restart the C/I FSSs if you want its values changed.
13. The following items refer to the DYNALLOC statement:
    - Do not add, change, or delete DYNALLOC statements for the spool data sets or the JCT data set
    - On local processors, only a subset of the DYNALLOC statements are allocated:
      - Spool data sets
      - JES3JCT
      - JES3OUT
      - JESABEND
      - SYSABEND
      - SYSUDUMP
    - If a local processor becomes the global as a result of dynamic system interchange (DSI) processing, the remaining DYNALLOC statements are allocated:
      - IATPLBxx
      - JES3DRDS
      - JES3SNAP
      - Any user data sets
14. The following items refer to the GROUP statement:
    - Do not delete a GROUP statement or delete an EXRESC parameter from the GROUP statement that has active initiators unless you IPL the processors where the initiators are running.
    - Device fences that were active prior to the hot start with refresh are reallocated unless the definition of the device fence in the EXRESC or DEVPOOL parameters was changed. If the device fence definition was changed, the fence is reallocated after JES3 initialization completes.

- The EXRESC parameter has a dependency on the MAINPROC and SETNAME statements which are not processed during a hot start with refresh.
- The DEVPOOL parameter has a dependency on the SETNAME statement which is not processed during a hot start with refresh.

7.2.2 Summary of JES3 Restart Effect on Commands

Table 12 lists some JES3 commands and shows how they remain in effect over a warm, hot, or hot start with refresh.

Table 12 (Page 1 of 2). JES3 Commands That Remain in Effect After a Restart

Command	Operands	Warm Start	Hot Start	Hot Start with Refresh	Note
*MOFIFY C - Spool partition					
*MOFIFY CONFIG - Configuration change			X		
*MOFIFY E - Event tracing	ON OFF START STOP EXCL LIMIT DUMP DISPLAY TRAP FIND		X X X X	X X X X	1
*MOFIFY F - Functional subsystem			X	X	2
*MOFIFY G - Generalized main scheduling	G,ON G,OFF G,INIT G,ALLOC G,UNALLOC G,BAR G,JSPAN S,option C,ON C,OFF CHK SP		X X X X X X X X	 X	3
*MOFIFY J - Job status	H R C CP P	X X X X X	X X X X X	X X X X X	
*MOFIFY L - Deadline scheduling					
*MOFIFY M - MCS routing					
*MOFIFY MT - Writer output multitasking					
*MOFIFY N - DJC network		X	X	X	
*MOFIFY NJE - NJE communication					
*MOFIFY O - Console information	DLOG A= O=,ADEST= O=,DDEST= O=,AROUT= O=,DROUT= O=,SAVEMSG=		X X X X X X X	X	4

Table 12 (Page 2 of 2). JES3 Commands That Remain in Effect After a Restart

Command	Operands	Warm Start	Hot Start	Hot Start with Refresh	Note
*MOFIFY Q Job queue and spool data sets	P= SP= DRAIN USE HOLD STOP RELEASE CANCEL	X X X X X X X X	X X X X X X X X	X X X X X X X X	
*MOFIFY S - Setup	VU= VA= M= U= J= AL= ALWIO= MR= THWSSEP DEFERCT	X X	X X	X X	
*MOFIFY T - Remote job processing	T=,H T=,R T=,B= T=,JOB=R T=,JOB=C T=,P= L=,H L=,R L=,P= L=,A L=,M L=,SNAP L=,TRACE	X X X X X	X X X X X	X X X X X	
*MOFIFY U - Output service	N= CONS= All others	 X	 X	 X	
*MOFIFY V - Change online or offline status			X	X	5
*MOFIFY X - Use counts					
*CALL			X	X	
*START/*RESTART					
*SWITCH - Switch RJP console messages			X		1

Table 12 on page 155 Notes:

1. Changes are lost if you IPL the global processor.
2. For Hot Start with Refresh:
   - If FSS is inactive, changes will not remain in effect.
   - If FSS is active on local, changes will remain in effect.
   - If FSS is active on global and global is not IPLed, changes will remain in effect.
   - If the FSS was active on the global and the global was IPLed, the changes will not remain in effect.
3. S,MODE is the only option preserved across hot start with refresh.
4. All changes except what is specified for DLOG are lost if you IPL the global processor.
5. Changes are only maintained for JES3-managed devices.

7.2.3 Hot Start With Refresh - How Is It Done?

Most of the JES3 initialization statements deal with information that is kept only on the JES3 global processor in the private area of the JES3 global address space. OS/390 Version 2 Release 4 JES3′s hot start with refresh reads the initialization stream during a hot start without IPL and allows many of the “global only” parameters to be changed.

7.2.3.1 Hot Start With Refresh Initialization Statement Processing

For all JES3 start types, the information extracted from the initialization statements is written to spool and is called *intermediate text* Intermediate text is saved to a JES3 multi-record file (MRF) and is checkpointed. All information for the same statement is written to the same intermediate text file. For example, all CIPARM statement information is written to the same file. Note that sometimes information from multiple statements is written to the same file.

During a hot start with refresh, not all initialization statements or parameters are processed. For example, the intermediate text for the MAINPROC statement is not recreated during a hot start with refresh. The information from the last warm or cold start is used from checkpointed intermediate text files.

New intermediate text is created for the initialization statements and parameters that are processed during a hot start with refresh.

Prior to OS/390 Version 2 Release 4 JES3, the only way to change the JES3 initialization data was to perform a warm or cold start. If the warm or cold start failed, the global must be re-IPLed and another warm or cold start must be performed, at which point the initialization stream is read and new intermediate text is created. A hot start with refresh can be initiated by just restarting the JES3 global address space without an IPL.

If a hot start with refresh fails, JES3 reverts back to the previous intermediate text when *a hot start* on the global is performed.

To make sure that JES3 initialization does not leave the intermediate text in a state where a partial configuration exists (that is, where part of the configuration reflects the old information and part of it reflects the new information), *all* types of OS/390 Version 2 Release 4 JES3 initialization is changed to maintain consistent initialization data in new control blocks called the *old configuration data area*((IATY)OCF) and the *new configuration data area* ((IATY)NCF).

At the end of a successful initialization processing, just before the new initialization checkpoint record (ICP) is written to spool, the information from the NCF data area is moved to the appropriate JES3 control blocks. For example, information that resides in the SSVT is copied into the SSVT before the checkpointed part of the SSVT is put into the ICP. File directory blocks (FDB) and other information that resides in the checkpoint data area (CKP) is copied from the NCF to the CKP and written to the checkpoint data set, and the spool space for the old information can be freed. The old initialization checkpoint record is left unchanged if any initialization statement errors are encountered. If the failing JES3 initialization attempt is not the first in the life of the JES3 complex, JES3 may be hot started from the previous valid checkpointed initialization data.

IATXCFSV - Configuration Services: The OS/390 Version 2 Release 4 JES3 IATXCFGS macro provides functions related to the JES3 configuration:

- INITIALIZE - Initialize configuration services.
- READ_ACCESS - Obtain read access to the configuration.
- WRITE_ACCESS - Obtain write access to the configuration.
- RELEASE_ACCESS - Release access to the configuration.
- ADD_NEW_INFO - Add information associated with the new configuration to the new configuration data area (NCF). This information is committed to the new configuration when a COMMIT_NEW_INFO request is issued.
- COMMIT_NEW_INFO - Commit/harden the information associated with the new configuration by placing it in the appropriate JES3 control blocks and checkpoint areas.
- ADD_OLD_INFO - Add information associated with the old configuration to the old configuration data area (OCF). If the information represents a spool control block, the spool space is deleted (purged) when a DELETE_OLD_INFO request is issued.
- DELETE_OLD_INFO - Delete information associated with the old configuration. For example, if the information represents a spool control block, the spool space for the control block is purged.
- TERMINATE - Terminate configuration services.

The IATCFSRV module services the IATXCFGS macro.

JES3 Checkpoint Data Set Serialization: Even though a hot start with refresh changes mostly global-only parameters, there are some cases where the information needs to be read on a local processor, such as intermediate text for the MSGROUTE, MAINPROC, DYNALLOC, STANDARDS and the SUPUNITs information from the checkpointed part of the SSVT. (However, the only SUPUNIT used is the DTYPE-SYSMAIN SUPUNIT.)

C/I FSS′s address spaces also read some intermediate text files when they initialize, such as the intermediate text for the CIPARM, RESDSN, HWSNAME, and SYSOUT statements.

OS/390 Version 2 Release 4 JES3 changes the checkpoint data set serialization processing to allow correct processing of the configuration information during a hot start with refresh and *MODIFY,CONFIG command.

In prior releases of JES3, the only way to change the configuration information was to perform a warm or cold start. This guaranteed that there were no C/I FSSs active when the configuration was being changed, since an IPL is required to perform a warm or cold start and all local processors must have been disabled. When global is reinitialized with a warm or cold start, it rewrites the JES3 complex status record (CSR) with only the global′s main name. If a local attempts to come up when the global is performing a warm or cold start, the local issues the following message:

```
IAT3043 GLOBAL IS REINITIALIZING CHECKPOINT DATA. DO YOU NEED TO
TAKE CONTROL (YES NO OR RETRY)
```

The proper response to this message if the global is in the process of performing a cold or warm start is RETRY, which should be done after the global has completed the warm or cold start.

OS/390 Version 2 Release 4 JES3 global during a hot start with refresh only updates the CSR to reflect the fact that the global is in the process of initializing while the locals and C/I FSS address spaces continue to run. To prevent the JES3 locals and C/I FSS address spaces from accessing the configuration information during update or modify attempts, the OS/390 Version 2 Release 4 JES3 global, locals, and C/I FSS address spaces serialize through ENQ/DEQ macro service access to the checkpoint data set. (Read-write updates of a logical record in the checkpoint data set still require RESERVE EXCLUSIVE/DEQ major name SYSZIAT minor name ckptvolckptdsn to be issued for the data sets across the entire operation.)

JES3 issues a SYSTEMS scope ENQ with RNL=NO (guaranteed serialization - no RNL processing) to serialize access to the JES3 configuration information. The ENQ major name is SYSZIAT and the minor name is CONFIG.CHANGE.ckptvol.ckptdsn, where:

- ckptvol is the volume serial number of the primary JES3 checkpoint data set.
- ckptdsn is the name of the primary JES3 checkpoint data set.

Depending on the type of processing, ENQ′s use of resource type is:

- SHARED for cold, warm, DSI, hot, local starts, and C/I FSS starts
- EXCLUSIVE for hot starts with refresh and *MODIFY,CONFIG

Figure 88 shows an example of the JES3 checkpoint data set serialization when a hot start with refresh and a local start are proceeding in parallel.

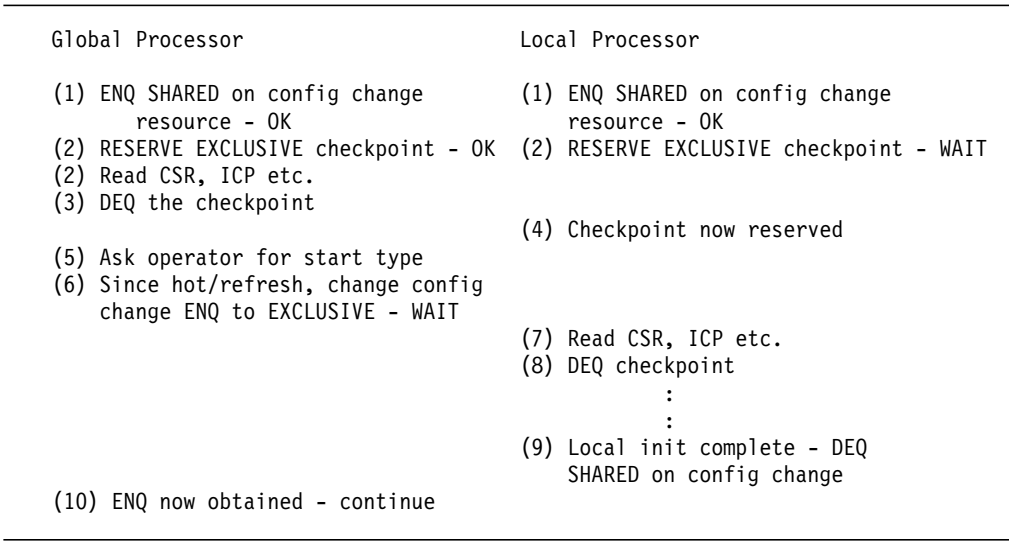

Figure 88. JES3 Checkpoint Data Set Serialization

SYSUNIT Table Changes: Each DEVICE statement that is specified in the initialization stream causes an entry to be created in the SYSUNITs table. The SYSUNITS table contains device allocation status for each device in the entire JES3 complex. The SYSUNITs table is located in CSA, and all of the entries in the table are contiguous prior to OS/390 Version 2 Release 4 JES3 systems. An entry in the contiguous SYSUNIT table is accessed either by using an index value (called a SYSUNITs index) to compute the offset of the desired entry, or else the table is scanned from the top until the requested entry is found or the end of the table is reached.

OS/390 Version 2 Release 4 JES3 restructures the SYSUNIT table to accommodate limited dynamic DEVICE statement changes through the hot start with refresh or the *MODIFY CONFIG operator command. The restructured SYSUNIT table allows it to be extended dynamically, and it will consist of one or more blocks of contiguous SYSUNITs entries. Each SYSUNITs table block is preceded by a header which contains the following information:

- A pointer to the next SYSUNITs table block
- Number of SYSUNITs entries in this block
- The high and low SYSUNITs index values in this block
- A pointer to the first free SYSUNITs entry and the number of free SYSUNITs entries in this block. New devices first try to use an available slot before creating a new SYSUNITs block. Each free slot contains a pointer to the next free slot.

Figure 89 shows the new SYSUNITs table structure.

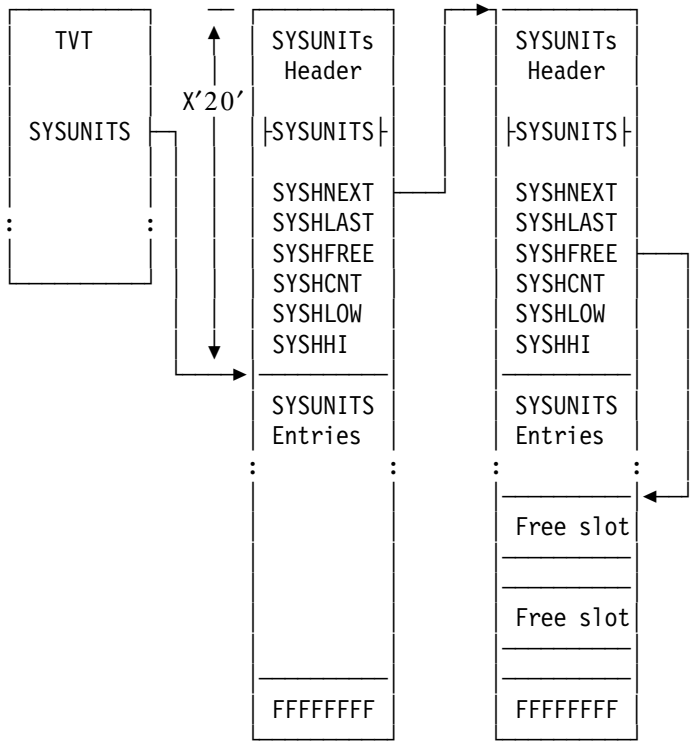

Figure 89. OS/390 Version 2 Release 4 JES3 Restructured SYSUNIT Table

When a device is added, the SYSUNITs table headers are examined to see if there are any free slots. If there is a free slot, the slot is removed from the free queue and used to build the new SYSUNITs entry. In addition, the number of free SYSUNITs entries in the block is decremented. If there are no free slots in any of the SYSUNITs block, a new SYSUNITs block of 20 entries is created and chained to the last block.

When a device is deleted, the SYSUNITs entry is cleared, added to the free chain in the SYSUNITs block, and the number of free SYSUNITs entries is incremented.

Note: Since execution devices cannot be added or deleted, the execution devices always appear in the first SYSUNITs block. This means that the MDS

which uses the SYSUNIT index to index into the SYSUNITs table still works without any changes. However, MDS or other functions running under JES3 which scan the SYSUNITs table are changed since it is possible that there are free SYSUNITs entries in the table, and these must be skipped over.

To make more JES3 code unaware of the structure of the SYSUNITs table, the IATXSYSU macro is provided to perform SYSUNIT's related functions. The functions of the IATCSYSU include:

- DIRECT_ACCESS to access a particular SYSUNIT's entry by index value. This function expands inline for fast access.
- SEQ_ACCESS to access the SYSUNITs entries sequentially. Free entries are skipped over. This function expands inline for fast access.

Dynamic Destination Queue Changes: The destination queue is a control block used by subsystem interface routines to route requests (represented by destination codes) to the JES3 routines responsible for servicing the requests.

The destination queue entry for FSS communication is a dynamic destination queue. Unlike normal destination queues which contain a pointer to the staging area chain for the function, the dynamic destination queue points to an FSS level destination queue table. This table is built in CSA and contains one entry for every FSS defined and is indexed using the FSS portion of the FSID. (JES assigns a FSS, and FSAs under the FSS, an identifier of the form *xxxxyyyy*, where *xxxx* corresponds to the controlling FSS identifier, and *yyyy* is a unique number for each the FSA.)

Each entry in the FSS level destination queue points to an FSA level destination queue table. Each FSA level destination queue table contains an entry for the FSS and each FSA associated with the FSS, and is indexed using the FSA portion of the FSID. Figure 90 shows the data structure.

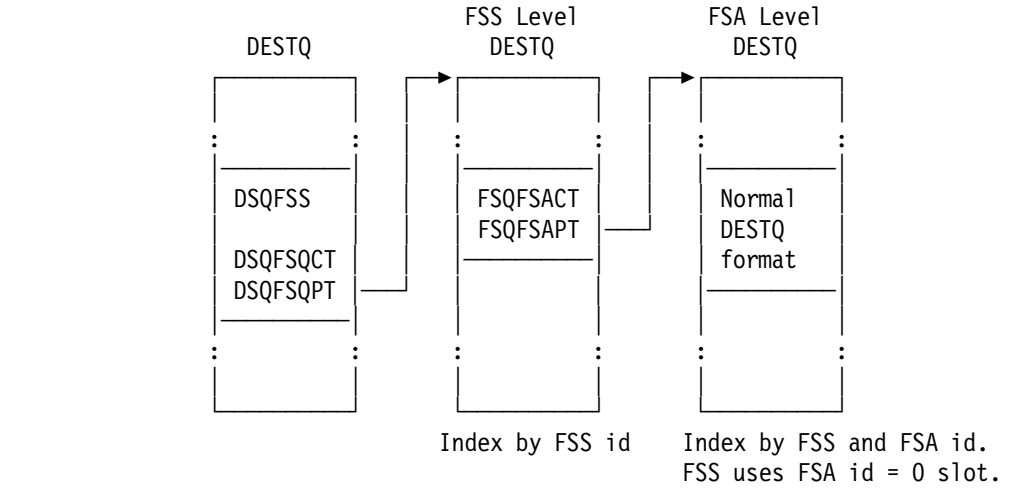

Figure 90. Pre-OS/390 Version 2 Release 4 JES3 Dynamic Destination Queue Structure

During a hot start with refresh, the number of FSSs can be increased by adding one or more FSSDEF statements to the initialization stream. OS/390 Version 2 Release 4 JES3 changes the dynamic destination queue to be expandable. The FSS destination queue entry is set to point to the first FSS-level destination queue. High and low FSS ID fields are added to the destination queue to show

which FSS IDs are represented by this table. If there are additional FSSs that cannot be represented by the first FSS-level destination queue, a dynamic destination queue extension is created to accommodate additional FSSs and the FSS destination queue entry is set to point to the dynamic extension.

Figure 91 shows the OS/390 Version 2 Release 4 JES3 format of the dynamic destination queue. Assume that prior to the hot start with refresh, 20 FSSs were defined. As a result, an FSS-level destination queue entry is created that is large enough to accommodate 20 FSSs.

In a hot start with refresh, the number of FSSs is increased to 30. As a result, a dynamic destination queue extension is created to accommodate FSSs 21 through 30, and chained to the first dynamic destination queue entry. In addition, an FSS level destination queue is created for these FSSs.

Note: During a hot start with refresh, the installation can also increase the number of FSA′s associated with an FSS by adding one or more DEVICE statements that specify the FSS on the FSSNAME parameter. This, however, does not require the FSA-level destination queues be expandable because during an FSS start, any existing FSA-level destination queue entries are freed and subsequently rebuilt. (FSAs cannot be added to an active FSS.)

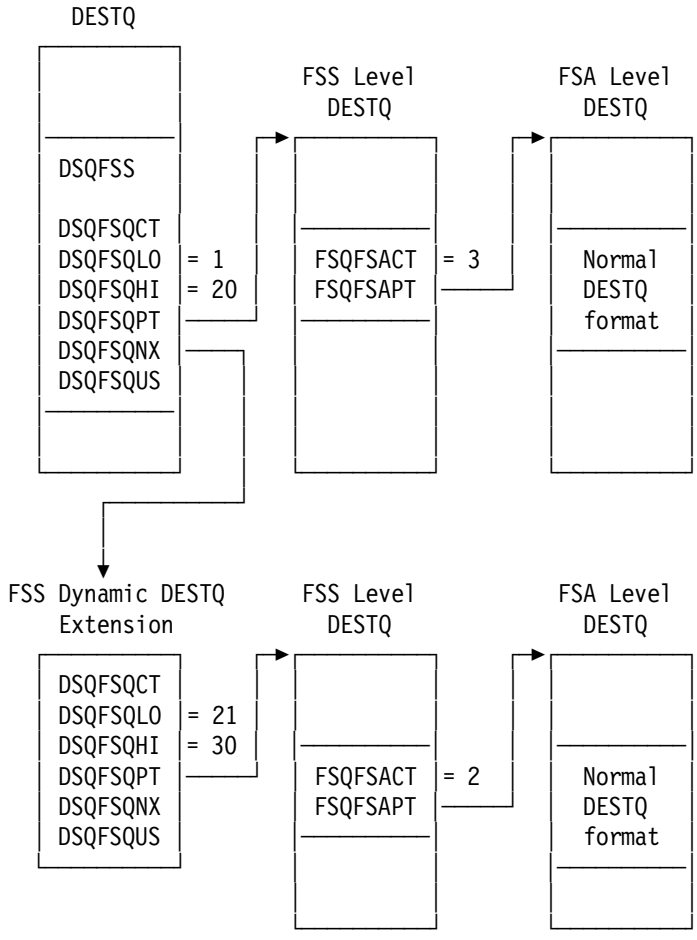

Figure 91. OS/390 Version 2 Release 4 JES3 Dynamic Destination Queue Structure

The JES3 DSQLOC macro is used to locate a destination queue entry. The JES3 DSQLOC macro is changed to support the OS/390 Version 2 Release 4 JES3 dynamic destination queue in a way that is transparent to its users.

7.3 JES3 Segmented Initialization stream

OS/390 Version 2 Release 4 JES3 allows the initialization stream to be segmented into different sections. Each section is a member in the partitioned data set specified on the JES3IN DD statement in the JES3 procedure, and is imbedded into the primary initialization stream with an INCLUDE initialization statement. Each INCLUDE initialization statement can refer only to one member of the partitioned data set. Up to four member levels can be used (the primary initialization stream member and up to three INCLUDE level members). Note that up to four levels of INCLUDE statement can also be used in the member specified on the *MODIFY,CONFIG,ADD= command.

Note: The INCLUDE statement cannot be used if the JES3IN data set is concatenated and the members to be included are in a data set other than the first in the concatenation.

INCLUDE statements cannot be placed in an initialization stream before the DYNALLOC statements that are in the stream.

The INCLUDE syntax is:

```
INCLUDE,MEMBER=member
```

Where member specifies the member name in the partitioned data set specified on the JES3IN DD statement in the JES3 procedure. This member is expected to contain initialization statements that are to be included in the initialization stream.

For example, if the members in the JES3IN data set are JES3INU2, A, B, and C, you can specify your INCLUDE initialization statements as follows:

```
JES3IN PDS Member        Member Contents
_________________        _______________
Member JES3INU2          ::::
                         DYNALLOC,...
                         INCLUDE,MEMBER=A
                         ::::
                         ENDINISH
                         ______
Member A                 INCLUDE,MEMBER=B
                         ______
Member B                 INCLUDE,MEMBER=C
                         ______
Member C                 Other JES3..
                         .initialization..
                         ..statements
                         ______
```

JES3 detects INCLUDE statement loops:

```
JES3IN PDS Member          Member Contents
_________________          _______________
Member JES3INU2            ::::
                           DYNALLOC,...
                           INCLUDE,MEMBER=A
                           ::::
                           ______
Member A                   INCLUDE,MEMBER=B
                           ______
Member B                   INCLUDE,MEMBER=A
                           ______
IAT3093 INCLUDE MEMBER A IGNORED — ALREADY ACTIVE
```

JES3 shows, in the JES3OUT data set or in the *MODIFY,CONFIG log in columns 73 through 80, the included member name to help determine from where the particular statements came. The following example comes from a *MODIFY,CONFIG log:

```
IAT8351 DSN = SYS1.PARMLIB , MEMBER = JES3INU2, PARMS =

   INCLUDE,MEMBER=AA                                              JES3INU2
   INCLUDE,MEMBER=AAA                                             AA
   INCLUDE,MEMBER=AAAA                                            AAA
   INCLUDE,MEMBER=AAAAA                                           AAAA
IAT3093 INCLUDE MEMBER AAAAA IGNORED - MAXIMUM NUMBER OF LEVELS HAS BEEN REACHED
```

7.4 Modify JES3 Configuration

OS/390 Version 2 Release 4 JES3 allows you to add SNA/RJP workstations and non-channel-attached FSS printers dynamically without having to restart JES3. To add new definitions, you code initialization statements (just as you would if you were restarting JES3), and specify the member containing the initialization statements on the *MODIFY,CONFIG command. The member must be in the data set specified on the JES3IN DD statement that was used to start JES3.

The following initialization statements are supported during *MODIFY,CONFIG processing:

- RJPWS to define SNA/RJP workstation characteristics.
- CONSOLE to define SNA/RJP console
- DEVICE to define SNA/RJP devices and non-channel-attached FSS-managed printers.

  Note: There is no order dependency among the SNA/RJP statements. Note also that there is *no* dynamic change support for COMPACT and COMMDEFN SNA/RJP initialization statements. The COMMDEFN statement (communication SSI definition records) specifies the optional APPLID= for the SNA/RJP name as specified to VTAM, P= password for the SNA/RJP VTAM application, and LU= the limit for the number of active sessions between SNA RJP and VTAM. The COMMDEFN statement can be changed using a hot start with refresh.

- FSSDEF to define writer FSS.

  Note: There is no order dependency for FSSDEF and related DEVICE statements.

- INTDEBUG which is the initialization debugging facility. The INTDEBUG statement specifies error message text and an index value. If the specified message text is issued the number of times indicated by the index value, JES3 issues a U005 JES3 user abnormal end and takes a storage dump.
- INCLUDE to include another initialization stream member.

All the *MODIFY,CONFIG changes are checkpointed so that they remain in effect over hot starts. However, if you perform a hot start with refresh, the changes are lost. Therefore, make sure you update your initialization stream *before* performing a hot start with refresh, warm, or cold start.

In addition, OS/390 Version 2 Release 4 JES3 is enhanced to allow the *MODIFY,T command to dynamically change the device group name for SNA or BSC RJP (G=), the number of times an incorrect password will be allowed (PL=), and the tracing option (TRACE=) parameters defined on the RJPWS and DEVICE initialization statements.

Also, the *MODIFY,W command is enhanced to allow the start option (DYNAMIC= or DYN=), the locally destined data set processing option (DGRPONLY= or DGRPY=), the peripheral data set information record sending (PDIR=), and the translate non-displayable characters (XLATE=) parameters defined on a SNA RJP DEVICE initialization statement to be dynamically changed.

7.4.1 *MODIFY,CONFIG Command

You may use the *MODIFY,CONFIG command to add dynamically the SNA RJP workstations and non-channel-attached FSS printers to the JES3 configuration. The following initialization statements are accepted by the *MODIFY,CONFIG processing:

- SNA RJP
  - RJPWS - SNA RJP workstation characteristics
  - CONSOLE - SNA RJP consoles
  - DEVICE - SNA RJP devices
- Non-channel-attached FSS printers
  - FSSDEF - functional subsystem definition
  - DEVICE - I/O device definition

If unsupported initialization statements are included, the error message IAT8346 STATEMENT stmt NOT VALID FOR DYNAMIC CHANGE is issued.

The syntax of command *MODIFY,CONFIG is as follows:

```
*MODIFY,CONFIG,ADD=mem_name[,LOG={YES|NO|ERR}][,P=XXXXXXXX]
*F
```

Where:

- ADD=mem_name specifies the member name to be read from the data set allocated to the JES3IN DDNAME in the JES3 procedure.
- LOG=YES/NO/ERR specifies whether you want to record each statement processed and any error message generated in a spinoff data set under JES3 JOB 0. (The log data set DDNAME is MODIFY.CONFIG.LOG.)
  - YES - Create the log data set.
  - NO - Do not create the log data set. All error messages are displayed on the issuing console.
  - ERR - Create the log data set only if errors occur.
- P=XXXXXXXX is a parameter string which is passed to the JES3 scan an initialization statement exit (IATUX15) as the statements are processed.

As an example, suppose you want to add dynamically the JES3 initialization statements for the sample functional subsystem (FSS) and FSS application (FSA) IAZSFSS found in SYS1.SAMPLIB. The printer definition is created into the JES3IN partitioned data set as member NEWFSS:

```
DEVICE,DTYPE=PRTAFP1,
  JNAME=IAZFSS,DGROUP=LOCAL,DGRPONLY=NO,
  JUNIT=(,LO6RSO08,UR,,,LO6RSO07,UR,,,SC50,UR,,,SC49,UR,,,SC43,UR),
  FSSNAME=WTRIAZF,DYNAMIC=NO,BURST=NO,FORMS=(YES,STANDARD),
  WC=(C,D),WS=(PM,D,CL,P)
FSSDEF,TYPE=WTR,FSSNAME=WTRIAZF,PNAME=WTRIAZF,TERM=YES
```

Next, issue the *MODIFY,CONFIG command to activate the printer definition:

```
*F CONFIG ADD=NEWFSS LOG=NO
*506 IAT8337 CONFIRM "*F CONFIG" COMMAND (CONTINUE(U) OR CANCEL)
506,U
 IEE600I REPLY TO 506 IS;U
 IAT8350 CONFIGURATION MODIFICATION IS COMPLETE - NO ERRORS
*I D D=IAZFSS
 IAT8562 IAZFSS        (AV )          F=STD CH=(GT15) DYN=N
 IAT8562 IAZFSS                       H=Y B=N M=A
 IAT8562 IAZFSS                       L=0+ PG=0 CK=200P DGRPY=N MK=C
 IAT8562 IAZFSS                       NPRO=90 MODE=FSS FSS=WTRIAZF
 IAT8562 IAZFSS                       C=STD3 PM=LINE,PAGE
 IAT8562 IAZFSS                       SS=C CM=(NONE,0)
 IAT8562 IAZFSS                       SETUPMSG=Y
 IAT8562 IAZFSS                       PDEFAULT=(NONE)
 IAT8562 IAZFSS                       WS=(PM,D,CL,P)
 IAT8562 IAZFSS                       WC=CD DGRP=LOCAL
 IAT8562 IAZFSS                       XLATE=Y
 IAT8567 LO6RSO07=      LO6RSO08=      SC50=
 IAT8567 SC49=          SC43=
```

Next, reissue the same *MODIFY,CONFIG command with the LOG=ERR option. JES3 issues the following messages:

```
*F CONFIG ADD=NEWFSS LOG=ERR
 IAT8348 WARNING LEVEL MESSAGE(S) ISSUED DURING INITIALIZATION STATEMENT PROCESSING
*507 IAT8337 CONFIRM "*F CONFIG" COMMAND (CONTINUE(U), CANCEL, OR LOG)
507,U
 IEE600I REPLY TO 507 IS;U
 IAT8069 MESSAGES WERE GENERATED - SEE LOG FOR DETAILS
 IAT8350 CONFIGURATION MODIFICATION IS COMPLETE - WARNING MESSAGES
```

The *MODIFY,CONFIG log contains:

```
IAT8351 *MODIFY,CONFIG - DATE = 1997.314, TIME = 18.00.49

IAT8351 DSN = SYS1.PARMLIB , MEMBER = JES3INTS, PARMS =

   DEVICE,DTYPE=PRTAFP1,
     JNAME=IAZFSS,DGROUP=LOCAL,DGRPONLY=NO,
     JUNIT=(,LO6RS008,UR,,,LO6RS007,UR,,,SC50,UR,,,SC49,UR,,,SC43,UR),
     FSSNAME=WTRIAZF,DYNAMIC=NO,BURST=NO,FORMS=(YES,STANDARD),
     WC=(C,D),WS=(PM,D,CL,P)
   FSSDEF,TYPE=WTR,FSSNAME=WTRIAZF,PNAME=WTRIAZF,TERM=YES
IAT3168 FSSDEF FOR WTRIAZF IGNORED, FSS ALREADY EXISTS

IAT8348 WARNING LEVEL MESSAGE(S) ISSUED DURING INITIALIZATION STATEMENT PROCESSING

IAT3219 DEVICE NAME IAZFSS   ALREADY EXISTS
```

Note: In OS/390 Version 2 Release 4 JES3, user exit 15 examine an initialization statement (IATUX15) is called in two additional environments, as compared with the previous releases. The first new exit point is when a *MODIFY,CONFIG command is being executed. The second is when a hot start with refresh is being executed. Exit 15 can perform the same processing in these two cases as it did previously during a cold or warm start.

7.4.2 *MODIFY,T and *MODIFY,W Command Enhancements

The following list summarizes the *MODIFY,T and *MODIFY,W command enhancements in OS/390 Version 2 Release 4 JES3:

- *MODIFY,T,T={wsnameALL},G=grpname
- *MODIFY,T,T={wsnameALL},PL=n
- *MODIFY,T,T={wsnameALL},TRACE={ONOFF}
- *MODIFY,W,devname,DYN={YESYNON} (or DYNAMIC=)
- *MODIFY,W,devname,DGRPY={YESYNON} (or DGRPONLY=)
- *MODIFY,W,devname,PDIR={ALLBDS}
- *MODIFY,W,devname,XLATE={YESYNON}

Following are restrictions and notes regarding the use of these commands:

- *MODIFY,T,...,G= is not allowed if the workstation is logged on.
- *MODIFY,T=ALL with G= is not allowed.
- *MODIFY,T,...,G={LOCALANYLOCAL} is not allowed.
- *MODIFY,W is not allowed if the workstation is signed on.

- Information on these commands is not preserved over a hot start.

Additionally, the following changes are made to *INQUIRY,D command output to support the preceding enhancements:

- *INQUIRY,D,D= displays XLATE and PDIR information.
- *INQUIRY,D,D= can display non-logged-on RJP workstations.
- *INQUIRY,D,T= is enhanced to display PL information.
- Completion messages for *INQUIRY,D,D= and *INQUIRY,D,T= are provided:
  - *INQUIRY,D,D= - IAT8500 INQUIRY ON DEVICES COMPLETE
  - *INQUIRY,D,T= - IAT8608 INQUIRY ON RJP COMPLETE

The following example shows a *MODIFY,W command and its result for the FSS writer that was defined with the *MODIFY,CONFIG command:

```
*MODIFY W IAZFSS XLATE=N
 IAT8334 MODIFY COMPLETE FOR DEVICE IAZFSS
*INQUIRY D D=IAZFSS
 IAT8562 IAZFSS        (AV )          F=STD CH=(GT15) DYN=N
 IAT8562 IAZFSS                       H=Y B=N M=A
 IAT8562 IAZFSS                       L=0+ PG=0 CK=200P DGRPY=N MK=C
 IAT8562 IAZFSS                       NPRO=90 MODE=FSS FSS=WTRIAZF
 IAT8562 IAZFSS                       C=STD3 PM=LINE,PAGE
 IAT8562 IAZFSS                       SS=C CM=(NONE,0)
 IAT8562 IAZFSS                       SETUPMSG=Y
 IAT8562 IAZFSS                       PDEFAULT=(NONE)
 IAT8562 IAZFSS                       WS=(PM,D,CL,P)
 IAT8562 IAZFSS                       WC=CD DGRP=LOCAL
 IAT8562 IAZFSS                       XLATE=N
 IAT8567 L06RS007=      L06RS008=      SC50=
 IAT8567 SC49=          SC43=
 IAT8500 INQUIRY ON DEVICES COMPLETE
```

7.5 OS/390 Version 2 Release 4 JES3 Global Restart Performance

There are occasions when the global processor must be restarted, such as when you have to implement configuration changes that cannot be made dynamically, or to apply service. While the global is down, jobs in execution may experience delays when they request global services such as allocating spool space or opening SYSOUT data sets.

Each address space requiring JES3 services has dependencies on the global in order to process work. The loss of the global can potentially cause the entire sysplex to stop functioning until the global is again restarted.

To minimize the delay that jobs in execution may experience, JES3 global restart time is reduced. The most time-consuming part of the global initialization is the job validation phase. Job validation includes:

- Validating the control blocks associated with a job (JCT, JDAB, JDS, JOBTAT, JMR)
- Reallocating the spool space associated with a job and making sure no two jobs are using the same spool space

- Determining where the job was prior to the restart, and building a RQ for the job

Although JES3 uses multiple functions (FCTs) to do job validation, and the JCT data set is read a track at a time, job validation still takes a long time for large job queues because of the number of job control blocks that must be read in and validated. JES3′s spool resident control blocks are chained together and the chain anchor for some control blocks cannot be found at the top of the chain, which dictates a predefined read-in order.

For example, to find a job′s spool data sets′ space allocation information for job validation, JES3 obtains the address of the job data set control block (JDS) from the JCT, and then reads in the first JDS record, loops through each JDS entry, and checks whether the data set has its own track allocation table (JBT). If a data set has a JBT, it is read in and validated. If the job has chained JDS records, all the records are read in one by one and the entries in each record are validated.

Each spool resident control belonging to a job is stored in a single physical record (buffer). The number of control blocks associated with a job, and the number of I/Os required to access them, is as follows:

- Job track allocation table (JBTAT) - One or more buffers (typically only one buffer).
- Job data set control block (JDS) - One or more buffers (mostly only one buffer).
- Data set track allocation table (DSTAT) - A job can have zero or more JDS entries with a DSTAT. Each DSTAT requires one or more buffers (typically only one buffer).
- Job data accounting block (JDAB) - One buffer.
- Job management record (JMR) - One buffer.

Figure 92 on page 170 shows a typical JES3 job control block structure and the control block read order during job validation.

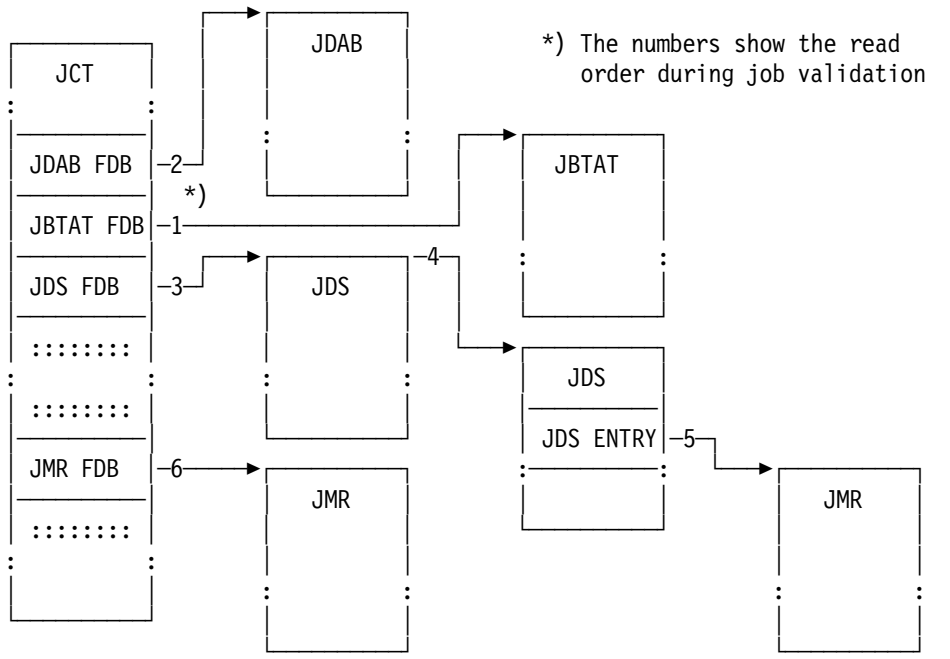

Figure 92. JES3 Control Block Read Order during Job Validation

In JES3 releases prior to OS/390 Version 2 Release 4, I/O for these control blocks is performed serially. That is, each control block is read and processed before the I/O for the next control block is initiated. Therefore, if four control blocks are required to be read fo a job, the amount of time spent in job validation for that job is the amount of time it takes to complete four I/Os.

To improve JES3 global restart time in OS/390 Version 2 Release 4 JES3, the I/Os for many job control blocks are initiated in parallel. Under ideal conditions, all parallel I/Os complete at about the same time and the I/O response time is about the same as it takes to complete a single I/O.

For example, the first and only control blocks to be read in for a simple job are JBTAT, JDAB, JDS, and JMR. If I/O for these control blocks are initiated in parallel, the time spent in job validation for that job is reduced approximately by the time it takes to complete three single I/Os.

I/O for the following control blocks is initiated in parallel:

- JDAB, JMR, JBTAT and the first JDS buffer
- The second JDS buffer and all DSTATs in the first JDS buffer
- The third JDS buffer and all data set TATs in the second JDS buffer, and so on

Table 12 shows the relationship between the number of JDS buffers and the number of logical I/O requests that are performed. It assumes that a job has one JDAB, one JMR, and one JBTAT.

Table 13. JES3 Initialization Statement Summary			
Number of JDS Buffers	**Serial I/O (Old Way)**	**Parallel I/O (New Way)**	**Improvement %**
1	4	1	75%
2	5	2	60%
3	6	3	50%
5	8	5	37%
10	13	10	30%
20	23	20	15%
50	53	50	6%

Note: Jobs that have spinoff data sets benefit more from I/O parallelism since the DSTATs for all spinoff JDS entries in a JDS buffer can be read in simultaneously with the next JDS buffer in the chain.

7.6 OS/390 Version 2 Release 4 JES3 MVS JCL COMMAND Statement

OS/390 Version 2 Release 4 JES3 allows the commands specified on the MVS JCL COMMAND statement in a job stream to be executed.

In-stream operator commands can be entered through:

- JCL command statement

  The JCL command statement can be used to enter an MVS operator command through the input stream only on a JES2 system. The syntax of the in-stream JCL command statement is:

  ```
  //  command  [operands] [comments]
  ```

  An in-stream command is usually executed as soon as it is read. Therefore, the command is not synchronized with the execution of any job or step in the input stream.

- JCL COMMAND statement

  The COMMAND statement is the preferred way to enter in-stream operator commands. JES3 also supports the COMMAND statement starting with OS/390 Version 2 Release 4 JES3. The syntax of the JCL COMMAND statement is:

  ```
  //[name] COMMAND 'command command-operand' [comments]
  ```

  An in-stream command entered through the JCL COMMAND statement is executed during JCL conversion on the system where the JCL conversion takes place. Therefore, the command will not be synchronized with the execution of any job or job step in the input stream.

  Note: On a JES3 system, the system does not record in a job's JESMSGLG data set any commands you enter with the COMMAND statement.

  A COMMAND statement can appear anywhere in the job after the JOB statement.

- JES3 command statement

  The JES3 command statement can be used to enter JES3 operator command through the input stream. The syntax of the JES3 command statement is:

```
//**command-verb[,operand]...
```

JES3 executes an in-stream command as soon as it is read. Therefore, the command will not be synchronized with the execution of any job or step in the input stream.

The *command-verb* on a JES3 command statement can be: CALL, CANCEL, DELAY, DISABLE, ENABLE, ERASE, FAIL, FREE, INQUIRY, MESSAGE, MODIFY, RESTART, SEND, START, SWITCH, or VARY. Do not specify a *DUMP or *RETURN command on a JES3 command statement.

Note: The JES3 command statement must be placed before the first JOB statement in the input stream, if jobs are also being submitted. JES3 treats any JES3 command statements that follow the JOB statement as comment statements.

The COMMAND= and AUTH= parameters on the OS/390 Version 2 Release 4 JES3 CIPARM initialization statement control what jobs can use COMMAND statements and which commands they can issue:

```
CIPARM,PARM=(optionlist),PARMID=id
     [,REGION=nnnnx]
     [,AUTH=ALL|SYS|IO|CONS|INFO]
     [,COMMAND=DISPLAY|EXECUTE|IGNORE|VERIFY]
```

The AUTH=ALL|SYS|IO|CONS|INFO parameter specifies which commands will be accepted through COMMAND JCL statements in the job stream. The groups include:

- SYS - system commands
- IO - input/output commands
- CONS - console commands
- INFO - information commands (such as display)
- ALL - all operator command types

The COMMAND=DISPLAY|EXECUTE|IGNORE|VERIFY parameter specifies the disposition of commands entered through COMMAND JCL statements in the job stream:

- DISPLAY to display and execute the command
- EXECUTE to execute the command
- IGNORE to ignore the command
- VERIFY to display the command and ask the operator whether the command should be executed

The following example shows the JESMSGLG, JESJCL, and JESYSMSG data sets after a job's execution when the job's JCL includes COMMAND statements:

```
 IAT6140 JOB ORIGIN FROM GROUP=ANYLOCAL, DSP=IR , DEVICE=INTRDR  , 0000
IRR010I  USERID VAIN     IS ASSIGNED TO THIS JOB.
 11:08:55 IAT2000 JOB VAINCO   (JOB16323) SELECTED SC49     GRP=A
 11:08:55  ICH70001I VAIN     LAST ACCESS AT 11:08:54 ON TUESDAY, NOVEMBER 18, 1997
 11:08:55  IEF403I VAINCO - STARTED - TIME=11.08.55
 11:08:55  IEF404I VAINCO - ENDED - TIME=11.08.55
//VAINCO JOB '#1','E',MSGLEVEL=1,MSGCLASS=X,NOTIFY=&SYSUID,            *
//  CLASS=A
//*MAIN SYSTEM=SC49
//*      COMMAND  'SEND  ''THIS TEST MESSAGE WILL BE SENT TO USER VAIN
//*            WHEN THIS JOB IS CONVERTED'',USER=(VAIN)'
//SEND   COMMAND 'SEND ''THIS TEST A MESSAGE'',USER=(VAINI)'
//BDT    COMMAND '/I A'
//START  COMMAND 'S DEALLOC'
//TS EXEC PGM=IEFBR14
        1 //VAINCO JOB '#1','E',MSGLEVEL=1,MSGCLASS=X,NOTIFY=&SYSUID,            *
          //  CLASS=A
          //*      COMMAND  'SEND  ''THIS TEST MESSAGE WILL BE SENT TO USER VAIN
          //*            WHEN THIS JOB IS CONVERTED'',USER=(VAIN)'
          IEFC653I SUBSTITUTION JCL - '#1','E',MSGLEVEL=1,MSGCLASS=X,NOTIFY=VAIN,CLASS=A
        2 //SND    COMMAND 'SEND ''THIS TEST A MESSAGE'',USER=(VAINI)'
        3 //BDT    COMMAND '/I A'
        4 //ST     COMMAND 'S DEALLOC'
        5 //TS EXEC PGM=IEFBR14
ICH70001I VAIN     LAST ACCESS AT 11:08:54 ON TUESDAY, NOVEMBER 18, 1997
IEF142I VAINCO TS - STEP WAS EXECUTED - COND CODE 0000
```

Note that the commands entered on the COMMAND statement are not recorded in the JESMSGLG data set.

The next example shows the messages recorded in the SYSLOG for the preceding job execution:

```
N 0000000 SC49     97322 11:08:54.56 JES3     00000290  IRR010I  USERID VAIN     IS ASSIGNED TO THIS JOB.
N 0000000 SC49     97322 11:08:54.62 JES3     00000090  IAT6100 (JOB15337) JOB VAINCO   (JOB16323), PRTY=01, ID=VAIN
N 0020000 SC49     97322 11:08:54.77 JES3CI   00000090  ICH70001I VAIN     LAST ACCESS AT 20:23:25 ON MONDAY, NOVEMBER 17, 1997
MR8000000 SC49     97322 11:08:54.80 JES3CI   00000090  IEFC165I // SEND 'THIS TEST A MESSAGE',USER=(VAINI) 831
 R                                831 00000090
NC0000000 SC49     97322 11:08:54.81 INSTREAM 00000290  SEND 'THIS TEST A MESSAGE',USER=(VAINI)
MR8000000 SC49     97322 11:08:54.82 JES3CI   00000090  IEFC165I // /I A 833
 R                                833 00000090
NC0000000 SC49     97322 11:08:54.83 INSTREAM 00000290  /I A
MR8000000 SC49     97322 11:08:54.84 JES3CI   00000090  IEFC165I // S DEALLOC 835
 R                                835 00000090
N 0000000 SC49     97322 11:08:54.85 JES3     00000090  108548  WTSCPLX3 SC49     MCS CN(28)   -I A
NC0000000 SC49     97322 11:08:54.85 INSTREAM 00000290  S DEALLOC
NR0000000 SC50     97322 11:08:54.89 BDT5O    00000090  BDT8672 BDTJOB JOBNAME  JESJOB   DAP   FROM     TO       ELAPSED
NR0000000 SC50     97322 11:08:54.89 BDT5O    00000090  BDT8673 0001   SNA               SNA.
N 0000000 SC49     97322 11:08:55.04 JES3     00000090  IAT2000 JOB VAINCO   (JOB16323) SELECTED SC49     GRP=A
N 0020000 SC49     97322 11:08:55.05 VAINCO   00000090  ICH70001I VAIN     LAST ACCESS AT 11:08:54 ON TUESDAY, NOVEMBER 17, 1997
N 0000000 SC49     97322 11:08:55.08 VAINCO   00000090  IEF403I VAINCO - STARTED - TIME=11.08.55
N 0000000 SC49     97322 11:08:55.11 VAINCO   00000090  IEF404I VAINCO - ENDED - TIME=11.08.55
N 0000000 SC49     97322 11:08:55.15 JES3     00000090  IAT2001 OS MVS WAITING FOR WORK - SC49
NI0000000 SC49     97322 11:08:55.16 JES3     00000290  SE 'IAT6108 JOB VAINCO   (JOB16323) ENDED, COMP CODE = ZERO ',USER=(VAIN),LOGON
N 0000000 SC49     97322 11:08:55.18 JES3     00000090  IAT6100 ( DEMSEL ) JOB DEALLOC  (JOB16324), PRTY=15, ID=STC
N 0080000 SC49     97322 11:08:55.61 JES3     00000090  ICH70001I STC      LAST ACCESS AT 11:08:55 ON TUESDAY, NOVEMBER
N 0020000 SC49     97322 11:08:55.73 DEALLOC  00000290  IEF695I START DEALLOC  WITH JOBNAME DEALLOC  IS ASSIGNED TO USER
N 0000000 SC49     97322 11:08:55.74 DEALLOC  00000090  IEF403I DEALLOC - STARTED - TIME=11.08.55
N 0000000 SC49     97322 11:08:55.77 DEALLOC  00000090  IEF404I DEALLOC - ENDED - TIME=11.08.55
```

JCL conversion for the example job was done in a C/I FSS. The SYSLOG listing shows that the COMMAND statement commands are issued from the CIFSS address space, and that the command execution is not synchronized with the execution of the job.

Note: Even if the in-stream commands enter the system from JES3 address spaces, the command will be executed with the job submitter's RACF user security token (UTOKEN). The UTOKEN information for the example job includes:

- Port of Entry - INTRDR (JOB was submitted from TSO)
- Submitting User ID - VAIN
- Submitting Group ID - SYS1
- Submitter Node - WTSCPLX3
- Execution Node - WTSCPLX3
- Session type - INTERNAL READER BATCH JOB

7.7 OS/390 Version 2 Release 4 JES3 Statistics

OS/390 Version 2 Release 4 JES3 collects statistics about different JES3 functions such as initialization, connect/restart, and output service. The IATXSTAT macro provides the statistics collection services. These statistics are contained in control blocks called the statistics data area (SDA).

In OS/390 Version 2 Release 4 JES3, statistics are collected for:

- Initialization
  - Total initialization time
  - Times for different phases of initialization (job validation, consoles initialization, RJP initialization)
  - Number of jobs in queue (total, DJC, in output service)
  - I/O counts performed during job validation
- Restart/Connect Processing
  - MDS restart time and number of jobs
  - Total connect time for each system
  - Times for different phases of connect for each system
- Output Service
  - Output service restart time

The statistics data area (SDA), created by using the IATXSTAT macro, can be dumped using the dump core (DC) utility:

```
*S DC OPTION=SDA

*** STATISTICS DATA AREA ENTRIES ***

SDA ENTRY - 09A00460 - INITIALIZATION
00000000-E2C4C1C5 00010000 00000170 C9D5C9E3 *SDAE........INIT*
00000010-C9C1D3C9 E9C1E3C9 D6D54040 40404040 *IALIZATION      *
00000020-40404040 40404040 40404040 00000000 *            ....*
00000030-00000000 00000000 00000000 00000000 *................*
00000040-2001000 00000000 AF949A9E 312AC104 *.........m....A.*
00000050-AF949AA8 FB5D1805 AF949A9E D3C5F503 *.m.y.)...m..LE5.*
00000060-AF949AA0 60A7C206 AF949AA0 624DDD00 *.m..-xB..m...(..*
00000070-AF949AA0 A04EDB04 AF949AA0 A32E1303 *.m...+...m..t...*
00000080-AF949AA6 0ADE0A03 00000000 00000000 *.m.w...........*
00000090-00000000 00000000 AF949AA6 1427F905 *.........m.w..9.*
000000A0-AF949AA8 E14E5906 AF949AA6 15675405 *.m.y.+...m.w....*
000000B0-AF949AA6 25328703 AF949AA6 28F8C105 *.m.w..g..m.w.8A.*
000000C0-AF949AA6 3F155204 AF949AA6 41D2DA04 *.m.w.....m.w.K..*
000000D0-AF949AA6 BCF91504 AF949AA6 BF169903 *.m.w.9...m.w..r.*
000000E0-AF949AA7 F8473005 AF949AA7 F8473505 *.m.x8....m.x8...*
000000F0-AF949AA7 FBC9AF04 AF949AA8 1947E305 *.m.x.I...m.y..T.*
00000100-AF949AA8 9048E300 AF949AA8 F7120406 *.m.y..T..m.y7...*
00000110-AF949AA8 FAF4A103 0007063F 0007209C *.m.y.4.........*
00000120-000719C6 002D0000 000005A4 00000005 *...F.......u....*
00000130-00000084 0000050B 00000007 00000686 *...d...........f*
00000140-00001870 00000686 000005A4 00000000 *.......f...u....*
00000150-000005A4 00000000 000005E3 00000001 *...u.......T....*
00000160-00000744 00000009 000001A1 00000000 *................*
```

The IATYSDA1 macro maps the initialization SDA entry. For example, the highlighted fields tell that last JES3 restart was a hot start, it started at 12:22:25.266348 on 11/18/1997, and ended at 12:22:36.580306 on the same day. There was 1507 jobs in the job queue and one DJC net.

```
SDA ENTRY - 09A005D0 - RESTART (CONNECT)
00000000-E2C4C1C5 00010000 000008E0 D9C5E2E3 *SDAE........REST*
00000010-C1D9E340 4DC3D6D5 D5C5C3E3 5D404040 *ART (CONNECT)   *
00000020-40404040 40404040 40404040 00000000 *            ....*
00000030-00000000 00000000 00000000 00000000 *................*
00000040-AF949AAF 3873CC00 AF949AAF 3DAF4E05 *.m.......m....+.*
00000050-AF949AAF A4393305 00000084 00000000 *.m..u......d....*
00000060-00000000 00000000 00000000 00000000 *................*
00000070-AF949AB0 18F7C200 AF949AAF CAD2FD06 *.m...7B..m...K..*
00000080-00000000 00000000 00000000 00000000 *................*
::::::::-:::::::: :::::::: :::::::: :::::::: ::::::::::::::::::

SDA ENTRY - 09A00EB0 - OUTPUT SERVICE
00000000-E2C4C1C5 00010000 00000050 D6E4E3D7 *SDAE.......&OUTP*
00000010-E4E340E2 C5D9E5C9 C3C54040 40404040 *UT SERVICE      *
00000020-40404040 40404040 40404040 00000000 *            ....*
00000030-00000000 00000000 00000000 00000000 *................*
00000040-AF949AA9 387BA804 AF949AB6 B2333A03 *.m.z.#y..m......*
IAT7921 ISSUE START/CANCEL/RESTART DC REQUEST
```

The IATYSDA2 macro maps the statistics data area entry for restart (connect) related information, and IATYSDA3 correspondingly maps the output service entry. In the example, output service restarted at 12:22:36.830651 and was ready at 12:22:50.960692.

7.8 OS/390 Version 2 Release 4 JES3 64K Job Numbers

In OS/390 Version 2 Release 4 JES3, you may specify a job number range from one to 65534 to be assigned to jobs that enter the JES3 complex. The prior job number limit was 32,767. The OPTIONS initialization statement JOBNO= *joblim* parameter specifies the maximum number (up to 65534) of jobs that may be in the JES3 complex at any given time.

The size of the JCT data set may also set the limit on how many jobs can be in your JES3 complex simultaneously. If the JCT data set is larger than specified on the *joblim* parameter, JES3 uses only that portion required to hold the specified number of jobs. (The remainder of the JCT is used for write error recovery.)

Note: The larger the JCT data set, the longer JES3 initialization takes. For optimum performance, keep the size of the JCT data set as small (*joblim* as low) as possible or as close to the maximum number of jobs required.

64K job number changes to JES3 JES3 initialization statements include:

- OPTIONS JOBNO=({lowest1},{highest9999},{joblim9999})

  Where:

  lowest Specifies a number within the range 1 to 65530 to indicate the lowest job number. This number must be at least 4 job numbers below the value of *highest* in the range.

  highest Specifies a number within the range 5 to 65534 to indicate the highest job number. This number must be at least 4 job numbers higher than the value of *lowest* in the range.

  joblim Specifies the maximum number of jobs that may be in the JES3 complex at any given time (within the range 1 to 65534).

- GROUP JSPAN={nnnnnALL}

  Where:

  nnnnn Specifies the number from 1 to 65534 of jobs in this group to be examined in selecting a job to be scheduled.

- Operator commands
- Messages

7.8.1 Considerations When Using over 32767 Job Numbers

JES3 stores and manipulates job numbers in several formats:

- 2-byte binary job numbers (for example, JCTJOB, RQJOBNO).
- 4-byte binary job numbers (for example, LETJOBNO)
- 4-byte mixed job numbers (for example, JDABJNUM).

The 2-byte job number field has been treated as a 15-bit signed integer in the range of -32767 to 32767, that normally only contains values in the range 0 to 32767. In OS/390 Version 2 Release 4 JES3 all 2-byte job number fields

containing job numbers are 16-bit unsigned integers in the range 0 to 65535. In hexadecimal representation a job number was prior OS/390 Version 2 Release 4 JES3 in the range X'0000' to X'7FFF' and is in OS/390 Version 2 Release 4 JES3 in the range X'0000' to X'FFFF'.

The change from 15-bit signed to 16-bit unsigned integers for 2-byte job numbers must be properly handled in user exits to prevent undesired results. For example, transferring a value from a 2-byte job number field to a 4-byte job number field can be accomplished using the following sequence of instructions:

- The job number rage is between 1 and 32767:

```
Old_Techn DS  OH
          LH    Rx,2_byte_jobno_field    Load 2-byte value
          ST    Rx,4_byte_field          Save as 4-byte value
```

- The job number rage may exceed 32767:

```
Techn_1  DS  OH
         SLR   Rx,Rx                    Clear Rx
         ICM   Rx,3,2_byte_jobno_field  Load 2-byte value in..
*                                       .bytes  2,3 or Rx
         ST    Rx,4_byte_field          Save as 4-byte value

Techn_2  DS  OH
         ICM   Rx,3,2_byte_jobno_field  Load 2-byte value in..
*                                       .bytes 2,3 of Rx
         N     Rx,=AL4(MAXIMUM_JOB_NUMBER_MASK) Clear unused bytes
         ST    Rx,4_byte_field          Save as 4-byte value

Techn_3  DS  OH  - Indexing required to access 2_byte_jobno_field
         LA    Ry,Index_value            Set indexing value
         LH    Rx,2_byte_jobno_field(Ry) Load 2-byte value in..
*                                       .bytes 2,3 of Rx
         N     Rx,=AL4(MAXIMUM_JOB_NUMBER_MASK) Clear unused bytes
         ST    Rx,4_byte_field          Save as 4-byte value

Techn_4  DS  OH
         ICM   Rx,12,2_byte_jobno_field Load 2-byte value in
*                                       .bytes 0,1 of Rx
         SRL   Rx,JOB_NUMBER_SHIFT Shift bytes 0,1 to 2,3 of Rx
         ST    Rx,4_byte_field Store value in 4-byte field

* IATYTVT equates:

MAXIMUM_JOB_NUMBER_MASK EQU X'FFFF'
JOB_NUMBER_SHIFT        EQU 16
```

Comparison of binary job numbers must also be handled differently when 16-bit unsigned integers are used:

```
Old_Compare  DS    OH
             L     Rx,4_byte_field    Load compare value
             :::
             CH    Rx,2_byte_field    Check for equivalence
             BC    EQ,some_where      BifEQ
```

With 16-bit unsigned job numbers, the above technique does not work when the job number is greater than 32767. This is because CH treats the 2_byte_field as a 15-bit signed number and propagates the sign in forming a 4-byte value for comparison. The above technique must be changed to:

```
New_Compar_1 DS OH
             L     Rx,4_byte_field    Load compare value
             :::
             CLM   Rx,3,2_byte_field  Check for equivalence.
             BC    EQ,some_where      BifEQ

New_Compar_2 DS OH - Indexing required to access 2_byte value
             L     Rx,4_byte_field    Load compare value
             LA    Rz,Index_value     Set indexing value
             LH    Ry,2_byte_field(Rz) Load 2_byte value
             N     Ry,=AL4(MAXIMUM_JOB_NUMBER_MAX) Normalize to 0 - 65534
             CR    Rx,Ry              Check for equivalence.
             BC    EQ,some_where      BifEQ
```

Using CLM allows the contents in bytes 2 and 3 of the register and the 2-byte field to be compared as unsigned values in the range of 0 to 65534. Obviously, converting the 16-bit integer to a non-negative 4-byte integer allows for proper comparison, at the expense of using an additional register.

All JES3 code referencing to job numbers and JES3 macros containing job numbers and job counts have been changed. Comments have been added to affected data area mapping fields to make it obvious that these fields should be treated as unsigned halfwords, even though the field definition has not been changed (still is DC H'..'). Note that the only control block change required by the 64K job number support is to the IATYJQX field JQXPCHN (priority chain) which is changed from an array of 16 halfwords to an array of 16 fullwords.

7.9 OS/390 Version 2 Release 4 JES3 Migration

OS/390 Version 1 Release 1 JES3 was the first release of JES3 that allowed the global processor and the local processor to coexist with mixed levels of JES3 as long as the global is at the highest level in the JES3 complex. For example, the global could be at the OS/390 Release 4 JES3 level, and any locals in the JES3 complex can be at any other level down to, and including, MVS/ESA System Product JES3 Version 5 Release 2.1. IBM recommends that installations migrate to the same level of JES3 for all processors as soon as possible (even if it is possible to have a mixture of JES3 levels at or above MVS/ESA SP JES3 5.2.1 in the same complex).

Note that if the JES3 global is at MVS/SP JES3 3.1.3 through MVS/ESA SP JES3 5.2.1, all the locals and the global in the same JES3 complex must be at the same JES3 level.

Table 14 describes the start types that are allowed when first migrating to OS/390 Version 2 Release 4 JES3 from a previous level of JES3.

Releases

Table 14 (Page 1 of 2). OS/390 Release 4 JES3 Global Restart Types Required from Previous

Previous JES3 Release	**JES3 Global Start Types Required for Release 4 Activation**
MVS/ESA SP JES3 5.2.1, OS/390 Release 1 JES3, or OS/390 Release 3 JES3	Hot, DSI, Warm, or Cold
MVS/SP JES3 3.1.3 through MVS/ESA SP JES3 5.1.1	Warm or Cold

Table 14 (Page 2 of 2). OS/390 Release 4 JES3 Global Restart Types Required from Previous	
Previous JES3 Release	**JES3 Global Start Types Required for Release 4 Activation**
Pre-MVS/SP JES3 3.1.3	Cold

Notes:

1. An IPL and CLPA is required for all types of starts on the processor that is to become the OS/390 Version 2 Release 4 JES3 global.

2. A hot start with refresh cannot be performed as the first OS/390 Version 2 Release 4 JES3 initialization.

7.9.1 OS/390 Release 4 JES3 Migration through DSI

Dynamic system interchange (DSI) is the JES3 process that enables a local processor to become the global processor with minimal disruption and without loss of work. Once DSI is complete, JES3 on the old global processor can be reinitialized as a local processor without an intervening IPL.

The following lists the steps required to migrate to OS/390 Release 4 JES3 through DSI:

1. IPL a local processor to run OS/390 Version 2 Release 4.

2. Start OS/390 Release 4 JES3 as a local. The lower-level global will fail the local′s connect attempt with the IAT2640 CONNECT REJECTED FOR local_nm - JES3 INCOMPATIBLE message.

3. Initiate a DSI dialog on the local.

4. Once the DSI is complete, start JES3 on the old global as a local.

The following example shows a console log for a DSI migration dialog. In the example, SC50 system is running OS/390 Release 3 JES3, and a OS/390 Release 4 system is IPLed on SC49. OS/390 Release 4 JES3 system is started on SC49, DSI is initiated on SC49, and finally OS/390 Release 3 JES3 is restarted as a local on SC50:

```
|||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
17.11.24 SC49     VAINI     S JES3,DCEV=
17.11.25 SC49     JES3      IAT3040 STATUS OF JES3 PROCESSORS IN COMPLEX
17.11.25 SC49     JES3      IAT3040 L06RS007(  ), L06RS008(  ), SC50    (UP), SC49    (  )
17.11.25 SC49     JES3      IAT3040 SC43    (  )
17.11.25 SC49     JES3      IAT3011 SPECIFY JES3 START TYPE
17.11.25 SC49     JES3     *088 IAT3011 (C, L, H, HA, HR, HAR, W, WA, WR, WAR, OR CANCEL)
17.11.36 SC49     VAINI     88,L
17.11.36 SC49               IEE600I REPLY TO 088 IS;L
17.11.36 SC49     JES3      IAT4030 0001 SPOOL DATA SET  IN USE
17.11.36 SC49     JESXCF    IXZ0001I CONNECTION TO JESXCF COMPONENT ESTABLISHED,
17.11.36 SC49     JESXCF             GROUP WTSCPLX3 MEMBER SC49
17.11.38 SC49     JES3      IAT3100 JES3 OS240    LOCAL   START ON 1997.323 AS SC49
17.11.39 SC50     JES3      IAT2640 CONNECT REJECTED FOR SC49     - JES3 INCOMPATIBLE
17.13.02 SC49     VAINI    *X DSI
17.13.02 SC49     JES3     *IAT0915 DSI - REVIEW LOCAL DSI PROCEDURE FOR SC49
17.13.10 SC49     VAINI    *S DSI
17.13.10 SC49     JES3     *IAT0910 DSI - DISABLE OLD GLOBAL PROCESSOR
17.13.02 SC49     VAINI    RO SC50,*X DSI
17.13.30 SC50     JES3     *IAT0920 DSI - CHECK GLOBAL DSI PROCEDURE FOR SC50
17.13.42 SC49     VAINI    RO SC50,*S DSI
17.13.42 SC50     JES3      IAT6370  FSS CI7  ABEND  REQUESTED BY DSI  REJECTED - MAIN SC49 NOT CONNECTED
17.13.42 SC50     JES3      IAT0905 DSI - STARTED FOR SC50
17.13.43 SC50     IEESYSAS  IXZ0003I CONNECTION TO JESXCF COMPONENT BROKEN
17.13.43 SC50     IEESYSAS           GROUP WTSCPLX3 MEMBER SC50
17.13.43 SC50     JES3      IEF404I JES3 - ENDED - TIME=17.13.43
17.13.54 SC49     VAINI    *S DSI
17.13.54 SC49     JES3     *IAT0900 DSI - SWITCH GLOBAL DEVICES
```

```
17.13.58 SC49     VAINI    *S DSI
17.13.58 SC49     JES3      IAT0905 DSI - STARTED FOR SC49
17.13.59 SC49     IEESYSAS  IXZ0003I CONNECTION TO JESXCF COMPONENT BROKEN
17.13.59 SC49     IEESYSAS           GROUP WTSCPLX3 MEMBER SC49
17.13.59 SC49     JES3      IAT3040 STATUS OF JES3 PROCESSORS IN COMPLEX
17.13.59 SC49     JES3      IAT3040 LO6RS007(  ), LO6RS008(  ), SC50    (DS), SC49    <DS>
17.13.59 SC49     JES3      IAT3040 SC43    (  )
17.14.00 SC49     JES3      IAT4030 0001 SPOOL DATA SET  IN USE
17.14.00 SC49     JES3      IAT4075 MAXIMUM NUMBER OF JOBS IS LIMITED TO 0006686 BY JCT DATA SET SIZE
17.14.07 SC49     JESXCF    IXZ0001I CONNECTION TO JESXCF COMPONENT ESTABLISHED,
17.14.07 SC49     JESXCF             GROUP WTSCPLX3 MEMBER SC49
17.14.57 SC49     IEESYSAS  IEA630I  OPERATOR SYSJ3R02 NOW ACTIVE,   SYSTEM=SC49    , LU=JESXCF
17.14.58 SC49     IEESYSAS  IEA630I  OPERATOR SYSJ3N02 NOW ACTIVE,   SYSTEM=SC49    , LU=JESXCF
17.14.59 SC49     JES3     *IAT3100 JES3 OS240 SYSTEM  HOTSTART ON 1997.323 AS SC49
17.14.59 SC50     IEESYSAS  IEA631I  OPERATOR SYSJ3N03 NOW INACTIVE, SYSTEM=SC50    , LU=JESXCF
17.14.59 SC49     JES3     *IAT7921 ISSUE START/CANCEL/RESTART DC REQUEST
17.14.59 SC50     IEESYSAS  IEA631I  OPERATOR SYSJ3R03 NOW INACTIVE, SYSTEM=SC50    , LU=JESXCF
17.15.10 SC49     VAINI    *S JSS
17.15.16 SC49     JES3      IAT2002 LSTOR=32000 ALLOC=000000000 AVAIL=32000 SC49
17.15.16 SC49     JES3      IAT2645 ****** SC49     CONNECT COMPLETE ******
17.15.17 SC49     JES3      S JES3CI.CI7,,,(JES3,9,IATINFC)
17.15.17 SC49     JES3      IAT6372 FSS CI7      START-UP ON SC49     SUCCESSFUL
17.15.17 SC49     JES3CI    IEF403I JES3CI - STARTED - TIME=17.15.17
17.15.18 SC49     JESXCF    IXZ0001I CONNECTION TO JESXCF COMPONENT ESTABLISHED,
17.15.18 SC49     JESXCF             GROUP WTSCPLX3 MEMBER CI7
17.15.19 SC49     JES3      IAT4452 CI FSS CI7      IS READY FOR WORK
17.15.38 SC49     VAINI    RO SC50,S JES3
17.15.38 SC50     JES3      IAT3040 STATUS OF JES3 PROCESSORS IN COMPLEX
17.15.38 SC50     JES3      IAT3040 LO6RS007(  ), LO6RS008(  ), SC50    <DS>, SC49    (UP)
17.15.38 SC50     JES3      IAT3040 SC43    (  )
17.15.38 SC50     JES3      IAT3011 SPECIFY JES3 START TYPE
17.15.38 SC50     JES3     *942 IAT3011 (C, L, H, HA, W, WA, WR, WAR, OR CANCEL)
17.16.02 SC49     VAINI     942,L
17.16.03 SC50     JESXCF    IXZ0001I CONNECTION TO JESXCF COMPONENT ESTABLISHED,
17.16.03 SC50     JESXCF             GROUP WTSCPLX3 MEMBER SC50
17.16.04 SC50     JES3      IAT3100 JES3 OS130    LOCAL   START ON 1997.323 AS SC50
17.16.10 SC49     JES3      IAT2007 GMS CONNECT - JOB VAINI    (JOB16338) - DEMAND SELECT    - SC50
:::::::::::::     ::::      ::::::: ::: :::::::: - ::: ::::::   :::::::::: - :::::: ::::::    - ::::
17.16.11 SC49     JES3      IAT2007 GMS CONNECT - JOB VTAM44   (JOB15794) - DEMAND SELECT    - SC50
17.16.11 SC49     JES3      IAT2007 GMS CONNECT - JOB SYSLOG   (JOB15793) - DEMAND SELECT    - SC50
17.16.11 SC49     JES3      IAT2002 LSTOR=32000 ALLOC=000000000 AVAIL=32000 SC50
17.16.11 SC49     JES3      IAT4452 CI FSS CI6      IS READY FOR WORK
17.16.11 SC49     JES3      IAT4452 CI FSS CI5      IS READY FOR WORK
17.16.11 SC50     JES3      IAT2645 ****** SC50     CONNECT COMPLETE ******
||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
```

7.9.2 Migration Considerations for 64K Job Numbers

Before the 64K job number support can be activated, all processors in the JES3 complex must be at the OS/390 Release 4 JES3 level. There are no coexistence PTFs for prior releases of JES3 to tolerate a higher than 32K job number range. Thus you should not use 64K job numbers unless you are sure that you will not fall back to a prior release of JES3.

Once all systems in the JES3 complex are at the OS/390 Release 4 JES3level, you can *decrease the low job number or increase the high job number* over a hot start with refresh.

You cannot delete job numbers from the job number range specified on the JOBNO parameter during a hot start with refresh. This requires a warm start. Therefore, at least a JES3 warm start is needed if you need to fall back to a prior release of JES3 and greater than 32K job numbers were being used.

7.9.3 Suggested Migration Scenario

A migration scenario from a JES3 complex currently running a level of JES3 prior to OS/390 Release 4 JES3 might be:

- Migrate all JES3 systems in your complex to at least the MVS/ESA SP JES3 5.2.1 level, if not already at that level.

- Migrate at least one JES3 processor to OS/390 Version 2 Release 4 JES3 global by the allowable restart type, as shown in Table 14 on page 178.
- Migrate all JES3 locals to the OS/390 Release 4 JES3 level prior to:
  - 64K job number and activation
  - Segmenting the initialization stream
- Activate 64K job number support as required. Note that fallback to a pre-OS/390 Release 4 JES3 is not supported and may cause some functions to fail if done.

  The 64K job numbers may require a larger JCT data set than is currently in use. The JCT data set must be allocated as a single extent data set on cylinder boundaries. *OS/390 JES3 Initialization and Tuning Guide* contains information on how to calculate the size of a JCT entry and determine the size of the JCT data set.

  A 3380 device holds 39 unmodified JCT entries on a track, which amounts to 585 entries per cylinder. For a 3390 device, the numbers are 42 entries per track and 630 entries per cylinder. Using these numbers, once you have determined the number of jobs that you wish to allow in the JES3 complex concurrently, you can calculate the size of the JCT data set.

  Note: If you have segmented the JES3 initialization stream, you can only hot start the pre-OS/390 Release 4 JES3 as the global. For warm or cold starts, you should have an initialization stream without INCLUDE initialization statements.

  Note: A pre-OS/390 Release 4 JES3 tolerates the OS/390 Release 4 JES3 CIPARM initialization statement parameters AUTH= and COMMAND= during a hot start. However, it ignores the MVS JCL COMMAND statements in a job′s JCL. You can only get the JCL COMMAND statements processed on a OS/390 Release 4 JES3 system.

7.9.4 OS/390 Release 4 JES3 Customization and Operation Notes

This section describes some of the actions that the JES3 system programmer may be required to perform to ensure that JES3 user exits and modifications are properly customized for OS/390 Version 2 Release 4 JES3.

7.9.4.1 64K Job Numbers Note

Handling 2-byte job numbers internally as 16-bit unsigned integers in the 64K job number support allows a large number of JES3 control blocks to remain unchanged and eliminates a cold start that would be otherwise required to migrate to the new release. Additionally this allows prior releases of JES3 to coexist in a JES3 with OS/390 Release 4 JES3, as long as the job number range is contained between 1 and 32767.

Installations should change their JES3 user exits and modifications to handle the 2-byte job number fields as 16-bit unsigned numbers even if they do not plan to use job numbers greater than 32767. This includes examining all JES3 user exits and other user modifications for occurrences of macros:

- AJOBNUM (Job number macro)
- IATXJCT (JCT access services)
- IATXJOB (Convert job number and job ID)
- IATXJQE (Obtain the address of a JQE)
- IATXSDM (Invoke spool initialization routine)
- NCBTAFND (Find NCB by jobname or Job ID)

These macros call services with job numbers as input and thus are usually referencing fields with job numbers.

7.9.4.2 User Initialization Statement Note

The hot start with refresh support has changed the OS/390 Release 4 JES3 initialization processing in many areas.

During JES3 initialization, many initialization statements are converted into internal form and written to a file on spool. Each spool file is assigned a file identification (ID) through the IATYSPL macro, which contains a file identification to locate information on subsequent hot starts. Prior to this release, JES3 assigns these file IDs based on location of IATYSPL within the IATYINT macro. Remember that the IATYSPL macro is used on the ITREADs and ITWRITEs to access intermediate spool text files.

In OS/390 Release 4 JES3, a new ID= parameter on IATYSPL is used to assign the file IDs that are defined in the IATYITXT macro. The IATYITXT contains 10 file IDs (24 through 2D) for user initialization statements.

7.9.4.3 IWASPOUT Macro Note

The IWASPOUT JES3 macro, used to write messages during JES3 initialization, allows a new COMPRESS=Y|N parameter to be specified. The COMPRESS=Y option requests blank compression for initialization messages. The default is COMPRESS=N. OS/390 Release 4 JES3 has started to use COMPRESS=Y on many of its own IWASPOUT macros.

7.9.4.4 User Exit 15 Note

In OS/390 Release 4 JES3, user Exit 15 (IATUX15), Examine an Initialization Statement, is called in two additional environments. The first exit point is when a *MODIFY,CONFIG command is being executed. The second is when a hot start with refresh is being executed. Exit 15 can perform the same processing in these two cases as in the prior releases of JES3. The statement can be allowed, or it can be discarded, to be later processed by JES3 when control is returned from the user exit.

7.9.4.5 AWAIT Macro Note

The AWAIT JES3 macro, used to wait for completion of one or more events in JES3, has been modified for OS/390 Release 4 JES3 by adding the DATA= parameter on the invocation of the macro. This allows up to 16 bytes of data to be associated with a particular AWAIT request. Through a combination of the existing REASON= parameter and the added DATA= parameter, you can more easily determine the AWAIT event that the FCT is waiting to occur.

The exploiters of the enhanced AWAIT macro are:

- Spool services, to specify the name of control block being read or written (for example, JDAB, JDS, OSE)
- IATXCSF service, to put the name of the function being invoked as a subtask (for example, ALOAD-modname, VERIFY-UNLOAD)
- AENQ service, to put the name of the resource being serialized

The formatted AWAIT data can be viewed in the IPCS VEBX JES3 'OPTION=FCT' output, along with the AWAIT reason code.

7.9.4.6 IATXCSF Macro Note

The IATXCSF macro, used to call a JES3 subtask function, has the DESC= keyword added to it. The DESC= keyword allows a description of the function being invoked as a subtask to be described. The data can be 16 bytes in length.

7.9.4.7 ACALL Macro Note

The ACALL macro, used to call another section of executable code, has two additional changes to it in OS/390 Release 4 JES3.

- A new sub-parameter on the EP parameter to generate an ADCON during the in-line expansion. With the A sub-parameter, an inline ADCON is generated:

```
          ACALL ENTER=(DMREADIO,A)
+         L     R15,=A(DMREADIO)
+         L     R14,ASAVE
+         BALR  R14,R14
```

- A new EPLOC parameter specifies the field that contains the address of the target routine:

```
          ACALL EPLOC=VIWRDRTN
+         L     R15,VIWRDRTN
+         L     R14,ASAVE
+         BALR  R14,R14
          :::::
VIWRDRTN DC    A(Read_routine_address)
```

7.9.4.8 IATYGLOB Macro Note

The IATYGLOB mapping macro, new as of OS/390 Release 4, contains constants that identify a variety of product levels, maintenance levels, and so on. Many of these fields existed in prior releases of JES3 in macro IATYEQU, but OS/390 Release 4 JES3 moves them to IATYGLOB. The following fields are included in to the OS/390 Release 4 JES3 IATYGLOB macro:

Table 15 (Page 1 of 2). IATYGLOB Macro Fields

Field Name	Variable Type	Value	Definition
&J3LEVEL	Character	OS 2.4.0	Current release
&J3REL	Character	OS240	JES3 release
&J3PLVL	Arithmetic	8	Product level for HJS6604
&J3FMID	Character	HJS6604	FMID of current release
&J3_OS390_ PROGNO	Character	5647-A01	Program number
&J3_SUPP_ RELEASE_COUNT	Arithmetic	3	Number of supported down-level releases
&J3_SUPP_ RELEASE1	Character	HJS5521	Supported release
&J3_SUPP_ RELEASE2	Character	HJS6601	Supported release
&J3_SUPP_ RELEASE3	Character	HJS6603	Supported release

Table 15 (Page 2 of 2). IATYGLOB Macro Fields			
Field Name	**Variable Type**	**Value**	**Definition**
&J3_REL_YEAR	Character	1997	Year release is available

7.9.4.9 Configuration Serialization Note

Due to the changed serialization against the JES3 configuration information in OS/390 Release 4 JES3, additional operator procedures are required when JES3 cannot obtain the serialization. The serialization can fail either from the global or the local. The following describes various recovery procedures that might be used depending on the ENQ circumstances:

1. ENQ fails from the global.

   If the global cannot obtain the ENQ during a hot start with refresh, JES3 will issue the following two messages:

   ```
   IAT3072 mainname WAITING FOR EXCLUSIVE USE OF THE CONFIGURATION
   IAT3073 ISSUE "CANCEL" TO CANCEL WAIT FOR mainname
   ```

   These messages remain highlighted, and are issued after waiting 15 seconds for the ENQ. CANCEL is allowed so that the JES3 operator can cancel the hot start with refresh and perform a JES3 hot start.

2. ENQ fails from the local JES3 or C/I FSS address space.

   If the local or C/I FSS cannot obtain the ENQ during a local or C/I FSS start, JES3 will issue the following message:

   ```
   IAT3072 xxxx WAITING FOR SHARED USE OF THE CONFIGURATION
   ```

   This message is not highlighted, but is retrievable through the DISPLAY R,L command. It is issued after waiting 60 seconds for the ENQ. There is no CANCEL allowed, because there is no alternative.

Depending on the specific situations, the following procedures can help in determining the correct response for the situation:

- To find out who has the ENQ, issue either:

  ```
  D GRS,RES=(SYSZIAT,*)
  D GRS,RES=(SYSZIAT,CONFIG.CHANGE*)
  ```

The reasons why the global might not be able to get the ENQ are as follows:

- The JES3 local is in the process of being started.-

  In this case, make sure no IAT3011 messages for locals are outstanding. If you have previously replied to message IAT3011, then the local might be initializing or hung.

  Reply CANCEL and perform a hot start of the JES3 global, or issue the FORCE command to cancel JES3 on the local.

- C/I FSS is in process of being started.

  The FSS might be attempting to request a JES3 global service (for example, dynamic allocation or SYSOUT allocation). This can occur if the FSS is just starting when JES3 is terminated. In this case, reply CANCEL and perform a hot start, or issue the CANCEL command to cancel the C/I FSS address space.

7.9.4.10 Problem Determination and Diagnosis Notes

OS/390 Release 4 JES3 provides additional aids for problem determination and diagnosis.

Dump Core Enhancements: The dump core utility has been enhanced in this release to view information contained in the spool or checkpoint data sets relating to JES3 initialization.

OS/390 Release 4 JES3 adds option ITX= to the *START DC command to view intermediate text files either completely or in a particular initialization statement or sub-statement. Because of the volume of data that might be generated, directing the output to a specific printer (OUT=) is recommended on the Dump Core command.

An example:

```
||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
    *S,DC,OPTION=(ITX=CIPARM)

    *** INTERMEDIATE TEXT - CIPARM

    RECORD NUMBER = 0000001 - CIPARM
    00000000-F0F1F4F0 F0F1F4F4 F0F0F0F9 F9F9F1F1 *0140014400099911*
    0000001O-C5F0F0F0 F1F1C1F0 F9F9F9D2 00000000 *E00011A0999K....*
    0000002O-000000                              *...            *

    RECORD NUMBER = 0000002 - CIPARM
    00000000-F0F2F4F0 F0F1F4F4 F0F0F0F9 F9F9F1F1 *0240014400099911*
    0000001O-C5F0F0F0 F1F1C5F0 F9F9F9D2 00000000 *E00011E0999K....*
    0000002O-000000                              *...            *
||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
```

The highlighted field in the example output, (F1) is the command disposition. Prior to OS/390 Release 4 JES3, this field was always set to 3 by JES3, indicating that the MVS converter is to ignore commands read from the input stream. In OS/390 Release 4 JES3, the values for the command disposition are:

0 Execute command

1 Display and execute command

2 Display and request disposition

3 Ignore command.

Dump Initialization Checkpoint Record: To dump the Initialization Checkpoint Record that contains the checkpointed portions of various control blocks such as the TVT, and SVT use the following dump code option:

```
*S,DC,OPTION=ICP
```

7.9.4.11 JES3 Monitoring Facility Notes

JMF FCT I/O Statistics Display: JES3 Monitoring Facility (JMF) has been enhanced in OS/390 Release 4 JES3 to report I/O information that is collected for each FCT. This includes:

- Number of multi-record file read I/Os
- Number of multi-record file write I/Os
- Number of single-record file read I/Os
- Number of single-record file buffers read
- Number of single-record file write I/Os
- Number of single-record file buffers written

JMF formats only non-zero values for the statistics. JMF also adds a dashed line between the FCT displays to make it easier to read, as in the following example:

An example:

```
||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
--------------------------------------------------------------------------------------------------------------------------------------
5  DSP NAME IS JSAM         DEVICE NAME IS **NONE**     FCT PRTY IS 250
          FCT ON FCT CHAIN                100.00 % OF SAMPLES
          FCT IN NUC MODE                 100.00 % OF SAMPLES                        100.00 % OF FCT
          FCT IN AUX MODE                    .00 % OF SAMPLES                           .00 % OF FCT
  I/O ACTIVITY
          SRF READ I/O'S    =       187
          SRF READ BUFFERS  =       187
          SRF WRITE I/O'S   =       189
          SRF WRITE BUFFERS =       189
  ACTIVITY UNDER IATNUC TASK
          ::::::::::::::::::              ::::::::::::::::     ::::::::::::::::   ::::::::::::::
||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||
```

JMF AWAIT Analysis: JMF uses the AWAIT reason codes specified on the AWAIT macro to determine what event the FCT is waiting for. The specific code to determine AWAITs has been removed, and AWAIT reason codes have been added to many AWAITs. As a result, the AWAIT report has a more detailed explanation of why JES3 is AWAITing. Examples:

```
AWAIT IS WAITING FOR JESREAD I/O COMPLETION - JBT
AWAIT IS WAITING FOR AWRITE I/O COMPLETION - JDS
AWAIT IS WAITING FOR JESREAD I/O COMPLETION - ISP
```

The default AWAIT= value on *CALL JMF has been changed from 15 to 45.

JMF Bottleneck Report: The same bottlenecks are reported as in prior releases of JES3. However, the description text is reformatted to allow for the longer AWAIT descriptions. The AWAIT reason codes themselves are used to identify bottlenecks; JMF no longer has specific JMF code to identify which AWAITs are bottlenecks.

JMF SMF Record Changes: The JMF SMF 84 record is changed to be compatible with previous releases. A new version number (3) is used because of the changes:

- FCT section - has an FCT extension that is between the FCT section and the AWAIT section that follows. This extension contains the FCT I/O statistics.
- AWAIT section - sections that do not have reason codes are unchanged. AWAITs that do have reason codes are mapped differently. Flags that identify the type of AWAIT are no longer set.
- AWAIT bottleneck section - the format of the AWAIT bottleneck section has been changed to contain more information.

7.9.4.12 SMF 43 Record Note

OS/390 Release 4 JES3 includes a small changes to the SMF 43 record type (JES SMF SUBSYSTEM START RECORD). Bit SMF43REF in byte SMF43RST (Start Record Flags) has been added to indicate a hot start with refresh was requested.

Chapter 8. Standalone Dump Enhancements

The performance and usability of standalone dump (SADMP) have been enhanced in release 4 of OS/390. The enhancements have been made in

- Setup of dump data sets
- Operation of SADMP
- Copying dump data set
- SADMP program changes

The above enhancements assist a site in reducing the system outage time caused by the need to take an SADMP.

The Release 4 enhancements along with previous release enhancements are making SADMP easier to use and also reducing the impact on a production environment caused by the length of time it takes to perform an SADMP.

One such previous enhancement was the ability to dump more than 64K tracks of data over multiple volumes. An SADMP of 86,000 tracks from one system were dumped over 3 dump data sets on different volumes during the writing of this redbook. This is documented here as some sites have reported not being able to do this.

8.1 Setup of Dump Data Sets

The first enhancement is in the allocation and formatting of dump data sets. Allocation and formatting is performed by the REXX utility AMDSADDD. AMDSADDD can be found in SYS1.SAMPLIB and can be run using TSO/E or a batch job. In release 4 of OS/390 AMDSADDD only formats the first line of an dump data set when it is being defined or cleared enabling much faster formatting of dump data sets.

For example, formatting of a 400 cylinder dump data set running the utility from TSO/E took 3 seconds. Running the same formatting using OS/390 Release 3 SADMP took 80 seconds.

SADMP formats the remainder of the dump data set as it now performs format writes instead of update writes. Format writes have the additional benefit of providing better performance to Redundant Array of Independent Disks (RAID) devices.

Figure 93 on page 188 shows the allocation of a new dump data set running AMDSADDD from TSO/E. This allocates a 400 cylinder cataloged dump data set called 'SYS1.SC64.SADMP' on a 3390 with a volser of SBOX04.

Figure 94 on page 189 shows the same allocation using a batch job.

```
 Enter TSO or Workstation commands below:

 ===> ex 'sys1.samplib(amdsaddd)'

 What function do you want?
 Please enter DEFINE if you want to allocate a new dump data set
 Please enter CLEAR if you want to clear an existing dump data set
 Please enter REALLOC if you want to reallocate and clear an existing
        dump data set
 Please enter QUIT if you want to leave this procedure
 Please enter REALLOC if you want to reallocate and clear an existing
        dump data set
 Please enter QUIT if you want to leave this procedure
define

 Please enter VOLSER or VOLSER(dump_data set_name)
sbox04(sys1.sc64.sadmp)
 Please enter the device type for the dump data set
 Device type choices are 3380 or 3390 or 9345
3390
 Please enter the number of cylinders
400
 Do you want the dump data set to be cataloged?
 Please respond Y or N
y
 IKJ56650I TIME-05:05:33 PM. CPU-00:00:10 SERVICE-276202 SESSION-01:04:16
 MAY 15,1997

 Initializing output dump data set with a null record:
 Dump data set has been successfully initialized

 Results of the DEFINE request:
 Dump data set Name   : SYS1.SC64.SADMP
 Volume               : SBOX04
 Device Type          : 3390
 Allocated Amount     : 400
```

Figure 93. TSO/E screen output from AMDSADDD

AMDSADDD will prompt you for:

- Type of function
- Volser or Volser and dump data set name

  If you provide a dump data set name, it must have the text “SADMP” as either part of the name or as an entire data set qualifier.

  If you do not provide a dump data set name, AMDSADDD will use the name SYS1.SADMP
- Device type
- Number of cylinders
- Whether you want the dump data set cataloged or not.

None of the fields will default except the dump data set name on the volser field.

```
//FERGMICS JOB (1111,2222),'MIKE',CLASS=A,MSGCLASS=X,
//         NOTIFY=FERGMIC,MSGLEVEL=(1,1)
//*
//S1     EXEC  PGM=IKJEFT01
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 EXEC 'SYS1.SAMPLIB(AMDSADDD)' 'DEFINE SBOX04(SYS1.SC64.SADMP) 3390 -
400 Y'
/*


READY
 EXEC 'SYS1.SAMPLIB(AMDSADDD)' 'DEFINE SBOX04(SYS1.SC64.SADMP) 3390 400 Y'
IKJ56650I TIME-05:23:26 PM. CPU-00:00:00 SERVICE-4417 SESSION-00:00:01 MAY
15,1997
Initializing output dump data set with a null record:
Dump data set has been successfully initialized

Results of the DEFINE request:

  Dump data set Name   : SYS1.SC64.SADMP
  Volume               : SBOX04
  Device Type          : 3390
  Allocated Amount     : 400

READY
END
```

Figure 94. JCL and Output from AMDSADDD

The batch invocation requires all the parameters to be entered in the SYSTSIN with a space between each parameter. The "-" is used as the continuation character in the SYSTSIN.

Note: Dump data sets formatted with OS/390 Release 4 SADMP *cannot* be used to take a dump on an OS/390 Release 3 system or below. Dump data sets formatted with OS390 Release 3 SADMP and below "can" be used to take a dump on an OS/390 Release 4 system.

8.2 Creation of the SADMP program

SADMP can be created on either a DASD volume or a tape volume. The size of SADMP program has increased from previous releases. The JCL in Figure 95 on page 190 creates SADMP on DASD volume 2049 using the default space allocation of 588 4K blocks. The IPL parameter provides the DASD device address for placement of the SADMP program. The OUTPUT parameter provides the address of the first volume that will be used for the dump. If using a 4 digit address for IPL or OUTPUT parameters, it must be preceded by a slash to distinguish it from a device type.

To specify a tape device for the SADMP program or first output device, specify a "T" instead of a "D" on the IPL or OUTPUT parameter.

Figure 95 on page 190 creates SADMP program:

- On DASD device address 2049 with a volser of SBOX09

- SADMP will dump the minimum virtual storage (LSQA and selected system subpools) for all of the address spaces
- Use 8A0, a 3278 as the console
- Look for the first dump data set on DASD device address 204A.

```
//FERGMICS JOB (1111,2222),'MIKE',CLASS=A,MSGCLASS=X,
//         NOTIFY=FERGMIC,MSGLEVEL=(1,1)
//*
//S1     EXEC  PGM=AMDSAOSG
//SYSLIB    DD DSN=SYS1.MACLIB,DISP=SHR
//          DD DSN=SYS1.MODGEN,DISP=SHR
//GENPRINT  DD SYSOUT=*
//GENPARMS  DD *
          AMDSADMP IPL=D/2049,                                    X
               VOLSER=SBOX09,                                     X
               MINASID=ALL,                                       X
               CONSOLE=(8A0,3278),                                X
               OUTPUT=D/204A
           END
/*
```

Figure 95. JCL to create SADMP program on DASD

SADMP is created in a data set called SYS1.PAGEDUMP.

8.3 Operation of SADMP

SADMP is initiated by performing an IPL from the device address where SYS1.PAGEDUMP resides. This is the same as previous releases of the SADMP program.

The messages issued by SADMP have been changed to provide more meaningful information. There is also a new message issued during the virtual dump phase.

The first changed message you may see is

```
AMD096A 204A SBOX10 SYS1.SADMP
        IS VALID, HOWEVER IT MAY ALREADY CONTAIN DATA FROM A PREVIOUS DUMP
        05/15/97 16:15:45 TITLE: TEST SADMP OVER 64K TRACKS AND MULTIVOLUME
        REPLY 'R' TO REUSE OR 'U' TO USE ANOTHER DATA SET.
```

This message now provides the title and information from the “dummy header record” to assist you in deciding whether or not this dump data set should be overwritten.

The next changed message is AMD005I. This message has 3 new variations.

```
AMD005I DUMPING OF REAL STORAGE COMPLETED (MINIMAL)
```

This indicates address spaces 1 to 4 and common storage have been dumped.

```
AMD005I DUMPING OF REAL STORAGE COMPLETED (SUMMARY)
```

This indicates that a set of selected system address spaces such as consoles and catalog as well as the current primary, secondary and home spaces of the actively executing work on all of the processors have been dumped.

```
AMD005I DUMPING OF REAL STORAGE COMPLETED (IN_USE)
```

This indicates that all other allocated storage in the system has been dumped.

Following the above messages you will see:

```
AMD095I REAL DUMP XX% COMPLETED. TOTAL MEGABYTES DUMPED: YY
AMD095I VIRTUAL DUMP XX% COMPLETED. TOTAL MEGABYTES DUMPED: YY
```

XX is the percentage of address spaces processed out of the total number of address spaces in the system being dumped. YY is the number of megabytes. You may see multiple of these messages issued where XX does not change but YY does.

The frequency that these messages are issued has been reduced to 30 seconds. In previous releases the messages could have been issued as often as every 2 seconds.

There is a new message AMD108I which provides information on the progress of the virtual dump phase.

```
AMD108I DUMPING OF SUMMARY ADDRESS SPACES COMPLETED
AMD108I DUMPING OF SWAPPED IN ADDRESS SPACES COMPLETED
```

If during the dump the dump data set fills up you will receive the following message as in previous releases prompting you for another device address which contains a dump data set.

```
AMD001A SPECIFY OUTPUT DEVICE ADDRESS(N)
```

“N” is a file sequence number. For example, if this is the second device you are dumping to then N would be 2.

Another usability enhancement has been made by dumping real storage in a prioritized order. This attempts to capture the most useful diagnostic data first so if the SADMP is terminated before it completes, either by an operator or error, the partial captured dump will be more likely to contain data needed to diagnose the original failure.

8.4 Copying Dump Data Sets

The recommendation in previous releases of SADMP to use the Interactive Problem Control System (IPCS) COPYDUMP command to copy dump data sets has now been removed. If a dump data set needs to be copied to another system to be viewed, perhaps because IPCS is not installed on the failing system, IEBGENER is quite acceptable. Although IEBGENER could be used in the past, it was recommended against.

8.5 SADMP Program Changes

A number of changes have been made in SADMP some of which have already been mentioned to improve the elapsed time that takes to perform an SADMP. The remainder of performance improvements are

- Various instruction path lengths have been reduced in the virtual dump phase
- The size of SADMP working storage has been doubled

This allows longer Channel-Command Word (CCW) chains to be built which will reduce the number of missed output CCW chain tack ons which in turn will increase the output data rate.

- Defer reading non-page-table pages and sort them into page slot order while performing the virtual phase of the dump process

  This allows:

  – Multiple pages to be read with one I/O

  – Reduced overhead locating areas of paged-out storage

  – Paged out storage to be read in DASD physical order

    Which has the effect of:

    - Avoiding DASD seeks where possible

    - Avoiding the need to locate records.

- Read a whole DASD track when a single page must be read immediately

  This yields up to 50% performance improvement of the virtual dump phase when there are a lot of page-table pages to be read from DASD.

- Various instruction path length reductions in the virtual dump phase

8.6 Migration in a Parallel Sysplex

As mentioned in section 8.1, “Setup of Dump Data Sets” on page 187 dump data sets created using OS/390 release 4 AMDSADDD REXX utility cannot be used by an OS/390 Release 3 or below SADMP program. However an OS/390 Release 3 or below formatted dump data set can be used by the OS/390 Release 4 SADMP program.

In a Parallel Sysplex environment the best approach is to:

- Continue to use your existing formatted dump data sets
- Have 2 versions of the SADMP program
  1. An existing version
  2. An OS/390 Release 4 version

As each system in the sysplex is migrated to OS/390 Release 4 operators will need to be advised to use the OS/390 Release 4 version of SADMP.

Once all systems in the Sysplex have been migrated, you can discard your older version of the SADMP program. If more dump data sets are required these can now be initialised with the OS/390 release 4 version of AMDSADDD.

Chapter 9. Transaction Dump

Transaction dump is a new mechanism for obtaining multiple, concurrent, unformatted, problem state dumps. Transaction dump is similar to SYSMDUMP and can be thought of as SYSMDUMP with a user interface in the form of extensions to the existing functions. SYSMDUMP still functions in the same manner as it has in the past.

The need for this type of diagnostic tool has increased with the introduction of the Parallel Syplex environment where dumps from a number of different address spaces on different systems in the sysplex are required.

Transaction dump provides the following additional function over SYSMDUMP:

- Automatic allocation of data sets
- Dump title
- Dump index
- Symptom record
- Incident token
- Remote dump
- Asynchronous processing
- Completion notification
- Supervisor state

9.1 IEATDUMP Macro

IEATDUMP is a new macro introduced in OS/390 release 4 to request a transaction dump. This macro issues SVC 51. Although this macro can be issued by unauthorized callers, use of the “remote” parameter is restricted to authorized callers. That is callers in supervisor state, PSW key 0-7 or APF-authorized.

9.1.1 IEATDUMP parameters

The parameters are divided into 3 groups.

1. Required
2. Optional
3. Parameters with defaults

A suffix of *AD* on a parameter indicates that, if the parameter is specified as an address in a register, the register contains a pointer field. For example, HDRAD. Where the parameter has a corresponding parameter without the AD suffix, and the parameter is specified as an address in a register, the register contains an address of a character field. For example, the corresponding parameter of HDRAD is HDR.

The following parameters are required:

- DCBAD, DCB, DSNAD or DSN

This parameter describes the output data set for the transaction dump. The transaction dump can be written to either a preallocated or a dynamically allocated data set. If a preallocated data set is used the DCBAD or DCB parameter is used and provides a previously opened DCB. If the data set is to be dynamically allocated the DSNAD or DSN parameter is used and provides a name pattern for the data set.

- HDRAD or HDR

  This parameter provides the dump title.

The following parameters are optional:

- IDXAD or IDX

  This parameter provides the name of the dump index data set. The data set must be the name of an existing IPCS dump directory.

- SYMRECAD or SYMREC
- INTOKENAD or INTOKEN
- REMOTEAD or REMOTE
- PROBDESCAD or PROBDESC
- LISTAD or LIST
- SUBPLSTAD or SUBPLST
- DSPLISTAD or DSPLIST
- ECBAD or ECB
- RETCODE
- RSNCODE

The following parameters have default values:

- ASYNC

  The default is *ASYNC=NO*

- MF

  The default is *MF=S*

- PLISTVER

  The default is *PLISTVER=IMPLIED_VERSION*

- SDATA

  The default is *SDATA=DEFS*

9.2 Automatic Allocation of Data Sets

Transaction dump processing stores data in dump data sets that you pre-allocate manually, or that are allocated automatically, as needed. Transaction dump processing supports automatic allocation of dump data sets at the time the system writes the dump to DASD. The dump is allocated from the generic resource SYSALLDA. When the system writes a dump, it allocates a data set of the correct size. Automatic allocation is not multi-volume or striped. If automatic allocation fails, message IEA820I is issued, and the dump is deleted.

9.2.1 Dump Title

The application has control of the name of the data sets created by the automatic allocation function, and you can select a name-pattern to allo for dump data set organization according to your needs. The name is determined through an installation-supplied pattern on the DSN(AD) keywo in the IEATDUMP macro. A set of symbols is available so that you can include the following kinds of information in the names of your automatically allocated dump data sets:

- System name
- Sysplex name
- Job name
- Local and GMT time and date

You can specify a name-pattern to generate any name acceptable under normal OS/390 data set name standards.

To allocate dump data sets automatically, the caller′s and/or DUMPSRV address space must have authority to allocate new data sets. Do the following:

- Associate the caller′s and/or DUMPSRV address space with a user ID.
- Authorize caller′s or DUMPSRV user ID to create new dump data sets using a naming convention.

The SUMDUMP or SUM option on the IEATDUMP macro causes Transaction dump to produce a summary dump, which can be formatted using the IPCS SUMDUMP subcommand. IBM strongly recommends that you view the SUMDUMP output prior to investigating the usual portions of the dump. Each summary dump record is indicated by the header ″----tttt---- RECORD ID X′nnnn′,″ where tttt is the title for the type of SUMDUMP output, and nnnn is the hexadecimal record identifier assigned to the type.

Note: During Transaction dump processing, the system sets some tasks in the requested address space non-dispatchable; non-dispatchable tasks in the dump may have been dispatchable at the time of the problem.

9.2.2 Dump Index

The dump index record is a dump directory record describing the dump. This record is written to the dump directory allocated to IPCSDDR DD statement in the current job step.

The summary dump is formatted by the IPCS VERBEXIT SUMDUMP subcommand and has an index which describes what the summary contains.

Chapter 10. Enhancements in TSO/E

Before OS/390 Release 4, the following commands only worked on the system on which they were issued by a user, operator, or system programmer:

- The operator SEND command
- The TSO/E SEND command
- The TSO/E PARMLIB command
- The TSO/E CALL and TEST commands

With OS/390 Release 4, these commands also work in a Parallel Sysplex. The following sections describe how they operate and their expected results.

10.1 Operator SEND Command

The operator SEND command is an authorized command. It is used to:

- Send a message to one or more terminal users
- Save a message in the SYS1.BRODCAST data set
- List, delete, or send a particular message from the notices section of SYS1.BRODCAST
- List all messages in the notices section of SYS1.BRODCAST
- Communicate with console operators and other terminals in operator mode or extended MCS console mode

In OS/390 Release 4, this command has been enhanced by a new parameter called ROUTE which allows the system operator to send a message to all logged-on users on any system that belongs to one sysplex.

This parameter is optional and can have one of the following values:

***** The asterisk causes the message being sent to be routed to all active systems in the sysplex. This is the default when the parameter ALL is specified.

systemname This specifies the system name (1 to 8 characters) that receives and processes the command.

groupname This specifies that the command is to be routed to a named subset of systems in the sysplex.

For information on how to define group names in a sysplex, see *OS/390 MVS Planning: Operations*.

Note: Messages to be routed to a *specific* user or a named number of users find their way within a sysplex by invoking JESXCF.

Using the operator SEND command to send this message to all currently logged on users in a sysplex, type in:

```
SEND 'System IPL will be done at 7:00 pm',ROUTE=*
```

Using the operator SEND command to send a message to all logged-on users on a specific system in the sysplex (for example, SC63), type in:

```
SEND 'System IPL will be done at 7:00 pm',ROUTE=SC63
```

Using the operator SEND command to send a message to four specific users who are currently logged-on, type in:

```
SEND 'Please quit TSO immediately',USER=(JOHN,PAUL,GEORGE,RINGO)
```

In the last example, it does not matter on which of the available TSO systems in the sysplex the four users are logged on. The SEND command finds them and routes the message to them.

10.2 TSO/E Send Command

The TSO/E SEND command can be used by a TSO/E user to send messages to

- Another TSO terminal user
- A specific operator
- The master console operator
- A specific operator console.

In OS/390 Release 4, this command allows you to send messages independent of which system within the sysplex the user is logged on to. Therefore, as a normal end user, you do not need to know on which specific OS/390 system the user you want to send a message to is logged on. TSO/E automatically determines where the user is logged on within the sysplex and routes the message to that system by invoking JESXCF.

The syntax of the TSO/E SEND command has not been changed, which means that you can use the command as you already know it. For example, if you specify the following:

```
SEND 'please send the updated data sets to my userid' USER(FRITZ JOE)
```

then both users can be on different systems in the sysplex and still receive the message if they are logged on. system.

Figure 96 on page 199 shows how this command finds its way through the sysplex.

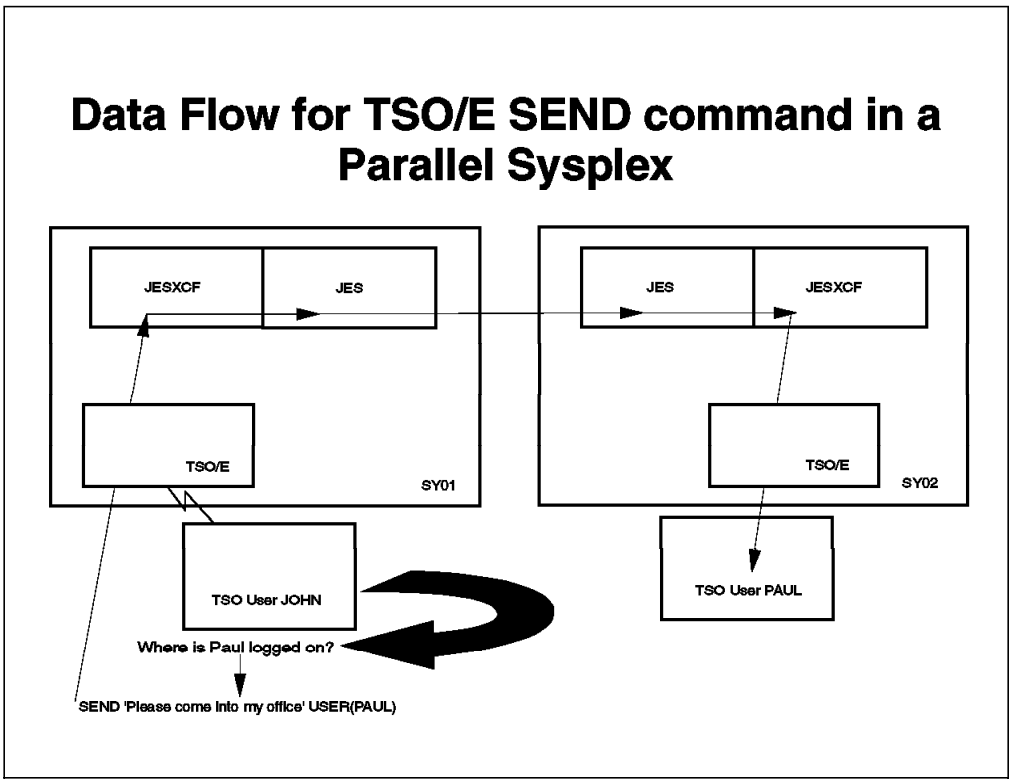

Figure 96. Data Flow for TSO/E SEND Command in a Parallel Sysplex

10.3 TSO/E PARMLIB Command

The PARMLIB command is an authorized command used to:

- Display the specifications in the active IKJTSOxx member of SYS1.PARMLIB
- Dynamically change the active member without a new IPL
- Check the syntax of any IKJTSOxx member of SYS1.PARMLIB

In OS/390 Release 4, the TSO/E PARMLIB command has been enhanced to list or update anything on all systems within a Parallel Sysplex.

This is done via the new ROUTE parameter. This parameter is optional and can have one (and only one) of the following values, which must be in parentheses:

***** The asterisk specifies that the command will be routed to all active systems in the sysplex.

systemname This specifies the system name (1 to 8 characters) that will receive and process the command.

groupname This specifies that the command will be routed to a named subset of systems in the sysplex.

For information on how to define group names in a sysplex, see *OS/390 MVS Planning: Operations*.

An example of the PARMLIB LIST command in the sysplex looks like the following, and it shows the system programmer all information from member(s) IKJTSOnn in SYS1.PARMLIB:

```
PARMLIB LIST ROUTE(*)
```

The next example shows you the table of authorized commands of all systems within a defined sysplex.

```
PARMLIB LIST(AUTHCMD) ROUTE(*)
```

With the following example, you can change the active IKJTSOxx member in SYS1.PARMLIB to IKJTSO02 for system SC63:

```
PARMLIB UPDATE(02) ROUTE(SC63)
```

This makes a change from IKJTSO00 to IKJTSO02 on the running TSO/E system.

Figure 97 shows the messages your receive after entering the TSO/E PARMLIB command. The changed parameter here is the allowance of the user SEND command from ON to OFF, issued from an authorized user (for example, an operator).

```
 IKJ712I DEFAULT VALUES WERE USED FOR ALLOCATE
 IKJ712I DEFAULT VALUES WERE USED FOR HELP
 IKJ712I DEFAULT VALUES WERE USED FOR TEST
 IKJ712I DEFAULT VALUES WERE USED FOR PLATCMD
 IKJ712I DEFAULT VALUES WERE USED FOR PLATPGM
 IKJ55121I SYSTEM DEFAULTS WERE UPDATED ON SYSTEM SC63 USING PARMLIB
MEMBER IKJTSO02 OF DATASET SYS1.PARMLIB
```

Figure 97. TSO/E PARMLIB UPDATE(SEND) ROUTE(SC63) Command

The operator for example issues the command on system SC64 and routes it via the TSO/E PARMLIB command to system SC63, which is also within the sysplex.

The result after issuing the

```
PARMLIB LIST(SEND) ROUTE(*)
```

command can be seen in Figure 101 on page 204.

The ROUTE parameter is only valid for the TSO/E PARMLIB LIST and UPDATE commands, and not for the TSO/E PARMLIB CHECK command.

When the update of multiple systems takes place and one or more systems fail, then no backout will be done on the other systems that completed that TSO/E PARMLIB UPDATE command successfully.

Responses to a TSO/E PARMLIB LIST command are grouped , to give you a better understanding of the command responses from a number of systems. In Figure 98 on page 201, Figure 99 on page 202 and Figure 100 on page 203 you can see that all systems that have the same parmlib settings are shown in one output screen.

If there are equal settings within a sysplex, the result looks like the output shown in Figure 98 on page 201.

```
parmlib list route(*)
 TSO/E PARMLIB SETTINGS :

   SYS1.PARMLIB(IKJTSO00) on volume SBOX01
   Activated by **IPL** on 1997-05-08 at 23:44:43 from system SC63
   Applies to :    SC63

   SYS1.PARMLIB(IKJTSO00) on volume SBOX01
   Activated by **IPL** on 1997-05-09 at 00:15:05 from system SC64
   Applies to :    SC64



           THE FOLLOWING ARE THE OPTIONS FOR THE ALLOCATE STATEMENT:
                 DEFAULT   (OLD)


           CURRENT SETTINGS FOR CONSOLE:

           INITUNUM              9999
           INITSNUM              9999
           MAXUNUM              39999
           MAXSNUM              19999



           CURRENT SETTINGS FOR HELP:

           LANGUAGE: ENU  DATASET(S):
                            SYS1.HELP


           THE FOLLOWING ARE THE PARMLIB OPTIONS FOR SEND:

           OPERSEND(ON)
           USERSEND(ON)
           SAVE(ON)
           CHKBROD(OFF)
           LOGNAME(BRODCAST)
           USEBROD(ON)
           MSGPROTECT(OFF)
           SYSPLEXSHR(ON)
           OPERSEWAIT(ON)
           USERLOGSIZE(1,2)

           CURRENT SETTINGS FOR TEST:

           TSO COMMANDS SUPPORTED:
```

Figure 98. TSO/E PARMLIB LIST ROUTE(*) Command with Equal Settings in a Sysplex

```
TEST SUBCOMMANDS/LOAD MODULES SUPPORTED:


THE FOLLOWING ARE THE PARMLIB OPTIONS FOR TRANSREC:

DAPREFIX: USERID    SPOOLCL:  B          CIPHER:  YES
VIO:      VIO       LOGSEL:   LOG        LOGNAME: MISC
OUTWARN:  100000    ,100000              SYSOUT:  *
OUTLIM:   9999999
USRCTL:   NAMES.TEXT
SYSCTL:   CTS.NAMES.TEXT


THE FOLLOWING ARE THE NODE/SMF PAIRS FOR TRANSREC:
WTSCPLX1/  SC42     WTSCPLX9/  SC43     WTSCPLX1/  SC47
WTSCPLX3/  SC49     WTSCPLX3/  SC50     WTSCPLX1/  SC52
WTSCPLX1/  SC53     WTSCPLX1/  SC54     WTSCPLX1/  SC55

THERE ARE NO PLATCMD COMMANDS DEFINED TO THE SYSTEM


THERE ARE NO PLATPGM PROGRAMS DEFINED TO THE SYSTEM
CURRENT PARMLIB SETTINGS FOR AUTHCMD

ADYOPCMD  ADA       EXITEDIT  RACTRACE  RACTR     RECEIVE
TRANSMIT  XMIT      LISTB     LISTBC    SE        SEND
PARMLIB   IKJPRMLB  TESTAUTH  TESTA     RMM       MVPXDISP
LIBLIST   Q         QCBXA     QCBESA    DITTO     DITTOA
TL        RACONVRT  SYNC      LISTD     LISTDS    RACONVRT
CONSPROF  AD        ADDSD     AG        ADDGROUP  AU
ADDUSER   ALG       ALTGROUP  ALD       ALTDSD    ALU
ALTUSER   BLKUPD    CO        CONNECT   DD        DELDSD
DG        DELGROUP  DU        DELUSER   LD        LISTDSD
LG        LISTGRP   LISTDATA  LU        LISTUSER  RALT
RALTER    RDEF      RDEFINE   RDEL      RDELETE   RE
REMOVE    RL        RLIST     RACLINK   RVARY     PASSWORD
PW        PE        PERMIT    SHCDS     SETCACHE  SETR
```

(continued)

Figure 99. TSO/E PARMLIB LIST ROUTE(*) Command with Equal Settings in a Sysplex

```
          SETROPTS  SINFO     SR        SEARCH    BACKUP    OPINFO
          OPSTAT    SRSTAT    WSSTAT    FREMIGID  HCEMCS    ARMDREG
          IRRDPI00

          CURRENT PARMLIB SETTINGS FOR AUTHPGM

          AMASPZAP  CC        EXITEDIT  IEBCOPY   IEBGENER  IFCEREP1
          Q         QCBXA     QCBESA    CLUPGTRN  CLUPGTPR  CLUPVTPD
          CLUPGLOC  EDGHSKP   EDGUTIL   EDGRPTD   EDGAUD    ICHDSM00
          ICHUT100  ICHUT200  ICHUT400  IRRDSC00  IRRUT100  IRRUT200
          IRRUT400  ITPENTER  CSFDAUTH  IXCMIAPU  SINFO     IEARELCN
          ARMDREG   IRRDPI00

          CURRENT PARMLIB SETTINGS FOR AUTHTSF

          IEBCOPY   CLUPGTRN  CLUPGTPR  ICQASLI0  ICQASLC0  CLUPVTPD
          CLUPGLOC  IHVUUSD   ITPENTER  CSFDAUTH  EQQMINOR  EDGHSKP
          EDGUTIL   EDGRPTD   EDGAUD    IKJEFF76


          CURRENT PARMLIB SETTINGS FOR NOTBKGND

          OPER      OPERATOR  TERM      TERMINAL
READY
```

PARMLIB LIST ROUTE(*) Command with Equal Settings in a Sysplex (continued)

Figure 100. TSO/E

If there are different settings within a sysplex, the result looks like the output shown in Figure 101 on page 204.

```
parmlib list(send) route(*)
 TSO/E PARMLIB SETTINGS :

   SYS1.PARMLIB(IKJTSO02) on volume SBOX01
   Activated by HARJANS on 1997-05-22 at 17:06:14 from system SC64
   Applies to :    SC63



           THE FOLLOWING ARE THE PARMLIB OPTIONS FOR SEND:

           OPERSEND(ON)
           USERSEND(OFF)               <== difference
           SAVE(ON)
           CHKBROD(OFF)
           LOGNAME(BRODCAST)
           USEBROD(ON)
           MSGPROTECT(OFF)
           SYSPLEXSHR(ON)
           OPERSEWAIT(ON)
           USERLOGSIZE(1,2)

   ----------------------------------- ----------------------------

   SYS1.PARMLIB(IKJTSO00) on volume SBOX01
   Activated by HARJANS on 1997-05-22 at 17:05:26 from system SC64
   Applies to :    SC64



           THE FOLLOWING ARE THE PARMLIB OPTIONS FOR SEND:
           OPERSEND(ON)
           USERSEND(ON)                <== difference
           SAVE(ON)
           CHKBROD(OFF)
           LOGNAME(BRODCAST)
           USEBROD(ON)
           MSGPROTECT(OFF)
           SYSPLEXSHR(ON)
           OPERSEWAIT(ON)
           USERLOGSIZE(1,2)
```

Figure 101. TSO/E PARMLIB LIST ROUTE(*) Command with Different Settings in a Sysplex

In a single system, the output can look as seen in Figure 102 on page 205 and Figure 103 on page 206.

```
parmlib list

           PARMLIB INFORMATION FROM MEMBER IKJTSOOO


           THE FOLLOWING ARE THE OPTIONS FOR THE ALLOCATE STATEMENT:
                 DEFAULT   (OLD)


           CURRENT SETTINGS FOR CONSOLE:

           INITUNUM              9999
           INITSNUM              9999
           MAXUNUM              39999
           MAXSNUM              19999

           CURRENT SETTINGS FOR HELP:

           LANGUAGE: ENU  DATASET(S):
                           SYS1.HELP


           THE FOLLOWING ARE THE PARMLIB OPTIONS FOR SEND:

           OPERSEND(ON)
           USERSEND(ON)
           SAVE(ON)
           CHKBROD(OFF)
           LOGNAME(BRODCAST)
           USEBROD(ON)
           MSGPROTECT(OFF)
           SYSPLEXSHR(ON)
           OPERSEWAIT(ON)
           USERLOGSIZE(1,2)

           CURRENT SETTINGS FOR TEST:

           TSO COMMANDS SUPPORTED:


           TEST SUBCOMMANDS/LOAD MODULES SUPPORTED:
           THE FOLLOWING ARE THE PARMLIB OPTIONS FOR TRANSREC:

           DAPREFIX: USERID    SPOOLCL:  B          CIPHER:  YES
           VIO:      VIO       LOGSEL:   LOG        LOGNAME: MISC
           OUTWARN:  100000    ,100000              SYSOUT:  *
           OUTLIM:   9999999
           USRCTL:   NAMES.TEXT
           SYSCTL:   CTS.NAMES.TEXT
```

Figure 102. TSO/E PARMLIB LIST Command in a Single System Environment

```
THE FOLLOWING ARE THE NODE/SMF PAIRS FOR TRANSREC:
WTSCPLX1/  SC42     WTSCPLX9/  SC43     WTSCPLX1/  SC47
WTSCPLX3/  SC49     WTSCPLX3/  SC50     WTSCPLX1/  SC52
WTSCPLX1/  SC53     WTSCPLX1/  SC54     WTSCPLX1/  SC55
THERE ARE NO PLATCMD COMMANDS DEFINED TO THE SYSTEM


THERE ARE NO PLATPGM PROGRAMS DEFINED TO THE SYSTEM


CURRENT PARMLIB SETTINGS FOR AUTHCMD

ADYOPCMD  ADA       EXITEDIT  RACTRACE  RACTR     RECEIVE
TRANSMIT  XMIT      LISTB     LISTBC    SE        SEND
PARMLIB   IKJPRMLB  TESTAUTH  TESTA     RMM       MVPXDISP
LIBLIST   Q         QCBXA     QCBESA    DITTO     DITTOA
TL        RACONVRT  SYNC      LISTD     LISTDS    RACONVRT
CONSPROF  AD        ADDSD     AG        ADDGROUP  AU
ADDUSER   ALG       ALTGROUP  ALD       ALTDSD    ALU
ALTUSER   BLKUPD    CO        CONNECT   DD        DELDSD
DG        DELGROUP  DU        DELUSER   LD        LISTDSD
LG        LISTGRP   LISTDATA  LU        LISTUSER  RALT
RALTER    RDEF      RDEFINE   RDEL      RDELETE   RE
REMOVE    RL        RLIST     RACLINK   RVARY     PASSWORD
PW        PE        PERMIT    SHCDS     SETCACHE  SETR
SETROPTS  SINFO     SR        SEARCH    BACKUP    OPINFO
OPSTAT    SRSTAT    WSSTAT    FREMIGID  HCEMCS    ARMDREG
IRRDPI00

CURRENT PARMLIB SETTINGS FOR AUTHPGM

AMASPZAP  CC        EXITEDIT  IEBCOPY   IEBGENER  IFCEREP1
Q         QCBXA     QCBESA    CLUPGTRN  CLUPGTPR  CLUPVTPD
CLUPGLOC  EDGHSKP   EDGUTIL   EDGRPTD   EDGAUD    ICHDSM00

ICHUT100  ICHUT200  ICHUT400  IRRDSC00  IRRUT100  IRRUT200
IRRUT400  ITPENTER  CSFDAUTH  IXCMIAPU  SINFO     IEARELCN
ARMDREG   IRRDPI00

CURRENT PARMLIB SETTINGS FOR AUTHTSF

IEBCOPY   CLUPGTRN  CLUPGTPR  ICQASLI0  ICQASLC0  CLUPVTPD
CLUPGLOC  IHVUUSD   ITPENTER  CSFDAUTH  EQQMINOR  EDGHSKP
EDGUTIL   EDGRPTD   EDGAUD    IKJEFF76


CURRENT PARMLIB SETTINGS FOR NOTBKGND

OPER      OPERATOR  TERM      TERMINAL
```

Figure 103. TSO/E PARMLIB LIST Command in a Single System Environment (continued)

In this case, you have to enter:

```
PARMLIB LIST
```

You do not need the ROUTE parameter because the command only works locally.

10.4 TSO/E CALL and TEST Commands

The TSO/E CALL command is used to load and execute a program that exists in an executable form. This program can be written by users, or it may be a system module such as a compiler, sort or utility program.

The TSO/E TEST command is used to test a program, command processor, or APPC/MVS transaction program for proper execution, and to locate programming errors. The program must be a member of a partitioned data set (PDS) or a partitioned data set extended (PDSE).

Prior to OS/390 Release 4, users had to keep the following points in mind when using the TSO/E CALL or TEST commands:

- They had to know the full data set name.
- The data set name of a load library had to be hard-coded in products, command lists (CLISTs), and executable files (for example, REXX files).
- The installation had to provide at least RACF READ access to those data sets.

With OS/390 Release 4, the TSO/E CALL and TEST commands are enhanced to use the standard load module search sequence if an asterisk (*) is specified by a user, instead of the full data set name.

Now an example of the CALL command looks like:

```
CALL  *(iebcopy)
```

If you are using the TEST command it looks as follows:

```
TEST  *(member)
```

Note: A module residing in the linklist chain and, for example, in the MLPA or PLPA cannot be tested in the way described; you have to also specify the related data set.

10.5 PARMLIB Concatenation

This feature allows you to specify those PARMLIB data sets that make a logical concatenation, similar to LPALSTxx and LNKLSTxx support for LPA and linklist libraries. This allows more flexibility in managing members that belong to the PARMLIB.

The PARMLIB command processor uses the new logical PARMLIB service called IEFPRMLB to dynamically allocate, unallocate, and read the logical PARMLIB. Therefore, it provides easier handling of changes that may occur in case of upgrading product level. It also allows you to separate your data from IBM supplied data.

10.6 TSO/E RECEVIVE Command Enhancements

The TSO/E RECEIVE command is used to receive transmitted files, which come over a network or from another user on the same system, and restore them to their original format.

There is one enhancement in the TSO/E part of OS/390 Release 4 which might be important in a case where RECFM=U is specified (meaning that when data sets with undefined record length are to be received). In that case, the specified partitioned data set fits into one extent.

The TSO/E RECEIVE command works by multiplying the data set size by factor 4.

10.7 REXX Enhancements

REXX is a general purpose interpreted programming language. REXX has the usual structured-programming instructions: IF, SELECT, DO WHILE, LEAV, for example. REXX also has a number of useful built-in functions which are ready to use.

One of these built-in functions is the DATE() function which returns, after calling (by default), the local date in the format *dd mon yyyy* (day, month, year, for example: 12 May 1997), with no leading zero or blank on the day.

You can also specify other parameters on the DATE function to obtain specific formats, for example:

C This means means century and returns the number of days, including the current day, since and including January 1 of the last year that is a multiple of 100 in the format: *ddddd* (no leading zeros).

For example, a call to DATE(C) on 12 March 1997 returns 35561, which is the number of days from 1 January 1900 to 12 March 1997.

Similarly, a call to DATE(C) on 2 January 2000 returns 2, the number of days from 1 January 2000 to 2 January 2000.

E This means European and returns the date in the format: *dd/mm/yy*. A call to DATE(E) on May 12 1997 returns 12/05/97 as a result.

J This means Julian and returns the date in the format: *yyddd*. So a call to DATE(J) on May 12 1997 returns 97132 as a result.

M This means Month and returns the full English name of the current month, for example: August.

With the TSO/E part of OS/390 Release 4, the DATE() function of REXX is enhanced so you can submit it with two parameters for conversion between two formats. For example, you can make a call such as:

```
DATE('E','97/05/12','O')
```

As a result of this ordered input date, you will get the European date format *12/05/97.*

Now consider another example. We use this function to convert the input string from the standard format *yyyymmdd* to a more user-readable format, for example, the one used in the United States:

```
DATE('U','19970512','S')
```

As a result of this standard input date, you will get the date format *05/12/97.*

Finally, let us look at a small REXX program which uses the DATE input as a variable and converts it to a more date-like format:

```
/* A small REXX Program to show the updated DATE() Function */

say 'enter date in format yyyymmdd'
say'The converted result looks like this:'
say''
pull datein
say(DATE('E',datein,'S'))
say''
exit
```

The result of this example is shown in Figure 104.

```
exec 'harjans.test1.rexx(date)' exec

 enter date in format yyyymmdd
19970508
 The converted result looks like this:
 08/05/97

 READY
```

Figure 104. REXX DATE() Function Example

The default values for both parameters are Normal. That means the format is *dd mon yyyy*.

10.8 Renaming Samples

In the SYS1.SAMPLIB you will find some TSO-related members where the names have changed in OS/390 Release 4 to meet the naming convention standard beginning with I*.

The members and their meaning are shown in Table 16.

Table 16 (Page 1 of 2). Changed SYS1.SAMPLIB Member Names

Old Name	New Name	Meaning
BCSTSMPE	IKJBSMPE	IKJEBLMT: BRODCAST limit CSECT. Function: The macro IKJBCAST generates a CSECT, named IKJEBLMT, that is used by the TSO RIM, IKJEFXSR, to set the number of records that are to be set aside for the SYS1.BRODCAST data set.
EDITSMPE	IKJESMPE	IKJEBEPD and IKJEBINS: EDIT data set characteristic CSECTS. Function: The macro IKJEDIT generates CSECT, IKJEBINS, that is linkedited with the EDIT command in SYS1.CMDLIB.
TSOSMPE	IKJTSMPE	IKJEFLP1: LOGON limit CSECT. Function: The macro IKJTSO generates CSECT, IKJEFLP1, that is used by the LOGON modules.

Table 16 (Page 2 of 2). Changed SYS1.SAMPLIB Member Names		
Old Name	**New Name**	**Meaning**
TSOANCH	IRXTSMPE	Usermod to change the number of entries in the IRXANCHR table to support a larger number of concurrent environments.
TSOREXX1	IRXREXX1	TSO/E REXX System Default Parameter Module example.
TSOREXX2	IRXREXX2	TSO/E REXX TSO Environment Parameter Module example.
TSOREXX3	IRXREXX3	TSO/E REXX ISPF Environment Parameter Module example.
TSOREXX4	IRXREXX4	TSO/E REXX Compiler Programming Table Module example.

10.9 User Log Allocation

For the SEND command of IKJTSOxx member in SYS1.PARMLIB, a new keyword is introduced in OS/390 Release 4 called USERLOGSIZE. This new keyword allows you to specify the default space allocation for the user log.

The USERLOGSIZE keyword has the following three parameters:

primary-quantity This specifies the number of (primary) tracks to be allocated for the user log. The default is 1 track.

second-quantity This specifies the number of (secondary) additional tracks to be allocated for the user log, if even more space is needed. The default is 2 tracks.

dir-block This specifies the number of directory blocks needed in the directory if the user log data set is to be a partitioned data set. The default is 20 directory blocks.

Following is an example of how the statement in the IKJTSOxx member can look: shown below:

```
SEND  LOGNAME(MYOWNLOG) USERLOGSIZE(5,3)
```

This means that a userlog will be created under the name userid.MYOWNLOG as a sequential file with a primary size of 5 tracks and a secondary size of 3 tracks. For additional information on this topic, see *OS/390 V2R4.0 TSO/E Customization*.

Therefore, for example, if you issue the PARMLIB LIST command with parameter SEND; you will get output as shown in Figure 105 on page 211.

```
parmlib list(send)
 TSO/E PARMLIB SETTINGS :

   SYS1.PARMLIB(IKJTSO00) on volume SBOX01
   Activated by **IPL** on 1997-05-16 at 15:40:29 from system SC63
   Applies to :    SC63



           THE FOLLOWING ARE THE PARMLIB OPTIONS FOR SEND:

           OPERSEND(ON)
           USERSEND(ON)
           SAVE(ON)
           CHKBROD(OFF)
           LOGNAME(BRODCAST)
           USEBROD(ON)
           MSGPROTECT(OFF)
           SYSPLEXSHR(ON)
           OPERSEWAIT(ON)
           USERLOGSIZE(1,2)
```

Figure 105. PARMLIB LIST(SEND) Output with USERLOGSIZE

10.10 ALLOCATE Command Enhancement

In previous releases, TSO/E only accepted day-of-year values in a range of 001 and 366 for the expiration date parameter (EXPDT) of the ALLOCATE command. That was not consistent with processing via job control language.

Now the value of day-of-year can be between 000 and 366 for the format *yyddd* or *yyyy/ddd*.

You can, for example, code this command:

```
alloc da(mytest.data) new recfm(f) unit(sysda) lrecl(80) blksize(800)
expdt(98000)
```

As a result, you get in option 3.4 of ISPF the following information, shown in Figure 106 on page 212.

```
                              Data Set Information
Command ===>

Data Set Name  . . . : HARJANS.MYTEST.DATA

General Data                          Current Allocation
 Volume serial . . . : SBOX02          Allocated blocks  . : 39
 Device type . . . . : 3390            Allocated extents . : 1
 Organization  . . . : NONE
 Record format . . . : F
 Record length . . . : 80
 Block size  . . . . : 800            Current Utilization
 1st extent blocks . : 39              Used blocks . . . . : 0
 Secondary blocks  . : 24              Used extents  . . . : 0

 Creation date . . . : 1997/05/17
 Expiration date . . : 1998/01/00
```

Figure 106. Data Set Information with Expiration Date

10.11 TMP Conditional Disposition Processing

In the past, when using entry point IKJEFT1A or IKJEFT1B (members of SYS1.LPALIB) in a TSO/E batch job, the conditional disposition specification for files allocated via SVC99 or DD statement in job control language has been ignored. The result was that data sets were kept even if their disposition requested that they should be deleted after a user abend.

With OS/390 Release 4, this is changed so that conditional disposition becomes effective only when the TSO/E batch job abends.

10.12 Dynamic LPA Toleration Support

With Dynamic LPA, authorized users can perform special functions against the Link Pack Area (LPA). You can read more about this in Chapter 5, “Dynamic LPA” on page 123. Previous to OS/390 Release 4, you had to do an IPL before LPA changes became effective.

This support in TSO/E makes any update immediately effective to a TSO/E module which resides in the LPA even when a new version of this module is added to the LPA. This is done with the LPA ADD statement in member PROGxx in SYS1.PARMLIB.

One example of this function is to apply maintenance to modules residing in the LPA and to make these changes effective without an IPL.

Chapter 11. DFSMShsm Enhancements

The following enhancements are now available with DFSMShsm Version 1 Release 4:

- DFSMShsm duplex tape
- IDCAMS alter without RECALL
- ABARS improvements
- CDS record level sharing
- Dump analysis elimination
- Serviceability improvements

11.1 DFSMShsm Duplex Tape

The duplex tape option provides an alternative to TAPECOPY processing for backup and migration cartridge-type tapes. Two tapes are created concurrently with one designated the original, the other the alternate. The intent is that the original tape be kept onsite, while the alternate can either be created in a remote tape library, or taken offsite.

Using the SETSYS DUPLEX command, you can specify that the duplex tape option be used for backup tapes only, migration tapes only, or both backup and migration tapes. If you use duplex creation for both migration and backup and you have been auto scheduling tape copying, you probably want to remove the auto scheduling. If you are duplexing either migration or backup tapes but not both, you can auto schedule tape copies based on the type indicated in the ARC0421I message.

Now, with Version 1 Release 4, DFSMShsm creates two HSM-owned tapes at creation time. It is an alternative to the TAPECOPY method, and the tapes used are compatible to those created by TAPECOPY.

This mechanism is very useful in customer installations where automatic tape libraries, such as the IBM 3494 or IBM 3495, are to be used. This is needed when it is necessary to prevent backed-up or migrated data sets from having two identical copies of the data sets.

Routing of the original and alternate tapes can be done to different SMS managed libraries via automatic class selection routines (also known as ACS routines).

11.1.1 System Definitions

Two new SETSYS DUPLEX keywords are added to the ARCCMDnn member of SYS1.PARMLIB. They are:

BACKUP Specifies whether duplex tape is created for backup datasets under HSM control

The options are:

(Y) SETSYS DUPLEX(BACKUP(Y)) specifies that duplex tape must be done for backup.

(N) SETSYS DUPLEX(BACKUP(N)) specifies that duplex tape must not be done for backup. This is the default.

MIGRATE Specifies whether duplex tape is created for migration data sets under HSM control

The options are:

(Y) SETSYS DUPLEX(MIGRATION(Y)) specifies that duplex tape must be done for migration.

(N) SETSYS DUPLEX(MIGRATION(N)) specifies that duplex tape must not be done for migration. This is the default.

11.1.2 Duplex Tape Data Set Names

The following names are used as original tape data set names:

```
prefix.HMIGTAPE.DATASET
prefix.BACKTAPE.DATASET
```

The following names are used as alternate tape data set names:

```
prefix.COPY.HMIGTAPE.DATASET
prefix.COPY.BACKTAPE.DATASET
```

If you are entering the DFSMShsm DELVOL command for a tape migration or backup volume, both the original and the alternate volume are returned as scratch.

The DELVOL command reassigns the volume as not managed or owned by DFSMShsm. An example of the DELVOL command for a tape migration level 2 volume is:

```
DELVOL TVOL01 MIGRATION(PURGE)
```

Another example of the DELVOL command for a tape backup volume is:

```
DELVOL TVOL02 BACKUP(PURGE)
```

After the command is issued in either way, both the original and the alternate volume no longer contain any valid data.

11.2 IDCAMS Alter without RECALL

Prior to DFSMShsm Version 1 Release 4, you were not able to alter the catalog entry of a migrated data set without doing a recall of that data set to migration level 0.

This meant that you had to do deal with such time-consuming details as : CPU load, I/O traffic, and waiting for functions to end.

Even if you had the appropriate data set residing on migration level 2, but the tape did not reside in an automated tape library, then you might have to wait for an even longer period of time until the operator mounts a specific tape on a drive. But you do not want to access the data set, you only want to alter the catalog entry.

DFSMShsm Version 1 Release 4 addresses this problem because when an IDCAMS ALTER command is issued to change either the storage or management classes, or both, of a migrated data set, the change is made without recalling the data set. Both DFSMShsm′s and catalog′s records are updated to reflect the change when entering the following:

```
ALTER datasetname STRORAGECLASS(storageclassname)
                  MANAGEMENTCLASS(managementclassname)
ALTER datasetname STRORAGECLASS(storageclassname)
ALTER datasetname MANAGEMENTCLASS(managementclassname)
```

11.3 ABARS Improvements

Aggregate Backup and Recovery Support called (ABARS) is a function originally introduced with DFHSM Version 2 Release 5. It allows you to group data together that belongs to an application, and not only, for example, to a volume.

In DFSMShsm Version 1 Release 4, ABARS processing is enhanced in the following ways:

- Output files generated by ABARS can now be stacked on a minimum number of cartridge volumes.

  In the past you needed at least three tapes for ABARS ABACKUP processing for:

  - Control files
  - Data files
  - Instruction/Activity log file

  Now if you have very few files belonging to an aggregate group, you may only need one cartridge. You may choose this option by setting the following new SETSYS command in member ARCCMDxx in SYS1.PARMLIB:

  ```
  SETSYS ABARSTAPES(STACK)
  ```

  The SETSYS ABARSTAPES(STACK) command allows you to place (stack) the ABACKUP output files from a single aggregate group on a minimum number of tape cartridges (as few as one). The stacking function applies only to tape cartridges. Attempts to redirect ABACKUP output files to DASD when the STACK option is in effect causes ABACKUP to fail.

  Stacked tapes can only be used for ARECOVER processing on processors running DFSMS/MVS Version 1 Release 4 or subsequent releases. The default, if you don′t specify the SETSYS ABARSTAPES command, is the STACK option. Therefore, the SETSYS ABARSTAPES(NOSTACK) command must be specified if you want to direct the ABACKUP output files to DASD, or if the output tapes need to be recovered on a processor running a DFSMS/MVS release prior to Version 1 Release 4.

  Note: When issuing an ABACKUP command, you can temporarily override the SETSYS ABARSTAPES value by using the STACK | NOSTACK parameters of the ABACKUP command.

- The number of concurrent active ABARS requests is enhanced from 15 in the previous releases to 64 in Release 4.

  To specify how many ABACKUP commands can run concurrently, use the MAXABARSADDRESSSPACE parameter. MAXABARSADDRESSSPACE specifies the maximum number of concurrent ABARS secondary address spaces that DFSMShsm supports, up to 64. When determining the value to specify for MAXABARSADDRESSSPACE, you should factor in such items as system load, resource contention, and aggregate backup and recovery activity. The overriding factor is the number of tape drives available. Aggregate backup uses at least one tape drive for each address space and

uses two if backing up user tape data sets or data sets residing on ML2 tape (the default for MAXABARSADDRESSSPACE is 1).

The value is set within the command

```
SETSYS MAXABARSADDRESSSPACE(n)
```

where *n* can be between 1 and 64.

Note: The restriction that you can not backup or recover the same aggregate name simultaneously using the same aggregate group name remains.

- Invocation of exit ARCBEEXT to DFSMSdss dump processing.

  ABACKUP processing is enhanced to extend the error conditions where this exit gets control. The exit gains control when DFSMSdss errors occur while dumping data sets on level 0 DASD in the INCLUDE list.

  The ARCBEEXT installation-wide exit is invoked using the SETSYS EXITON(BE) command, and is entered if either an E level DFSMSdss message is issued (which indicates that a data set has failed dump processing) or a W level DFSMSdss message is issued and DFSMSdss does not indicate that the warning is correct.

  Note: The EXITON(exitname) parameter lets you specify which installation exits (ARCBEEXT, ARCEDEXT, ARCM2EXT, ARCTVEXT) are active during aggregate backup processing.

  You can use the ABARS backup error exit to skip any data set associated with an I/O error, SDSP allocation error, data set noncataloged error, error during DFSMSdss dump processing, or an error enqueuing the ARCDSN resource, so that the data set is not backed up. This allows aggregate backup to complete without ending early should any of these errors occur.

- Allow the specification of a GDG base name in the ALLOCATE statement

  ABARS is enhanced so installations are able to specify a GDG base name in the ALLOCATE statement at a recovery site and have ABARS predefine the GDG base name during ARECOVER processing if it does not currently exist. You can do this even though you have not specified an associated GDS name in the selection data set.

  During ARECOVER processing, all GDG base names are defined prior to restoring any GDSs in the aggregate.

- Automatical deletion of the ABARS activity log during during roll-off processing.

  Currently, ABARS activity log data sets have to be deleted manually, but by using the SETSYS ABARSDELETEACTIVITY parameter, you can specify the option of having DFSMShsm automatically delete the ABARS activity log associated with the aggregate backup or the aggregate recovery version during aggregate roll off or EXPIREBV ABARSVERSIONS processing.

  With the following command:

  ```
  SETSYS ABARSDELETEACTIVITY(Y)
  ```

  you now have the option to delete the ABARS activity log during ABARS roll-off processing or EXPIREBV ABARSVERSIONS processing.

- Aggregate CPU Time in WWFSR.

  This enhancement allows ABARS to maintain the CPU time for processing ABACKUP and ARECOVER requests in the WWFSR control block.

Installations can then use the reported CPU time to assist in calculating the ″charge back″ costs for users′ ABARS requests.

A WWFSR is created at the completion of each individual ABACKUP and ARECOVER request and can be written as an SMF record if requested.

The ABR record maintains an ABACKUP CPU time and an ARECOVER CPU time.

If an ARECOVER request fails and is reissued with a valid RESTART data set, the CPU time in the WWFSR only reflects the processing time of the remaining data sets. The ABR record, however, accumulates the CPU times of each restart until the recovery of the aggregate is successful.

- Externalization of keywords TGTGDS and OPTIMIZE that ABARS passes to DFSMSdss.

  This can be done through the commands:

  ```
  SETSYS ARECOVERTGTGDS(option)
  ```

  and

  ```
  SETSYS ABARSOPTIMIZE(n)
  ```

  TGTGDS is used to provide greater flexibility in managing SMS managed generation data sets that are restored to level 0 DASD.

  OPTIMIZE allows installations to adjust performance when backing up level 0 data sets that are specified in the INCLUDE list.

  Valid values for *options* can be:

  - DEFERRED
  - ACTIVE
  - ROLLEDOFF
  - SOURCE

  Valid values for *n* can be 1, 2, 3 or 4. Each number means the number of cylinders read at a time.

11.4 CDS Record Level Sharing

DFSMShsm needs a number of control data sets called CDS in order to operate against the user or system data sets which are to be processed for backup/restore or migrate/recall. The types of data sets discussed in this section are the following:

- Migration control data set (MCDS)

  The MCDS manages all migrated data sets (migration level 1 and 2) and contains information about those data sets and on which volumes they reside (either as active or as migrated data sets). MCDS is required for DFSMShsm processing.

- Backup control data set (BCDS)

  The BCDS manages all backup copies and also dump copies. It also contains information about the backup and dump environment. MCDS is also a required data set for the DFSMShsm environment.

- Offline control data set (OCDS)

The OCDS manages all tape volumes that contain either backup data sets or migration data sets. It contains information about each data set residing on those tapes. Use of MCDS is not required, but highly recommended.

It is also recommended that you create backup copies of these three types of data sets.

Note: Do not allow DFSMShsm to manage these control data sets, because in case of failure, you would be unable to access them. All of these data sets work not only in a single system environment, but also in a Parallel Sysplex environment.

With DFSMS/MVS Version 1 Release 4, you can use these VSAM key-sequenced data sets (KSDS) in the record level sharing mode.

You can find more detailed information regarding the implementation of these control data sets in *DFSMS/MVS Version 1, Release 4 DFSMShsm Implementation and Customization Guide*.

VSAM record level sharing is a data set access mode which was first introduced with DFSMS/MVS Version 1 release 3. It takes advantage of the locking and capabilities of the the coupling facility in a parallel sysplex environment. With DFSMS/MVS Version 1 Release 4 in OS/390 Release 4, DFSMS/hsm also uses VSAM record level sharing to cache the migration, backup, and offline control data sets (MCDS, BCDS and OCDS) in the coupling facility.

You must use this technique when three or more systems within a Parallel Sysplex work with DFSMShsm, especially when DFSMShsm processes a large number of small data sets, because of contention situations which can negatively impact throughput and performance.

The implementation of VSAM record level sharing for DFSMShsm only uses the CACHE and LOCK structures of the coupling facility in a Parallel Sysplex.

The new parameter CDSSHR specifies whether record level sharing of control data sets will be done or not. This parameter may have one of the follwing values:

YES Specifies that DFSMShsm does multiple-processor serialization of the type requested by the CDSQ parameter (volume serialization with GRS instead of volume reserves) or the CDSR parameter (control data set serialization with volume reserves).

NO Specifies that DFSMShsm does no multiple-processor serialization, and that no other system should be able to concurrently processing the specific control data sets.

RLS Specifies that DFSMShsm performs multiple processor serialization using record level sharing. In that case, the control data sets are accessed in record level sharing mode.

11.5 Dump Analysis Elimination (DAE)

Dump analysis elimination (DAE) is a MVS function which enables elimination of duplicate dumps in or across MVS systems

If you use this function to prevent the system from taking more than one dump in a specific failure situation you have to use the command

```
SETSYS SYS1DUMP
```

in member ARCCMDxx of SYS1.PARMLIB.

Member ADYSETxx of SYS1.PARMLIB must be coded with keyword SUPPRESSALL. You can see the content of such a member in Figure 107.

```
DAE=START,RECORDS(400),
          SVCDUMP(MATCH,SUPPRESSALL,UPDATE),
          SYSMDUMP(MATCH,SUPPRESSALL,UPDATE),
          SHARE(DSN),DSN(SYS1.DAE),GLOBAL(DSN)
```

Figure 107. SYS1.PARMLIB(ADYSETxx) for Dump Analysis Elimination

If you work in a Parallel Sysplex environment, XCF and GRS must be enabled in order for the DAE data set to be shared across the sysplex, and also for dumps to be suppressed across MVS systems.

DAE does not suppress the following dumps:

- SYSABEND
- SYSUDUMP
- SYSMDUMP
- SNAP
- SLIP
- DUMP operator commands
- Dumps produced by DFSMShsm as a result of the TRAP command.

11.6 Serviceability Improvements

DFSMShsm Version 1 Release 4 also offers improvements related to serviceability, as follows:

- New functional statistic record types
- Aliases for DFSMShsm keywords

  This allows:

  - BACKUP to be an alias for BACKUPCONTROLDATASET on the LIST and HLIST command
  - MIGRAT to be an alias for MIGRATIONCONTROLDATASET on the LIST and HLIST command
  - UUT and NOUUT to be used for USERUNITTABLE and NOUSERUNITTABLE on the SETSYS command
- Changes in AUDIT DATASETCONTROLS to improve the restart function by saving all 44 characters of the data set name.

- DFSMSrmm no longer ships code for exit ARCTVEXT. This exit is linked to the DFSMSrmm programming interface EDGDFHSM, which notifies DFSMSrmm of tapes being released from DFSMShsm. EDGDFHSM checks to see if DFSMSrmm is in use.

  Now EDGTVEXT is called from DFSMShsm. If DFSMSrmm is not in use, it no longer calls EDGDFHSM. DFSMShsm always invokes EDGTVEXT before calling ARCEXT.

  ARCTVEXT is invoked based on the command SETSYS EXITON/EXITOFF. This improvement frees ARCTVEXT from DFSMSrmm allowing customers to use ARCTVEXT for tape management systems other than DFSMSrmm.

Chapter 12. DFSMSrmm Enhancements

The following enhancements are now available with DFSMSrmm Version 1 Release 4:

- Journal usage threshold
- Non-intrusive backup
- Recognition of external data managers
- Inventory management trial runs
- DFSMSrmm PDA trace

12.1 Journal Usage Threshold

The DFSMSrmm member EDGRMMxx of SYS1.PARMLIB includes many options for setting up DFSMSrmm, such as:

- Defining system options, such as the date format for reports and messages, default retention periods, and whether to notify volume owners when their volumes are ready to be released
- Preventing a range of tapes from being used on specific systems
- Defining pools, such as the range of shelves to use for a pool, and whether a pool has RACF tape profile processing
- Tailoring mount messages with either the volume's shelf location or the pool identifier
- Defining security classes for data sets and volumes
- Defining the DFSMSrmm running mode to determine when DFSMSrmm records volume usage and performs tape validation
- Defining how DFSMSrmm bypass label processing is performed
- Defining storage locations to DFSMSrmm

The options described in Figure 108 on page 222 are available in member EDGRMMxx. Their description and possible values can be found in *DFSMS/MVS V1R4 DFSMSrmm Implementation and Customization*.

```
OPTION  OPMODE(P)                  /* protect mode         */   -
        TPRACF(A)                  /* automatic            */   -
        NOTIFY(N)                  /* no notification      */   -
        SYSID(DG4)                 /* system name          */   -
        MAXHOLD(100)               /* number of records    */   -
        DSNAME(RMM.CONTROL.DSET)   /* control data set     */   -
        JRNLNAME(RMM.JOURNAL.DSET) /* journal              */   -
>>new>> JOURNALFULL(75)            /* journal percentage   */   -
                                   /* threshold            */   -
        CATRETPD(12)               /* catalog retention    */   -
        SMFAUD(248) SMFSEC(249)    /* SMF records          */   -
        DATEFORM(E)                /* European dates       */   -
        RETPD(5)                   /* default retention    */   -
        MAXRETPD(9999)             /* maximum retention    */   -
        SCRATCHPROC(RMMSCR)        /* scratch procedure    */   -
>>new>> BACKUPPROC(RMMBKUP)        /* backup procedure     */   -
        CDSID(MVS2)                /* control data set ID  */   -
        BLP(RMM)                   /* bypass label process */   -
        VRSJOBNAME(2)              /* retention by job name*/   -
        MSG(M)                     /* message mixed case   */   -
        MASTEROVERWRITE(LAST)      /* overwrite default    */   -
        ACCOUNTING(J)              /* Account information  */   -
        VRSCHANGE(VERIFY)          /* VRSchange information*/   -
        VRSMIN(1,FAIL)             /* VRS minimum          */   -
```

Figure 108. Options of SYS1.PARMLIB Member EDGRMMxx

The OPTION command has two new parameters, as shown in Figure 108:

BACKUPPROC(procname) Specifies the name of the procedure that you want to be started automatically when the journal percentage full threshold is reached. Specify a valid alphanumeric procedure name from 1 to 8 characters. There is no default value. If no name is specified, then no automatic start command is issued.

JOURNALFULL(nn) Specifies JOURNALFULL to define a percentage full threshold for the journal data set. When DFSMSrmm detects that the journal has reached this threshold, DFSMSrmm issues message EDG2107E. DFSMSrmm also issues message EDG2107E at DFSMSrmm startup if the journal has already reached the threshold specified at that time. If you specify a backup procedure name on the BACKUPPROC operand, the procedure is started automatically. Thereafter at intervals of 5%, or 1% after 90% full is reached, message EDG2108E is issued. Values can be in the range of 0-99. The default is 75(%). If you specify a value of 0, DFSMSrmm issues no warnings on that system. You can specify different threshold values for sharing systems.

If you have already DFSMSrmm Version 1 Release 3 installed, then these options are available via PTF. You might have a look at *APAR OW21337*.

12.2 Non-Intrusive Backup

In order to minimize the impact of RMM control data set backup and allow updates to these control data sets during control data set backup and journal backup, you must use the DFSMSdss concurrent copy function. Therefore a new parameter called BACKUP(DSS) is added to the RMM modules EDGHSKP and EDGBKUP. Using this parameter, DFSMSrmm invokes module ADRDSSU in SYS1.LINKLIB and requests a concurrent backup of the control data set.

If a concurrent copy session cannot be done, the backup continues, but updates to the control data set are not allowed.

The description of the BACKUP(DSS) parameter is as follows:

BACKUP(DSS) Specify this parameter whenever you want the non-intrusive version of BACKUP to be executed. To use this, you must have the hardware and software required to establish a concurrent copy session.

In this case, hardware means that you must have *at least* a storage control system IBM 3990 model 3 with extended platform. The IBM 3990 model 6 or IBM 9390 also provide concurrent copy. Without such a system, the backup will be intrusive and updates to the control data set are prevented until backup completes.

You should ensure that all systems are on the same level of code supporting this feature before it is used because backups taken using DFSMSdss cannot be restored by using a prior level of restore. However, they can be restored by using a combination of DFSMSdss and the EDGBKUP utility for forward recovery.

When you specify BACKUP(DSS) and you use concurrent copy, the DFSMSrmm journal is only reset if you also backup the journal data set. Either the journal or the journal backup is required to restore and forward-recover to a consistent point in time, because use of concurrent copy allows the control data set (and therefore the journal) to be updated during backup processing.

The DFSMSrmm backup utility makes calls to DFSMSdss. To invoke the DFSMSdss component, code BACKUP(DSS) in the EXEC parm statement of the appropriate job, as in Figure 109 on page 224, and also add a DSSOPT DD statement.

```
//EDGBKUP  EXEC PGM=EDGBKUP,PARM='BACKUP(DSS)'
//SYSPRINT DD SYSOUT=*
//BACKUP   DD DISP=(,CATLG),UNIT=TAPE,DSN=BACKUP.CDS(+1),
//         LABEL=(,SL)
//JRNLBKUP DD DISP=(,CATLG),UNIT=TAPE,DSN=BACKUP.JRNL(+1),
//         LABEL=(2,SL),VOL=REF=*.BACKUP
//DSSOPT   DD *
    CONCURRENT OPTIMIZE(1) VALIDATE
/*
```

Figure 109. DFSMSrmm Backup Job

The DFSMSdss DUMP command therefore used by DFSMSrmm is:

```
DUMP DS(INCLUDE(cds_name)) OUTDD(BACKUP) SHARE -
   CONCURRENT OPTIMIZE(1) VALIDATE
```

You can use EDGBKUP with DFSMSdss DUMP or the access method services REPRO command. When you use DFSMSdss concurrent copy, DFSMSrmm allows update to the DFSMSrmm control data set. Use the EDGBKUP program to:

- Back up the control data set to a non-VSAM data set or a VSAM cluster.
- Back up the journal.
- Restore the control data set by returning a backup of the control data set and, optionally, forward-recover it.
- Forward-recover a control data set by applying journal updates to a copy of the control data set.

The backup then can go directly to tape if DFSMSdss is used.

Serialization of the control data set means that further tape processing must wait. This waiting period includes:

- The timeframe of the whole backup when IDCAMS is used.
- The timeframe concurrent copy initialization takes if it is used together with DFSMSdss.
- The timeframe of the whole backup when DFSMSdss is used without the function of concurrent copy.

If you already have DFSMSrmm Version 1 Release 3 installed, then these options are also available via PTF. You might have a look at *APAR OW21337*.

12.3 Support for DFSMShsm Alternate Tape Processing

In the past, if an alternate backup or migration tape had to be used for recalling or recovering a user data set, DFSMSrmm checked that the full 44-character DFSMShsm data set name matched the name in the control data set of DFSMSrmm. However, since the data set name of the alternate copy always had the name *COPY* as its second qualifier, the check failed.

When APAR OW21257 is installed, DFSMSrmm only validates the last 17 characters of the data set name used by DFSMShsm when it is opening a tape

for input processing. This DFSMSrmm validation allows DFSMShsm to use the data set name of the original tape, rather than the data set name of the copy tape. If APAR OW21257 is not installed, you must change the data set name known to DFSMSrmm in order for DFSMSrmm to process the tape. Issue the RMM CHANGEVOLUME subcommands to ensure that DFSMSrmm does not prevent DFSMShsm from using the tape. You must issue these RMM subcommands because DFSMShsm does not communicate the name difference resulting from the true replacement to DFSMSrmm.

12.4 Recognition of External Data Managers

In past releases, DFSMSrmm kept track of statistics at the whole volume level. With this release, DFSMSrmm can be instructed to manage a tape volume so that only the *first* data set is tracked for details, and not the following ones. In SYS1.SAMPLIB, there is a member with an installation exit called EDGUX100 that you can update to perform this task.

The purpose for creating an update to this exit is that you may have other applications (external data managers to other tape management systems) that manage pools of tapes and the data sets on those tapes. Previously, these external data managers stored information about the data sets, and if DFSMSrmm was also in use, it duplicated the information in the control data set. This resulted in a need for more disk capacity (duplicated data set information) and longer backup times for the DFSMSrmm-owned control data set and journal.

Figure 110 shows how the new definitions in the exit are made. The first DC statements show that jobs with the name JOBWRITE, data sets beginning with AUVENJ.*, and a program with the name IEBDG define the data set inventory on tape.

```
EDGUX100 CSECT
EDMTAB   DS    0F                  start of dsr table
*
         DC    CL8'JOBWRITE'                         jobname
         DC    CL44'AUVENJ.*'                        data set name
         DC    CL8'IEBDG'                            PROGRAM NAME
*
         DC    CL8'*'                                jobname
         DC    CL44'*'                               data set name
         DC    CL8'DUMMYGEN'                         program name
*
         DC    CL8'DB2%%REC'                         jobname
         DC    CL44'HSMATH%.*'                       data set name
         DC    CL8'PROGRAM3'                         program name
*
         DC    CL8'DUMP*'                            jobname
         DC    CL44'RMMU001.T%%L*'                   data set name
         DC    CL8'*'                                program name
*
         DC    CL8'EDM END'        end of dsr table marker
```

Figure 110. Recognition of External Data Managers in RMM Exit EDGUX100

12.5 Inventory Management Trial Runs

It is now possible to make a trial run of inventory management vital record processing before DFSMSrmm makes these changes to the control data set or the journal. Such a trial run lets you see how the vital record specifications which were made would be processed in a real production run. The exception is that DFSMSrmm does not make any changes to the control data set or the journal.

Now an ACTIVITY file (which must be allocated), contains information on all changes to the control data set that would be made during vital record processing.

You can find a job with the allocation-related information for the ACTIVITY file in member EDGJHKPA of SYS1.SAMPLIB. This is shown in Figure 111.

```
//EDGJHKPA JOB HARJANS,MSGLEVEL=(1,1),NOTIFY=HARJANS
//*********************************************************************
//*                                                                   *
//*  THIS SAMPLE JCL CAN BE USED TO ALLOCATE THE FILES REQUIRED       *
//*  TO EXECUTE RMM INVENTORY MANAGEMENT FUNCTIONS                    *
//*                                                                   *
//*********************************************************************
//ALLOC    EXEC PGM=IEFBR14
//BACKUP   DD DSN=BACKUP.FILE.NAME,DISP=(NEW,CATLG),
//            UNIT=SYSDA,AVGREC=U,SPACE=(4096,(1000,500)),
//            LRECL=9000,RECFM=VB
//JRNLBKUP DD DSN=JOURNAL.BACKUP.FILE.NAME,DISP=(NEW,CATLG),
//            UNIT=SYSDA,AVGREC=U,SPACE=(4096,(1000,500)),
//            LRECL=9000,RECFM=VB
//MESSAGE  DD DSN=MESSAGE.FILE.NAME,DISP=(NEW,CATLG),
//            UNIT=SYSDA,AVGREC=U,SPACE=(4096,(10,10))
//REPORT   DD DSN=REPORT.FILE.NAME,DISP=(NEW,CATLG),
//            UNIT=SYSDA,AVGREC=U,SPACE=(4096,(50,30))
//ACTIVITY DD DSN=ACTIVITY.FILE.NAME,DISP=(NEW,CATLG),
//            UNIT=SYSDA,AVGREC=U,SPACE=(4096,(50,30))
//REPTEXT  DD DSN=REPORT.EXTRACT.FILE.NAME,DISP=(NEW,CATLG),
//            UNIT=SYSDA,AVGREC=U,SPACE=(4096,(1000,500))
```

Figure 111. RMM Job to Allocate an ACTIVITY File

You should then submit an inventory management job using Job Control Language, as shown in Figure 112, to request a trial run of vital record processing.

```
//HSKP     EXEC PGM=EDGHSKP,
//     PARM='VRSEL,VERIFY'
//SYSPRINT DD *
//MESSAGE  DD DISP=SHR,DSN=HSKP.MESSAGES
//REPORT   DD DISP=SHR,DSN=HSKP.VRS.REPORT
//ACTIVITY DD DISP=SHR,DSN=HSKP.ACTIVITY
```

Figure 112. Trial Run Job for Inventory Management

12.6 DFSMSrmm Problem Determination Aid (PDA) Trace

DFSMSrmm provides a PDA trace that gathers diagnostic information at specific internal module points. It records this kind of data in a storage circular file in the RMM address space and, optionally, also on a disk file.

The problem determination aid (PDA) facility gathers DFSMSrmm processing information to enable analysis to pinpoint module flow and resource usage related to DFSMSrmm problems. The PDA facility is required for IBM service because it traces module and resource flow. The PDA facility consists of in-storage trace, optional DASD log data sets, EDGRMMxx parmlib member options, and operator commands to control tracing.

DFSMSrmm accumulates problem determination information at specific module points in the form of trace data, and it records this data in main storage. At predetermined intervals, the trace data is scheduled for output to DASD. The DFSMSrmm trace recording function receives the trace data scheduled for output and writes this data to a file on DASD. The PDA facility consists of two separate log data sets. DFSMSrmm recognizes these log data sets by their DD names, EDGPDOX and EDGPDOY. Recording takes place in the data set defined by EDGPDOX. When that data set is filled, the two data set names are swapped, and recording continues on the newly renamed data set.

When this data set is filled, the names are again swapped, and the output switches to the other data set, thus overlaying the previously recorded data. The larger the data sets, the longer the period of time that is represented by the accumulated data.

The preferred implementation of the PDA facility is to establish a protocol that automatically copies the EDGPDOY data set to tape as a generation-data-group data set each time DFSMSrmm issues message EDG9117I. This practice provides a sequential history of trace data over time so that the data is available when needed for resolving problems.

Appendix A. WLM CDS Allocation using Panels

The WLM Administrative Application is invoked by executing “IWMARIN0.” This can be done by selecting the relevent option from the ISPF panel or by entering

ex ’sys1.sblscli0(iwmarin0)’

under TSO. The following entry panel is displayed.

```
  File  Help
------------------------------------------------------------------------------

Command ===> ____________________________________________________________


                        W   W   L      M   M
                        W   W   L      MM MM
                        W W W   L      M M M
                        WW WW   L      M   M
                        W   W   LLLLL  M   M

                Licensed Materials - Property of IBM

                5645-001 (C) Copyright IBM Corp. 1997.
                All rights reserved.

                          ENTER to continue
```

The “Choose Service Definition” panel provides you with choices about where the service definition data is coming from.

```
  File  Help
------------------------------------------------------------------------------

Command ===> ____________________________________________________________


              ____________________________________________
              ]         Choose Service Definition          ]
              ]                                            ]
              ] Select one of the following options.       ]
              ] 1   1.  Read saved definition                ]
              ]     2.  Extract definition from WLM        ]
              ]         couple data set                    ]
              ]     3.  Create new definition              ]
              ]                                            ]
              ]                                            ]
              ]                                            ]
              ____________________________________________
                          ENTER to continue
```

Since we will be using the service definition that was saved earlier key in **1** and depress the ENTER key.

```
  File  Help
------------------------------------------------------------------------

Command ===> ____________________________________________________________


             ___________________________________________
             ]         Choose Service Definition         ]
             ]                                           ]
             ] Select one of the following options.      ]
             ] 1   1.  Read saved definition             ]
      ___________________________________________________________________
      ]                       Read saved definition                       ]
      ]                                                                   ]
      ]                                                                   ]
      ] Data set name  . . . ’RONN.O39OR3.WLM.SERVICE.DEFNTION’             ]
      ]                                                                   ]
      ] Type a ?  for a list of previously defined service definition data ]
      ] sets.                                                             ]
      ]                                                                   ]
      ]                                                                   ]
      ___________________________________________________________________
```

A titled popup window appears, “Read saved definition.” Key in the name of the data set. Highlighted in this instance.

```
  File  Utilities  Notes  Options  Help
------------------------------------------------------------------------
Functionality LEVEL001         Definition Menu         WLM Appl LEVEL004
Command ===> ____________________________________________________________

Definition data set  . . : ’RONN.O39OR3.WLM.SERVICE.DEFNTION’

Definition name  . . . . . CICSpol   (Required)
Description  . . . . . . . OS/390 R3 Service Definition

Select one of the
following options. . . . . ___  1.  Policies
                               2.  Workloads
                               3.  Resource Groups
                               4.  Service Classes
                               5.  Classification Groups
                               6.  Classification Rules
                               7.  Report Classes
                               8.  Service Coefficients/Options
                               9.  Application Environments
                              10.  Scheduling Environments
```

The data is read, which can take a while if large and you are presented with the primary WLM administrative application panel; “Definition Menu.”

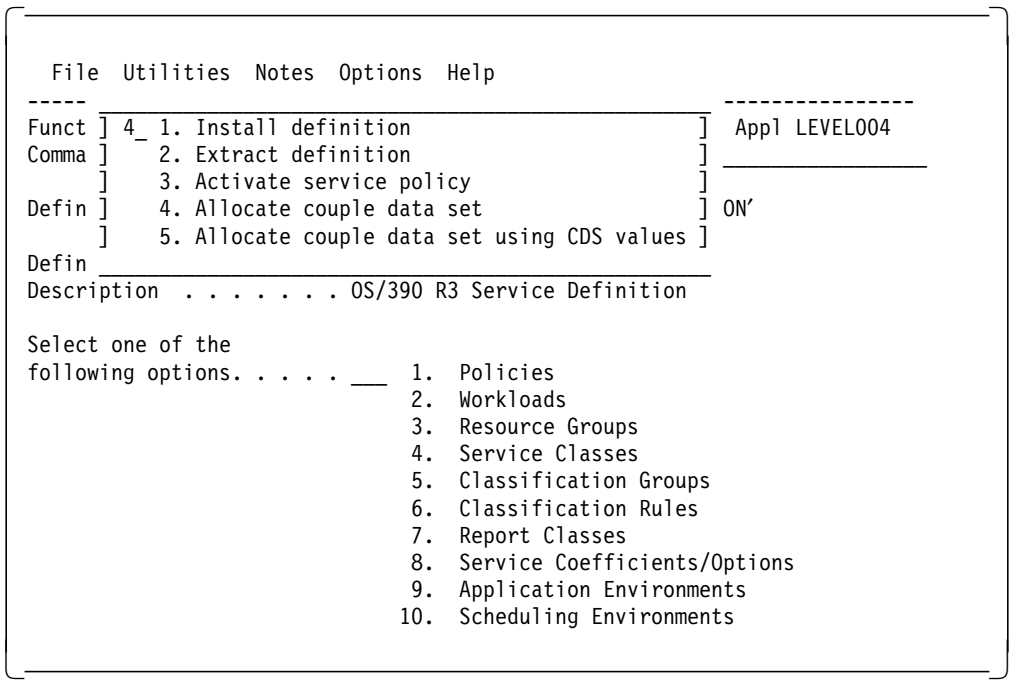
File Utilities Notes Options Help
Funct
Comma
Defin
Defin
1. Install definition
2. Extract definition
3. Activate service policy
4. Allocate couple data set
5. Allocate couple data set using CDS values
Appl LEVEL004
ON'
Description . . . . . OS/390 R3 Service Definition
Select one of the following options. . . . .
1. Policies
2. Workloads
3. Resource Groups
4. Service Classes
5. Classification Groups
6. Classification Rules
7. Report Classes
8. Service Coefficients/Options
9. Application Environments
10. Scheduling Environments

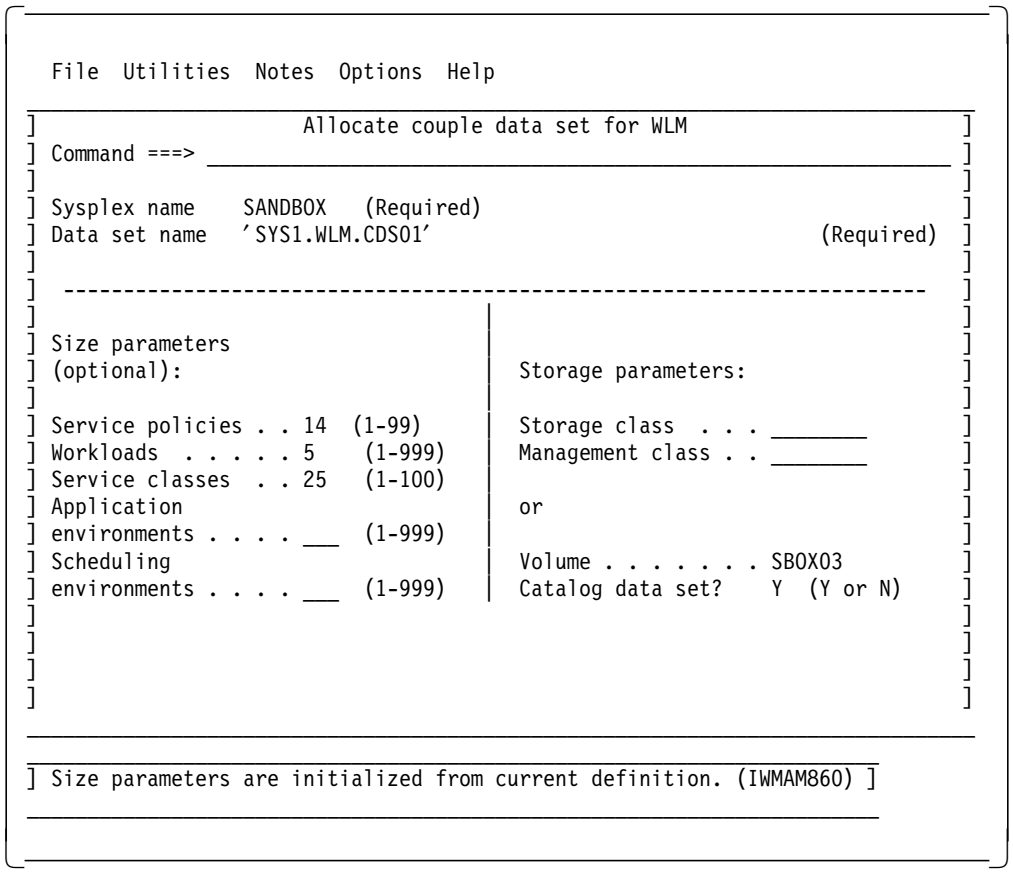
File Utilities Notes Options Help
Allocate couple data set for WLM
Command ===>
Sysplex name SANDBOX (Required)
Data set name 'SYS1.WLM.CDS01' (Required)
Size parameters (optional):
Service policies . . 14 (1-99)
Workloads 5 (1-999)
Service classes . . 25 (1-100)
Application environments . . . . ___ (1-999)
Scheduling environments . . . . ___ (1-999)
Storage parameters:
Storage class . . . ________
Management class . . ________
or
Volume SBOX03
Catalog data set? Y (Y or N)
Size parameters are initialized from current definition. (IWMAM860)

```
  File  Utilities  Notes  Options  Help
 ──────────────────────────────────────────────────────────────────────────────
 ]                      Allocate couple data set for WLM                       ]
 ] Command ===> ____________________________________________________________   ]
 ]                                                                             ]
 ] Sysplex name    SANDBOX   (Required)                                        ]
 ] Data set name   'SYS1.WLMR4.CDS01'                              (Required)  ]
 ]                                                                             ]
 ]  ------------------------------------------------------------------------   ]
 ]                                    |                                        ]
 ] Size parameters                    |                                        ]
 ] (optional):                        |  Storage parameters:                   ]
 ]                                    |                                        ]
 ] Service policies . . 10  (1-99)    |  Storage class  . . . ________          ]
 ] Workloads  . . . . . 35   (1-999)  |  Management class . . ________          ]
 ] Service classes  . . 100  (1-100)  |                                        ]
 ] Application                        |  or                                    ]
 ] environments . . . . ___  (1-999)  |                                        ]
 ] Scheduling                         |  Volume . . . . . . . SBOX03           ]
 ] environments . . . . ___  (1-999)  |  Catalog data set?    n  (Y or N)      ]
 ]                                                                             ]
 ]                                                                             ]
 ]                                                                             ]
 ]                                                                             ]
 ──────────────────────────────────────────────────────────────────────────────
```

```
  File  Edit  Confirm  Menu  Utilities  Compilers  Test  Help
 ──────────────────────────────────────────────────────────────────────────────
 EDIT       RONN.SPFTEMP1.CNTL                            Columns 00001 00072
 Command ===>                                               Scroll ===> PAGE
 ****** ***************************** Top of Data ******************************
 ==MSG> -Warning- The UNDO command is not available until you change
 ==MSG>           your edit profile using the command RECOVERY ON.
 000001 //*********************************************************************
 000002 //*                                                                   *
 000003 //* JCL to format the WLM couple data set by invoking IXCL1DSU, the *
 000004 //* XCF format utility.                                               *
 000005 //*                                                                   *
 000006 //* Couple data set allocation rules:                                 *
 000007 //*                                                                   *
 000008 //*   1. SYSPLEX name is required and is 1-8 characters.              *
 000009 //*   2. SYSPRINT DD is a required DD statement for format utility    *
 000010 //*      messages.                                                    *
 000011 //*   3. SYSIN DD is a required DD statement for format utility       *
 000012 //*      control statements.                                          *
 000013 //*   4. When allocating space for service definition extensions      *
 000014 //*      use the SVDEFEXT, SVDCREXT, SVAEAEXT and SVSEAEXT  @WLMPSTP*
 000015 //*      SYSIN keywords.                                              *
 000016 //*                                                                   *
 000017 //*      - SVDEFEXT is used to set aside space for extensions to      *
 000018 //*        the IWMSVDEF.  Values for SVDEFEXT range from 0 to 1024    *
 000019 //*        and are in terms of kilobytes.                             *
 000020 //*                                                                   *
 000021 //*      - SVDCREXT is used to set aside space for extensions to      *
 000022 //*        the IWMSVDCR.  Values for SVDEFEXT range from 0 to 1024    *
 000023 //*        and are in terms of kilobytes.                             *
```

```
EDIT       RONN.SPFTEMP1.CNTL                              Columns 00001 00072
Command ===>                                                  Scroll ===> PAGE
000024 //*                                                                    *
000025 //*      - SVAEAEXT is used to set aside space for extensions to    *
000026 //*        the IWMSVAEA.  Values for SVAEAEXT range from 0 to 1024  *
000027 //*        and are in terms of kilobytes.                  @WLMPSTP*
000028 //*                                                                    *
000029 //*      - SVSEAEXT is used to set aside space for extensions to    *
000030 //*        the IWMSVSEA.  Values for SVSEAEXT range from 0 to 1024  *
000031 //*        and are in terms of kilobytes.                  @WLMPRES*
000032 //*                                                                    *
000033 //*      Reference MVS/ESA Programming: Workload Management         *
000034 //*      Services for more details on SVDEFEXT, SVDCREXT,           *
000035 //*      SVAEAEXT and SVSEAEXT.                          @PQCPRES*
000036 //*                                                      @PQC1518*
000037 //*   5. MAXSYSTEM specifies the number of systems to be    @PQC1518*
000038 //*      uesd by the new couple data set.  If this          @PQC1518*
000039 //*      parameter is omitted, the format utility assigns   @PQC1518*
000040 //*      a default value of eight.  For WLM we make the     @PQC1518*
000041 //*      MAXSYSTEM value 32 so that it equal to the maximum @PQC1518*
000042 //*      number of systems allowed in a sysplex.            @PQC1518*
000043 //*      See OS/390 MVS Setting Up A Sysplex for more       @PQC1518*
000044 //*      details.                                           @PQC1518*
000045 //*                                                         @PQC1518*
000046 //*   6. When allocating the WLM couple data set WLM        @PSY0613*
000047 //*      provides a number of DATA-NAMEs for the ITEM       @PSY0613*
000048 //*      NAME(DATA-NAME) statement.  For example, POLICY is @PSY0613*
000049 //*      a DATA-NAME provided to specify the number of      @PSY0613*
```

```
EDIT       RONN.SPFTEMP1.CNTL                              Columns 00001 00072
Command ===>                                                  Scroll ===> PAGE
000050 //*      administrative policies that can be defined within @PSY0613*
000051 //*      the WLM couple data set.                           @PSY0613*
000052 //*                                                         @PSY0613*
000053 //*      When allocating the WLM couple data set the        @PSY0613*
000054 //*      DATA-NAMEs are dependent upon the MVS release that @PSY0613*
000055 //*      is running.                                        @PSY0613*
000056 //*                                                         @PSY0613*
000057 //*      For example, the APPLENV keyword was introduced in @PSY0613*
000058 //*      OS/390 Release 3.  If APPLENV is used in the job to@PSY0613*
000059 //*      allocate the WLM couple data set it must run on a  @PSY0613*
000060 //*      system running OS/390 Release 3 or later system.   @PSY0613*
000061 //*                                                         @PSY0613*
000062 //*      The following lists the DATA-NAMEs and at what     @PSY0613*
000063 //*      release they were introduced.  The DATA-NAMEs are  @PSY0613*
000064 //*      available to that release and any later release.   @PSY0613*
000065 //*      For example,  POLICY is available on all releases  @PSY0613*
000066 //*      and SCHENV is available on OS/390 R3 and R4.       @PSY0613*
000067 //*                                                         @PSY0613*
000068 //*      Release             DATA-NAME                      @PSY0613*
000069 //*      -------             ---------                      @PSY0613*
000070 //*      MVS/SP 5.1          POLICY                         @PSY0613*
000071 //*                          WORKLOAD                       @PSY0613*
000072 //*                          SRVCLASS                       @PSY0613*
000073 //*      MVS/SP 5.2          SVDEFEXT                       @PSY0613*
000074 //*                          SVDCREXT                       @PSY0613*
000075 //*      OS/390 Release 3    APPLENV                        @PSY0613*
```

```
EDIT       RONN.SPFTEMP1.CNTL                                Columns 00001 00072
Command ===>                                                    Scroll ===> PAGE
000076 //*                           SVAEAEXT                        @PSYO613*
000077 //*      OS/390 Release 4      SCHENV                          @PSYO613*
000078 //*                           SVSEAEXT                        @PSYO613*
000079 //*                                                           @PSYO613*
000080 //*      To run the job to allocate the WLM couple data set @PSYO613*
000081 //*      on the correct system use the JES2 /*JOBPARM or    @PSYO613*
000082 //*      JES3 //*MAIN statement to route control to the     @PSYO613*
000083 //*                                                         @PSYO613*
000084 //*      Reference MVS/ESA JCL User's Guide for more details@PSYO613*
000085 //*      on how to select a system.                         @PSYO613*
000086 //*                                                         @PSYO613*
000087 //*****************************************************************
000088 //STEP1    EXEC PGM=IXCL1DSU
000089 //STEPLIB  DD   DSN=SYS1.MIGLIB,DISP=SHR
000090 //SYSPRINT DD   SYSOUT=A
000091 //SYSIN    DD   *
000092    DEFINEDS SYSPLEX(SANDBOX)
000093             DSN(SYS1.WLMR4.CDS01)
000094             MAXSYSTEM(32)
000095             VOLSER(SBOX03) NOCATALOG
000096          DATA TYPE(WLM)
000097               ITEM NAME(POLICY) NUMBER(10)
000098               ITEM NAME(WORKLOAD) NUMBER(35)
000099               ITEM NAME(SRVCLASS) NUMBER(100)
000100 /*
```

Appendix B. Batch Execution Using a Scheduling Environment

Scenario

- There are two systems in the sysplex:
  - Development (SC63)
  - Production (SC64)
- SC64 processes a CICS ATM online workload and the ATM batch suite must be processed on the same system.
- A scheduling environment (BATCH_ATM) has been defined in the WLM service definition to set scheduling affinities for these batch jobs.

The first action is to activate the policy and switch the systems to goal mode.

```
V WLM,POLICY=RONNP1

IWM001I WORKLOAD MANAGEMENT POLICY RONNP1 NOW IN EFFECT


F WLM,MODE=GOAL

IWM007I SYSTEM SC63 NOW IN WORKLOAD MANAGEMENT GOAL MODE

ROUTE SC64 F WLM,MODE=GOAL

IWM007I SYSTEM SC64 NOW IN WORKLOAD MANAGEMENT GOAL MODE


D WLM,SYSTEMS

IWM025I  18.40.34  WLM DISPLAY 846
  ACTIVE WORKLOAD MANAGEMENT SERVICE POLICY NAME: RONNP1
  ACTIVATED: 1997/06/03  AT: 18:40:12  BY: RONN      FROM: SC63
  DESCRIPTION: Test OS/390 R4 WLM Enhancements
  RELATED SERVICE DEFINITION NAME: WLMR4_SD
  INSTALLED: 1997/06/03  AT: 18:37:08  BY: RONN      FROM: SC63
  WLM VERSION LEVEL: LEVEL004
  *SYSNAME*  *MODE*  *POLICY*  *WORKLOAD MANAGEMENT STATUS*
  SC63       GOAL    RONNP1    ACTIVE
  SC64       GOAL    RONNP1    ACTIVE
```

Display the current status of BATCH_ATM and the resource states on both systems.

```
D WLM,SCHENV=BATCH_ATM

IWM036I  10.27.45  WLM DISPLAY 011
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   NOT AVAILABLE ON ANY SYSTEM


D WLM,SCHENV=BATCH_ATM,SYSTEM=SC63

IWM037I  10.37.52  WLM DISPLAY 013
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   SYSTEM:                SC63
   STATUS:                NOT AVAILABLE
                       REQUIRED  CURRENT
    RESOURCE NAME      STATE     STATE
   *ATM_CICS           ON        RESET
   *DBRP               ON        RESET
   *PRD_SYS            ON        RESET

D WLM,SCHENV=BATCH_ATM,SYSTEM=SC64

IWM037I  10.38.15  WLM DISPLAY 015
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   SYSTEM:                SC64
   STATUS:                NOT AVAILABLE
                       REQUIRED  CURRENT
    RESOURCE NAME      STATE     STATE
   *ATM_CICS           ON        RESET
   *DBRP               ON        RESET
   *PRD_SYS            ON        RESET
```

Set the resource states ON for SC64 and redisplay the scheduling environment status.

```
SC64 F WLM,RESOURCE=PRD_SYS,ON
IWM039I RESOURCE PRD_SYS IS NOW IN THE ON STATE

SC64 F WLM,RESOURCE=ATM_CICS,ON
IWM039I RESOURCE ATM_CICS IS NOW IN THE ON STATE

SC64 F WLM,RESOURCE=DBRP,ON
IWM039I RESOURCE DBRP IS NOW IN THE ON STATE


D WLM,SCHENV=BATCH_ATM,SYSTEM=SC64
IWM037I  10.40.31  WLM DISPLAY 018
  SCHEDULING ENVIRONMENT: BATCH_ATM
   DESCRIPTION:           Production Batch For ATMs
   SYSTEM:                SC64
   STATUS:                AVAILABLE
                       REQUIRED  CURRENT
    RESOURCE NAME      STATE     STATE
    ATM_CICS           ON        ON
    DBRP               ON        ON
    PRD_SYS            ON        ON
```

Jobs sumitted from SC63 with the BATCH_ATM scheduling environment coded on the jobcard will be scheduled to process on SC64. As shown on the following display.

```
$DJ1077,LONG

$HASP890 JOB(RONNATM1)
$HASP890 JOB(RONNATM1)  STATUS=(AWAITING EXECUTION),
$HASP890                CLASS=W,PRIORITY=5,SYSAFF=(ANY),
$HASP890                HOLD=(JOB),CMDAUTH=(LOCAL),
$HASP890                OFFS=(),SECLABEL=,USERID=RONN,
$HASP890                SPOOL=(VOLUMES=(SBOX02),TGS=1,
$HASP890                PERCENT=0.0500),ARM_ELEMENT=NO,
$HASP890                CARDS=3,REBUILD=NO,
$HASP890                SRVCLASS=STC_HI,
$HASP890                SCHENV=BATCH_ATM,
$HASP890                SCHENV_AFF=(SC64),CC=()
```

At this point all the requirements for this job are met and it is initiated by WLM as class W has been designated a WLM managed class. This is the view from the syslog.

```
IWM034I PROCEDURE INIT STARTED FOR SUBSYSTEM JES2 915
APPLICATION ENVIRONMENT SYSBATCH
PARAMETERS SUB=MSTR
IEF196I         1 //INIT     JOB MSGLEVEL=1
IEF196I         2 //STARTING EXEC INIT
IEF196I
IEF196I           XX*************************************************************
IEF196I           XX*
IEF196I           XX*01* PROCEDURE NAME : INIT
IEF196I           XX*
IEF196I           XX*01* COPYRIGHT:
IEF196I           XX*          LICENSED MATERIALS - PROPERTY OF IBM
IEF196I           XX*          THIS MODULE IS "RESTRICTED MATERIALS OF IBM"
IEF196I           XX*          5740-XC6 (C) COPYRIGHT IBM CORP. 1987
IEF196I           XX*          SEE COPYRIGHT INSTRUCTIONS
IEF196I           XX*
IEF196I           XX*01* STATUS:  HBB3310
IEF196I           XX*
IEF196I           XX*01* FUNCTION:
IEF196I           XX*
IEF196I           XX*       The INIT procedure is used to start the MVS Initiator
IEF196I           XX*       by calling the module IEFIIC.
IEF196I           XX*
IEF196I           XX*       DPRTY=12  -  Sets the dispatching priority for the
IEF196I           XX*                    initiator.
IEF196I           XX*
IEF196I           XX*01* COMPONENT:
IEF196I           XX*
IEF196I           XX*       SC1B6 (Initiator)
IEF196I           XX*
IEF196I           XX*01* DISTRIBUTION LIBRARY:
IEF196I           XX*
IEF196I           XX*       SYS1.PROCLIB
IEF196I           XX*
IEF196I           XX*01* CHANGE ACTIVITY:
IEF196I           XX*
IEF196I           XX*************************************************************
IEF196I  STMT NO. MESSAGE
IEF196I         2 IEFC001I PROCEDURE INIT WAS EXPANDED USING SYSTEM
IEF196I LIBRARY SYS1.PROCLIB
IEF196I         3 XXIEFPROC   EXEC   PGM=IEFIIC,DPRTY=12

ICH70001I RONN     LAST ACCESS AT 09:56:35 ON WEDNESDAY, JUNE 4, 1997
$HASP373 RONNATM1 STARTED - WLM INIT  - SRVCLASS STC_HI   - SYS SC64
PROGRAM : IEFBR14  IN STEP : STP1     - COND CODE : 0000
TOTAL CPU (SECS)  BLKS READ/WRITE      TOTAL SWAPS
            0.02
$HASP395 RONNATM1 ENDED
$HASP099 ALL AVAILABLE FUNCTIONS COMPLETE
SE '10.39.40 JOB01049 $HASP165 RONNATM1 ENDED AT WTSCPLX1  MAXCC=0',
LOGON,USER=(RONN)
```

This is what the JES2 joblog and jobstream look like.

```
  J E S 2  J O B  L O G  --  S Y S T E M  S C 6 3  --  N O D E  W T S C P L X 1

JOB01049 ---- WEDNESDAY, 04 JUN 1997 ----
JOB01049  IRR010I  USERID RONN     IS ASSIGNED TO THIS JOB.
JOB01049  ICH70001I RONN     LAST ACCESS AT 09:56:35 ON WEDNESDAY, JUNE 4, 1997
RONNATM1  $HASP373 RONNATM1 STARTED - WLM INIT  - SRVCLASS STC_HI   - SYS SC64
JOB01049  PROGRAM : IEFBR14  IN STEP : STP1     - COND CODE : 0000
JOB01049  TOTAL CPU (SECS)  BLKS READ/WRITE      TOTAL SWAPS
JOB01049             0.02
RONNATM1  $HASP395 RONNATM1 ENDED
------ JES2 JOB STATISTICS ------
  04 JUN 1997 JOB EXECUTION DATE
            3 CARDS READ
           29 SYSOUT PRINT RECORDS
            0 SYSOUT PUNCH RECORDS
            2 SYSOUT SPOOL KBYTES
         0.00 MINUTES EXECUTION TIME
        1 //RONNATM1 JOB (POK,999),NORTHRUP,CLASS=W,MSGCLASS=T,NOTIFY=&SYSUID,
          //             SCHENV='BATCH_ATM',TYPRUN=HOLD
          IEFC653I SUBSTITUTION JCL - (POK,999),NORTHRUP,CLASS=W,MSGCLASS=T,
                                      NOTIFY=RONN,SCHENV='BATCH_ATM',TYPRUN=HOLD
        2 //STP1    EXEC PGM=IEFBR14
ICH70001I RONN     LAST ACCESS AT 09:56:35 ON WEDNESDAY, JUNE 4, 1997
IEF142I RONNATM1 STP1 - STEP WAS EXECUTED - COND CODE 0000
********************************************************************
********************************************************************
IEF373I STEP/STP1    /START 1997155.1039
IEF374I STEP/STP1    /STOP  1997155.1039 CPU    0MIN 00.02SEC SRB    0MIN 00.00SEC
IEF375I  JOB/RONNATM1/START 1997155.1039
IEF376I  JOB/RONNATM1/STOP  1997155.1039 CPU    0MIN 00.02SEC SRB    0MIN 00.00SEC
```

Appendix C. Special Notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (″vendor″) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM Advanced Function Presentation

Advanced Function Printing	AFP
AnyNet	APPN
CBIPO	CBPDO
CICS	CICS/ESA
CICS/MVS	DB2
DFSMS	DFSMS/MVS
DFSMSdfp	FFST
IBM	IMS
IMS/ESA	Intelligent Printer Data Stream
IP PrintWay	IPDS
MVS (logo)	MVS/ESA
NetSpool	OS/390
Parallel Sysplex	Print Services Facility
PrintWay	PSF
RACF	S/390
ServicePac	System/390
SystemPac	VTAM

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Inc.

Microsoft, Windows, Windows NT, and the Windows 95 logo
are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is
used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

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.

Appendix D. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 253.

Short Title	Title	Order Number
OS/390 Release 3 Implementation	*OS/390 Release 3 Implementation*	SG24-2067
OS/390 Release 2 Implementation MVS, SMP/E, SDSF, and RMF	*OS/390 Release 2 Implementation*	SG24-4834
Version 5 Implementation Guide	*MVS/ESA Version 5 Implementation Guide*	SG24-4584
MVS 5.1 Presentation Guide	*MVS/ESA SP 5.1.0 Technical Presentation Guide*	GG24-4137
Parallel Sysplex Perf.	*S/390 Parallel Sysplex Performance*	SG24-4356
HCD and Dynamic I/O Reconfiguration Primer	*MVS/ESA HCD and Dynamic I/O Reconfiguration Primer*	SG24-4037
Sysplex Migration Guide	*MVS/ESA Version 5 Sysplex Migration Guide*	SG24-4581
S/390 G3 Enterprise Server: Complex Systems Availability and Recovery Presentation Guide	*S/390 G3 Enterprise Server: CSAR Presentation Guide*	SG24-4911

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

D.3 Other Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook. A publication whose order number begins with the prefix **LY** is available to IBM-licensed customers only.

Note: A few publications in the following list are no longer available in hardcopy format and cannot be separately ordered. These publications are only available on CD-ROMs.

These publications are also relevant as further information sources:

- OS/390 OpenEdition

Short Title	Title	Order Number
OS/390 OpenEdition MVS User's Guide	*OS/390 OpenEdition MVS User's Guide*	SC28-1891
OS/390 OpenEdition MVS Command Reference	*OS/390 OpenEdition MVS Command Reference*	SC28-1892
OS/390 OpenEdition MVS Programming Tools	*OS/390 OpenEdition MVS Programming Tools*	SC28-1904
OS/390 OpenEdition MVS Messages and Codes	*OS/390 OpenEdition MVS Messages and Codes*	SC28-1908
OS/390 OpenEdition MVS Programming: Assembler Callable Services Reference	*OS/390 OpenEdition MVS Programming: Assembler Callable Services Reference*	SC28-1899
OS/390 OpenEdition MVS Planning	*OS/390 OpenEdition MVS Planning*	SC28-1890
OS/390 OpenEdition MVS File System Interface Reference	*OS/390 OpenEdition MVS File System Interface Reference*	SC28-1909
OS/390 OpenEdition MVS Using REXX and OpenEdition MVS	*OS/390 OpenEdition MVS Using REXX and OpenEdition MVS*	SC28-1905
OS/390 OpenEdition MVS Communications Server Guide	*OS/390 OpenEdition MVS Communications Server Guide*	SC28-1906

- I/O Configuration Management

Short Title	Title	Order Number
OS/390 HCD Planning	*OS/390 Hardware Configuration Definition Planning*	GC28-1750
OS/390 HCD User's Guide	*OS/390 HCD User's Guide*	SC28-1848
HCD Messages	*OS/390 HCD Messages*	GC28-1849
HCD Scenarios	*OS/390 HCD Scenarios*	SC28-1850

- RMF

Short Title	Title	Order Number
RMF Messages and Codes	*OS/390 RMF Messages and Codes*	GC28-1948
RMF Performance Management Guide	*OS/390 RMF Performance Management Guide*	SC28-1951
RMF User's Guide	*OS/390 RMF User's Guide*	SC28-1949
RMF Report Analysis	*OS/390 RMF Report Analysis*	SC28-1950
RMF Programmer's Guide	*OS/390 RMF Programmer's Guide*	SC28-1952

- Multi-System Configuration Management

Short Title	Title	Order Number
OS/390 Parallel Sysplex Systems Management	*OS/390 Parallel Sysplex Systems Management*	GC28-1861
OS/390 Parallel Sysplex Hardware and Software Migration	*OS/390 Parallel Sysplex Hardware and Software Migration*	GC28-1862
OS/390 Parallel Sysplex Application Migration	*OS/390 Parallel Sysplex Application Migration*	GC28-1863
OS/390 V2R4.0 MVS Setting Up a Sysplex	*OS/390 MVS Setting Up a Sysplex*	GC28-1779
OS/390 V2R4.0 MVS Sysplex Services Guide	*OS/390 MVS Programming: Sysplex Services Guide*	GC28-1771
OS/390 V2R4.0 MVS Sysplex Services Reference	*OS/390 MVS Programming: Sysplex Services Reference*	GC28-1772

- OS/390 Operating System

Short Title	Title	Order Number
OS/390 V2R4.0 MVS Auth Assembler Services Reference ALE-DYN	*OS/390 MVS Programming: Authorized Assembler Services Reference, Volume 1, ALE-DYN*	GC28-1764
OS/390 V2R4.0 MVS Auth Assembler Services Reference ENF-ITT	*OS/390 MVS Programming: Authorized Assembler Services Reference, Volume 2, ENF-ITT*	GC28-1765
OS/390 V2R4.0 MVS Auth Assembler Services Reference LLA-SDU	*OS/390 MVS Programming: Authorized Assembler Services Reference, Volume 3, LLA-SDU*	GC28-1766
OS/390 V2R4.0 MVS Auth Assembler Services Reference SET-WTO	*OS/390 MVS Programming: Authorized Assembler Services Reference, Volume 4, SET-WTO*	GC28-1767
OS/390 V2R4.0 MVS Extended Addressability Guide	*OS/390 MVS Programming: Extended Addressability Guide*	GC28-1769
OS/390 V2R4.0 MVS Assembler Services Guide	*OS/390 MVS Programming: Assembler Services Guide*	GC28-1762
OS/390 V2R4.0 MVS Assembler Services Reference	*OS/390 MVS Programming: Assembler Services Reference*	GC28-1910
OS/390 V2R4.0 MVS Auth Assembler Services Guide	*OS/390 MVS Programming: Authorized Assembler Services Guide*	GC28-1763
Introducing OS/390	*OS/390 Introduction and Release Guide*	GC28-1725
OS/390 V1R2.0 MVS JCL User's Guide	*OS/390 MVS JCL User's Guide*	GC28-1758
OS/390 V2R4.0 MVS JCL Reference	*OS/390 MVS JCL Reference*	GC28-1757
OS/390 V2R4.0 MVS Callable Services for HLL	*OS/390 MVS Programming: Callable Services for High-Level Languages*	GC28-1768
OS/390 V2R4.0 MVS Writing TPs for APPC/MVS	*OS/390 MVS: Writing Transaction Programs for APPC/MVS*	GC28-1775
OS/390 V2R4.0 MVS Planning: APPC/MVS Management	*OS/390 MVS Planning: APPC/MVS Management*	GC28-1807

Short Title	Title	Order Number
OS/390 V2R4.0 MVS IPCS Commands	*OS/390 MVS Interactive Problem Control System (IPCS) Commands*	GC28-1754
OS/390 V2R4.0 MVS IPCS User's Guide	*OS/390 MVS Interactive Problem Control System (IPCS) User's Guide*	GC28-1756
OS/390 V2R4.0 MVS IPCS Customization	*OS/390 MVS Interactive Problem Control System (IPCS) Customization*	GC28-1755
OS/390 V2R4.0 MVS Initialization and Tuning Guide	*OS/390 MVS Initialization and Tuning Guide*	SC28-1751
OS/390 V2R4.0 MVS Initialization and Tuning Reference	*OS/390 MVS Initialization and Tuning Reference*	SC28-1752
OS/390 V2R4.0 MVS Installation Exits	*OS/390 MVS Installation Exits*	SC28-1753
OS/390 V2R4.0 MVS Using the Functional Subsystem Interface	*OS/390 MVS Using the Functional Subsystem Interface*	SC28-1911
OS/390 V2R4.0 MVS Conversion Notebook	*OS/390 MVS Conversion Notebook*	GC28-1747
OS/390 V2R4.0 MVS Using the Subsystem Interface	*OS/390 MVS Using the Subsystem Interface*	SC28-1789
OS/390 V2R4.0 MVS System Commands Summary	*OS/390 MVS System Commands Summary*	GX22-0040
OS/390 V2R4.0 Planning for Installation	*OS/390 Planning for Installation Release 4*	GC28-1726
OS/390 V2R4.0 MVS System Commands	*OS/390 MVS System Commands*	GC28-1781
OS/390 V2R4.0 MVS System Management Facilities (SMF)	*OS/390 MVS System Management Facilities (SMF)*	GC28-1783
OS/390 V2R4.0 MVS Planning: Operations	*OS/390 MVS Planning: Operations*	GC28-1760
OS/390 V2R4.0 MVS Planning: Global Resource Serialization	*OS/390 MVS Planning: Global Resource Serialization*	GC28-1759
OS/390 V2R4.0 MVS System Data Set Definition	*OS/390 MVS System Data Set Definition*	GC28-1782
OS/390 V2R4.0 MVS Device Validation Support	*OS/390 MVS Device Validation Support*	GC28-1748
OS/390 V2R4.0 MVS System Messages, Vol 1 (ABA-ASA)	*OS/390 MVS System Messages, Volume 1 (ABA-ASA)*	GC28-1784
OS/390 V2R4.0 MVS System Messages, Vol 2 (ASB-ERB)	*OS/390 MVS System Messages, Volume 2 (ASB-ERB)*	GC28-1785
OS/390 V2R4.0 MVS System Messages, Vol 3 (GDE-IEB)	*OS/390 MVS System Messages, Volume 3 (GDE-IEB)*	GC28-1786
OS/390 V2R4.0 MVS System Messages, Vol 4 (IEC-IFD)	*OS/390 MVS System Messages, Volume 4 (IEC-IFD)*	GC28-1787
OS/390 V2R4.0 MVS System Messages, Vol 5 (IGD-IZP)	*OS/390 MVS System Messages, Volume 5 (IGD-IZP)*	GC28-1788
OS/390 V2R4.0 MVS Dump Output Messages	*OS/390 MVS Dump Output Messages*	GC28-1749
OS/390 V2R4.0 MVS System Codes	*OS/390 MVS System Codes*	GC28-1780
OS/390 V2R4.0 MVS Routing and Descriptor Codes	*OS/390 MVS Routing and Descriptor Codes*	GC28-1778
OS/390 V2R4.0 MVS Recovery and Reconfiguration Guide	*OS/390 MVS Recovery and Reconfiguration Guide*	GC28-1777

Short Title	Title	Order Number
OS/390 V2R4.0 MVS JES Common Coupling Services	*OS/390 MVS Programming: JES Common Coupling Services*	GC28-1770
OS/390 V2R4.0 MVS: Writing Servers for APPC/MVS	*OS/390 MVS: Writing Servers for APPC/MVS*	GC28-1774
OS/390 V2R4.0: MVS Writing Transaction Schedulers for APPC/MVS	*OS/390 MVS: Writing Transaction Schedulers for APPC/MVS*	GC28-1776
OS/390 MVS APPC/MVS Handbook for OS/2	*OS/390 MVS APPC/MVS Handbook for OS/2*	GC28-1746
OS/390 V2R4.0 MVS Product Registration	*OS/390 MVS Programming: Product Registration*	GC28-1729
OS/390 V2R4.0 MVS Product Management	*OS/390 MVS Product Management*	GC28-1730
OS/390 V2R4.0 MVS Diagnosis: Procedures	*OS/390 MVS Diagnosis: Procedures*	LY28-1082
OS/390 V2R4.0 MVS Diagnosis: Tools and Service Aids	*OS/390 MVS Diagnosis: Tools and Service Aids*	LY28-1085
OS/390 V2R4.0 MVS Diagnosis: Reference	*OS/390 MVS Diagnosis: Reference*	LY28-1084
OS/390 V2R4.0 MVS Planning: Workload Management	*OS/390 MVS Planning: Workload Management*	GC28-1761
OS/390 V2R4.0 MVS Workload Management Services	*OS/390 MVS Programming: Workload Management Services*	GC28-1773
OS/390 Information Roadmap	*OS/390 Information Roadmap*	GC28-1727
OS/390 V2R4.0 MVS Diagnosis: Tools and Service Aids	*OS/390 MVS Diagnosis: Tools and Service Aids*	LY28-1845
OS/390 Introduction and Release Guide	*OS/390 Introduction and Release Guide*	GC28-1725
SystemView for MVS Up and Running!	*SystemView for MVS Up and Running!*	GC28-1241
OS/390 Planning for Installation	*OS/390 V2R4.0 Planning for Installation*	GC28-1726
OS/390 Information Roadmap	*OS/390 V2R4.0 Information Roadmap*	GC28-1727

- SMP/E and Installation Manuals

Short Title	Title	Order Number
ServerPac Guide and Worksheet	*ServerPac Guide and Worksheet*	SC28-1244
MVS Packaging Rules	*Standard Packaging Rules for MVS-Based Products*	SC23-3695
OS/390 SMP/E Diagnosis Guide	*OS/390 System Modification Program Extended Diagnosis Guide*	SC28-1737
OS/390 SMP/E Messages and Codes	*OS/390 System Modification Program Extended Messages and Codes*	SC28-1738
OS/390 SMP/E User's Guide	*OS/390 System Modification Program Extended User's Guide*	SC28-1740
OS/390 SMP/E Command Reference	*OS/390 System Modification Program Extended Command Reference*	SC28-1805
OS/390 SMP/E Reference	*OS/390 System Modification Program Extended Reference*	SC28-1806

- JES2 Subsystem

Short Title	Title	Order Number
OS/390 JES2 Messages	*OS/390 JES2 Messages*	GC28-1796
OS/390 JES2 Commands	*OS/390 JES2 Commands*	GC28-1790
OS/390 JES2 Initialization and Tuning Guide	*OS/390 JES2 Initialization and Tuning Guide*	SC28-1791
OS/390 JES2 Initialization and Tuning Reference	*OS/390 JES2 Initialization and Tuning Reference*	SC28-1792
OS/390 JES2 Diagnosis	*OS/390 JES2 Diagnosis*	LY28-1086
OS/390 JES2 Introduction	*OS/390 JES2 Introduction*	GC28-1794
OS/390 JES2 Installation Exits	*OS/390 JES2 Installation Exits*	SC28-1793
OS/390 JES2 Macros	*OS/390 JES2 Macros*	SC28-1795
OS/390 JES2 Migration Notebook	*OS/390 JES2 Migration Notebook*	GC28-1797

- JES3 Subsystem

Short Title	Title	Order Number
OS/390 JES3 Conversion Notebook	*OS/390 JES3 Conversion Notebook*	GC28-1799
OS/390 JES3 Diagnosis	*OS/390 JES3 Diagnosis*	LY28-1090
OS/390 JES3 Initialization and Tuning Guide	*OS/390 JES3 Initialization and Tuning Guide*	SC28-1802
OS/390 JES3 Introduction	*OS/390 JES3 Introduction*	GC28-1808
OS/390 JES3 Initialization and Tuning Reference	*OS/390 JES3 Initialization and Tuning Reference*	SC28-1803
OS/390 JES3 Customization	*OS/390 JES3 Customization*	LY28-1089
OS/390 JES3 Messages	*OS/390 JES3 Messages*	GC28-1804
OS/390 JES3 Commands	*OS/390 JES3 Commands*	GC28-1798
OS/390 JES3 Diagnosis Reference	*OS/390 JES3 Diagnosis Reference*	LY28-1092

- ICKDSF

Short Title	Title	Order Number
ICKDSF R16 Refresh User's Guide	*ICKDSF R16 Refresh User's Guide*	GC35-0033

- Storage Environment

Short Title	Title	Order Number
DFSMS/MVS V1R4 General Information	*DFSMS/MVS Version 1 Release 4 General Information*	GC26-4900
DFSMS/MVS V1R4 Licensed Program Specifications	*DFSMS/MVS Version 1 Release 4 Licensed Program Specifications*	GC26-4903
DFSMS/MVS V1R4 Planning for Installation	*DFSMS/MVS Version 1 Release 4 Planning for Installation*	SC26-4919

Short Title	Title	Order Number
DFSMS/MVS V1R4 Installation Exits	*DFSMS/MVS Version 1 Release 4 Installation Exits*	SC26-4908
DFSMS/MVS V1R4 Managing Catalogs	*DFSMS/MVS Version 1 Release 4 Managing Catalogs*	SC26-4914
DFSMS/MVS V1R1 Checkpoint/Restart	*DFSMS/MVS Version 1 Release 1 Checkpoint/Restart*	SC26-4907
DFSMS/MVS V1R4 Using Data Sets	*DFSMS/MVS Version 1 Release 4 Using Data Sets*	SC26-4922
DFSMS/MVS V1R4 Macro Instructions for Data Sets	*DFSMS/MVS Version 1 Release 4 Macro Instructions for Data Sets*	SC26-4913
DFSMS/MVS V1R4 DFSMSdfp Advanced Services	*DFSMS/MVS Version 1 Release 4 DFSMSdfp Advanced Services*	SC26-4921
DFSMS/MVS V1R4 Access Method Services for ICF	*DFSMS/MVS Version 1 Release 4 Access Method Services for the Integrated Catalog Facility*	SC26-4906
DFSMS/MVS V1R4 Access Method Services for VSAM	*DFSMS/MVS Version 1 Release 4 Access Method Services for VSAM Catalogs*	SC26-4905
DFSMS/MVS V1R4 DFSMSdfp Storage Administration Reference	*DFSMS/MVS Version 1 Release 4 DFSMSdfp Storage Administration Reference*	SC26-4920
DFSMS/MVS V1R4 Using ISMF	*DFSMS/MVS Version 1 Release 4 Using the Interactive Storage Management Facility*	SC26-4911
DFSMS/MVS V1R4 Utilities	*DFSMS/MVS Version 1 Release 4 Utilities*	SC26-4926
DFSMS/MVS V1R4 DFM/MVS Guide and Reference	*DFSMS/MVS Version 1 Release 4 Distributed FileManager/MVS Guide and Reference*	SC26-4915
DFSMS/MVS V1R4 Program Management	*DFSMS/MVS Version 1 Release 4 Program Management*	SC26-4916
DFSMS/MVS V1R4 DFSMSdfp Diagnosis Guide	*DFSMS/MVS Version 1 Release 4 DFSMSdfp Diagnosis Guide*	LY27-9605
DFSMS/MVS V1R4 DFSMSdfp Diagnosis Reference	*DFSMS/MVS Version 1 Release 4 DFSMSdfp Diagnosis Reference*	LY27-9606
DFSMS/MVS V1R4 Using Magnetic Tapes	*DFSMS/MVS Version 1 Release 2 Using Magnetic Tapes*	SC26-4923
DFSMS/MVS V1R4 Using the Volume Mount Analyzer	*DFSMS/MVS Version 1 Release 2 Using the Volume Mount Analyzer*	SC26-4925
DFSMS/MVS V1R4 Implementing System-Managed Storage	*DFSMS/MVS Version 1 Release 4 Implementing System-Managed Storage*	SC26-3123
MVS/ESA SML: Leading a Storage Administration Group	*MVS/ESA Storage Management Library: Leading a Storage Administration Group*	SC26-3126
MVS/ESA SML: Managing Storage Groups	*MVS/ESA Storage Management Library: Managing Storage Groups*	SC26-3125
MVS/ESA SML: Managing Data	*MVS/ESA Storage Management Library: Managing Data*	SC26-3124
DFSMS/MVS V1R4 OAM Application Programmer's Reference	*DFSMS/MVS Version 1 Release 2 Object Access Method Application Programmer's Reference*	SC26-4917
DFSMS/MVS V1R4 OAM Planning, Installation, and Administration Guide for Object Support	*DFSMS/MVS Version 1 Release 4 Object Access Method Planning, Installation, and Storage Administration Guide for Object Support*	SC26-4918

Short Title	Title	Order Number
DFSMS/MVS V1R4 OAM Planning, Installation, and Storage Administration Guide for Tape Libraries	*DFSMS/MVS Version 1 Release 4 Object Access Method Planning, Installation, and Storage Administration for Tape Libraries*	SC26-3051
DFSMS/MVS V1R4 DFSMSrmm Guide and Reference	*DFSMS/MVS Version 1 Release 4 DFSMSrmm Guide and Reference*	SC26-4931
DFSMS/MVS V1R4 DFSMSrmm Implementing and Customizing	*DFSMS/MVS Version 1 Release 4 DFSMSrmm Implementing and Customizing Guide*	SC26-4932
DFSMS/MVS V1R4 DFSMSrmm Diagnosis Guide	*DFSMS/MVS Version 1 Release 4 DFSMSrmm Diagnosis Guide*	LY27-9615
DFSMS/MVS V1R4 DFSMShsm Managing Your Own Data	*DFSMS/MVS Version 1 Release 4 DFSMShsm Managing Your Own Data*	SH21-1077
DFSMS/MVS V1R4 DFSMShsm Storage Administration Guide	*DFSMS/MVS Version 1 Release 4 DFSMShsm Storage Administration Guide*	SH21-1076
DFSMS/MVS V1R4 Managing Data Availability	*DFSMS/MVS Version 1 Release 4 Managing Data Availability*	SC26-4928
DFSMS/MVS V1R4 DFSMShsm Implementing and Customizing	*DFSMS/MVS Version 1 Release 4 DFSMShsm Implementing and Customizing*	SH21-1078
DFSMS/MVS V1R4 DFSMShsm Diagnosis Guide	*DFSMS/MVS Version 1 Release 4 DFSMShsm Diagnosis Guide*	LY27-9607
DFSMS/MVS V1R4 DFSMShsm Diagnosis Reference	*DFSMS/MVS Version 1 Release 4 DFSMShsm Diagnosis Reference*	LY27-9608
DFSMS/MVS V1R4 DFSMShsm Storage Administration Reference	*DFSMS/MVS Version 1 Release 4 DFSMShsm Storage Administration Reference*	SH21-1075
DFSMS/MVS V1R4 DFSMSdss Storage Administration Guide	*DFSMS/MVS Version 1 Release 4 DFSMSdss Storage Administration Guide*	SC26-4930
DFSMS/MVS V1R4 DFSMSdss Storage Administration Reference	*DFSMS/MVS Version 1 Release 4 DFSMSdss Storage Administration Reference*	SC26-4929
DFSMS/MVS V1R4 DFSMSdss Diagnosis Guide	*DFSMS/MVS Version 1 Release 4 DFSMSdss Diagnosis Guide*	LY27-9609
Stand-Alone Services Overview	*Data Facility Data Set Services Version 2 Release 5 and DFSMS/MVS Version 1 Stand-Alone Services Overview*	SC26-0185
Remote Copy Administrator's Guide	*Remote Copy Administrator's Guide*	SC35-0169
DFSMS/MVS Network File System User's Guide	*DFSMS/MVS Network File System User's Guide*	SC26-7028
DFSMS/MVS Network File System Customization and Operation	*DFSMS/MVS Network File System Customization and Operation*	SC26-7029
DFSMS/MVS V1R4 Network File System Performance Tuning Guide	*DFSMS/MVS Version 1 Release 2 Network File System Performance Tuning Guide*	SC26-7019

- Security Server

Short Title	Title	Order Number
OS/390 Security Server (RACF) General User's Guide	*OS/390 Security Server (RACF) General User's Guide*	SC28-1917
OS/390 Security Server (RACF) System Programmer's Guide	*OS/390 Security Server (RACF) System Programmer's Guide*	SC28-1913

Short Title	Title	Order Number
OS/390 Security Server (RACF) Macros and Interfaces	*OS/390 Security Server (RACF) Macros and Interfaces*	SC28-1914
OS/390 Security Server (RACF) Command Language Reference	*OS/390 Security Server (RACF) Command Language Reference*	SC28-1919
OS/390 Security Server (RACF) Introduction	*OS/390 Security Server (RACF) Introduction*	GC28-1912
OS/390 Security Server (RACF) Messages and Codes	*OS/390 Security Server (RACF) Messages and Codes*	SC28-1918
OS/390 Security Server (RACF) Security Administrator's Guide	*OS/390 Security Server (RACF) Security Administrator's Guide*	SC28-1915
OS/390 Security Server (RACF) Auditor's Guide	*OS/390 Security Server (RACF) Auditor's Guide*	SC28-1916
OS/390 Security Server External Security Interface (RACROUTE) Macro Reference	*OS/390 Security Server External Security Interface (RACROUTE) Macro Reference*	GC28-1922
OS/390 Security Server (RACF) Callable Services	*OS/390 Security Server (RACF) Callable Services*	GC28-1921
OS/390 Security Server (RACF) Planning: Installation and Migration	*OS/390 Security Server (RACF) Planning: Installation and Migration*	GC28-1920
OS/390 Security Server (RACF) Support for MVS OpenEdition DCE, SOMobjects for MVS and SystemView	*OS/390 Security Server (RACF) Support for MVS OpenEdition DCE, SOMobjects for MVS and SystemView*	GC28-1924
OS/390 Security Server (OpenEdition DCE Security Server) Overview	*OS/390 Security Server (OpenEdition DCE Security Server) Overview*	GC28-1938

- TSO/E

Short Title	Title	Order Number
TSO/E Administration	*OS/390 TSO/E Administration*	SC28-1966
TSO/E REXX Reference	*OS/390 TSO/E REXX Reference*	SC28-1975
TSO/E General Information	*OS/390 TSO/E General Information*	GC28-1964
TSO/E Customization	*OS/390 TSO/E Customization*	SC28-1965
TSO/E Programming Guide	*OS/390 TSO/E Programming Guide*	SC28-1970
TSO/E Programming Services	*OS/390 TSO/E Programming Services*	SC28-1971
TSO/E CLISTs	*OS/390 TSO/E CLISTs*	SC28-1973
TSO/E Guide to SRPI	*OS/390 TSO/E Guide to the Server-Requester Programming Interface*	SC28-1976
TSO/E User's Guide	*OS/390 TSO/E User's Guide*	SC28-1968
TSO/E REXX User's Guide	*OS/390 TSO/E REXX User's Guide*	SC28-1974
TSO/E System Programming Command Reference	*OS/390 TSO/E System Programming Command Reference*	SC28-1972
TSO/E Command Reference	*OS/390 TSO/E Command Reference*	SC28-1969
TSO/E Messages	*OS/390 TSO/E Messages*	GC28-1978
TSO/E Primer	*OS/390 TSO/E Primer*	GC28-1967

- PSF/MVS

Short Title	Title	Order Number
AFP Printer Information	*Advanced Function Presentation: Printer Information*	G544-3290
AFP Printer Summary	*Advanced Function Presentation: Printer Summary*	G544-3135
Guide to Advanced Function Presentation	*Guide to Advanced Function Presentation*	G544-3876
IBM IP PrintWay Guide	*IBM IP PrintWay Guide*	S544-5379
IBM NetSpool Guide	*IBM NetSpool Guide*	G544-5301
OS/390 V1R3.0: Printing Softcopy BOOKs	*OS/390 V1R3.0: Printing Softcopy BOOKs*	S544-5354
PSF/MVS Messages and Codes	*Print Services Facility/MVS: Messages and Codes Version 2 Release 2 Modification 0*	S544-3675

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
(+45) 48 14 2207 (long distance charge)	Outside North America

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com
- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number
- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

Index

Special Characters

A

B

C

D

E

F

G

H

I

J

ITSO Redbook Evaluation

OS/390 Release 4 Implementation
SG24-2089-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redeval@vnet.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

_____

_____

_____

_____

_____

Printed in U.S.A.