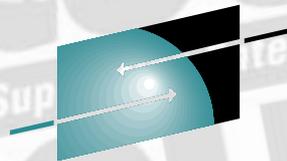


BY R.P. SHANNON

# Resource Affinity Scheduling in OS/390 R4

Resource Affinity Scheduling was introduced in OS/390 R4. It provides the ability to ensure that a job will execute on a system that has all the resources the job requires.



**SHARE**

Technology • Connections • Results

*This article is sponsored by SHARE Inc., a non-profit, voluntary association consisting of users of IBM information technology. SHARE's purpose is to bring together top IT professionals for learning, networking and advocacy purposes.*

*SHARE's next technical conference and technology exchange will be held in Washington D.C., August 16 - 21, 1998. For more information on SHARE Inc., or the upcoming conference, call 1-888-5SHARE5, email share-hq@share.org, or visit their web site at [www.share.org](http://www.share.org).*

**THROUGH** the introduction of Parallel Sysplex and the Workload Manager (WLM), IBM has invested heavily to improve transaction processing. Stogy, unglamorous batch processing was largely ignored until OS/390 R4 when WLM Managed Initiators and Resource Affinity Scheduling were introduced. Resource Affinity Scheduling provides a mechanism to direct batch jobs to specific systems based on the availability of resources. This article will examine Resource Affinity Scheduling and provide some tips on how to implement it in your shop.

## WHY NOT RUN BATCH EVERYWHERE?

The most common reasons for directing batch work to specific systems are that some resources are not available on every system, or that it was decided batch work should simply not run on a system during a particular time or perhaps not at all. The latter is often indicative of a system dedicated to online processing.

One of the tenets of a Parallel Sysplex is to configure systems symmetrically so that all hardware and software resources are available to all systems in the Sysplex. Why then would anyone choose to not configure symmetrically? There are few good reasons. Some devices, particularly older ones that attach with parallel channels, can only be attached to a fixed number of systems. The non-ESCON 3995 Optical Storage and "Round Tape" (3420) are examples. Other special purpose hardware such cryptographic or vector processors may only be installed on selected systems.

Software, too, can be tricky. OS/390, with its 70 or so elements, will be installed on every system in a Sysplex, but shops

may choose to not activate every element on every system. Although Parallel Sysplex Licensing Charges (PSLC) are the norm, there are some products that will only be licensed for a limited number of systems: compilers, 4GL report writers, utility programs, repositories such as MXG, and data base managers that are out of fashion to name a few. Add to the list a hodgepodge of products acquired through acquisitions and mergers, products that may not have the same favorable license agreements that you usually negotiate, and symmetry goes out the window. Small shops may limit licensing for budgetary reasons. Large shops may decide that symmetry alone cannot justify the expense of licensing software on dozens of systems.

## DIRECTING ACCESS TO RESOURCES

Routing work to resources is not a new concept. Over the years many elaborate job class and initiator schemes have been established to ensure batch jobs execute on the correct system. These schemes tend not to be very granular, and often initiators must support multiple job classes, which can introduce additional problems.

System Affinity, a JES2 construct, can be employed to identify the systems on which a job is eligible to execute. System Affinity is requested by coding a SYSAFF parameter on the /\*JOBPARM statement. The default system is ANY, which makes every system in the Multi-access Spool (MAS) eligible. JES3 provides similar capability on the /\*MAIN card. The drawback to SYSAFF or /\*MAIN is that the user's JCL must be changed to either add additional systems or to remove systems from eligibility.

## A NEW APPROACH

Resource Affinity Scheduling offers an alternative to prior schemes. Implementation requires four steps:

- ◆ identify resources
- ◆ associate resources with scheduling environments
- ◆ identify the systems on which the resources are available
- ◆ request a scheduling environment through JCL

The advantages are granularity and flexibility. Up to 999 scheduling environments can be created; however, I recommend restraint. JES2 support for the SCHENV statement was introduced in OS/390 R4. JES3 support is not available in R5 but should be forthcoming. SCHENV is ignored by APPC-initiated jobs, and if coded will cause a started task to fail. If SYSAFF is coded in conjunction with SCHENV, the eligible systems will be limited to those identified on the SYSAFF statement.

Resource Affinity Scheduling requires all systems sharing a JES2 multi-access spool (MAS) to be at JES2 OS/390 2.4 or higher. Additionally, either a \$ACTIVATE must be issued or a cold start performed. One must minimally run in WLM compatibility mode. This means the IPS/ICS can be used instead of Goal Mode, but a WLM couple dataset must be allocated and a WLM policy must be active.

Do not issue \$ACTIVATE until you have read WSC Flash WF9817A - "JES2 \$ACTIVATE Warning for OS/390 Release 4 and 5," as problems have been encountered in some shops.

## RESOURCES

A resource identifies something that is necessary for job execution. It can be hardware such as a vector processor, a product such as ADABAS, a construct such as a particular test database, or something abstract such as shift. A resource must be given a name. Using descriptive resource names such as PL1\_Compiler is recommended, but from a purely technical standpoint a name such as XXX will work just as well.

Resources can be in three states: ON, OFF, and RESET. ON indicates a resource is available; OFF indicates it is unavailable; and RESET indicates the state has not been set.

## SCHEDULING ENVIRONMENT

A scheduling environment is comprised of a set of resources. An environment may be comprised of a single resource, such as a COBOL Compiler or a group of resources such as SAS, a database, and a 3995 Optical Disk. Users request a scheduling environment by coding the SCHENV parameter on the jobcard. Users cannot directly request resources; they can only request a scheduling environment. If a nonexistent or invalid scheduling environment is requested, the job will fail with a JCL error. If the scheduling environment name contains an underscore (\_), the name must be in quotes. For example:

```
SCHENV='3420_TAPE'.
```

## GETTING STARTED

Resource Affinity Scheduling requires a Basic or Parallel Sysplex. Resource Affinity Scheduling is correctly advertised as available in WLM Compatibility Mode. However, what may not be obvious is that the scheduling environment is implemented as part of a WLM Service Definition. This requires a policy, workload, and service class to be defined even if you intend to run in Compatibility Mode. Scheduling environments are defined using the Workload Manager ISPF Dialog. The steps to create a scheduling environment include:

1. define the WLM Couple Datasets
2. create a service definition that includes the scheduling environment
3. install the service definition
4. activate the policy
5. modify WLM to define resource states

WLM Couple Datasets created prior to OS/390 R3 may not be large enough due to the addition of some new items. Use the IXCL1DSU utility to create new WLM Couple Datasets. A new item name, SCHENV, should be coded to ensure space for the scheduling environments.

## USING THE DIALOG

The WLM ISPF Dialog is used to create the scheduling environment. Select Option 3 "Create New Definition" to begin (Figure 1). Select Option 10 to create a scheduling environment (Figure 2). You are given the option to define scheduling environments first or resources first. We elected to define a scheduling environment (Figure 3) as the

dialog will lead into the Resource Definition Panel. A scheduling environment can be up to 16 alphanumeric characters, including national characters \$, #, or &, and can also contain imbedded underscores. An underscore cannot be the first or last character. We defined a scheduling environment called OSVS\_COBOL (Figure 4).

Under "Action" we entered an "a" (add) to define resources. In the Resource Definition we also entered an "a" under Action to define a resource (Figure 5). The resource we defined is OSVS\_COBOL, the same name as the scheduling environment (Figure 6). A scheduling environment can have the same name as one of its resources or a different name. In this case, we think the name of the environment will make it easy for the users to remember and request. After defining Resource OSVS\_COBOL we selected it for the OSVS\_COBOL scheduling environment (Figure 7).

We also defined and selected a resource named PROD\_ONLINE (not shown). Whenever production online work is running, the PROD\_ONLINE Resource state will be ON. On the test system the normal state will be OFF.

We returned to the Scheduling Environment panel and indicated the required state of the OSVS\_COBOL Resource is ON and the required state of PROD\_ONLINE is OFF (Figure 8). Jobs requesting the OSVS\_COBOL scheduling environment will not execute on systems where online work is running. Exiting the panel caused the scheduling environment to be created (Figure 9).

We repeated this exercise several times until we created four scheduling environments (Figure 10) and 10 resources (Figure 11). Figure 12 shows the resources and their required state for each scheduling environment.

Before we could install our definition we had to define a WLM policy, workload, resource group, and service class. Since we have no immediate plans to run in Goal Mode, we defined these items with minimal information. Our policy is called SHARE. The Workload name is DUMMY as is the Resource Group name. We defined one service class with an average response time of one hour. **Do not switch to Goal Mode using this policy!** We installed our definition by tabbing to the Utilities Option on the primary WLM panel and selecting option 1 "Install definition". The last step was to activate the service policy using Option 3 of the Utility pop-up "Activate service

policy”. WLM modify commands were subsequently issued to make the scheduling environments available.

**IDENTIFYING THE STATE OF RESOURCES**

The WLM does not dispatch agents to locate resources. Rather, it must be explicitly told which resources are available on which systems. Clearly the state of hardware resources such as cryptographic processors could be identified by automatically querying the systems; however, this capability currently does not exist. Instead, the state of resources is identified by issuing a Modify Command:

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

When a policy is activated its resources have a state of RESET. Scheduling environments that require any of the resources will have a status of 'NOT AVAILABLE'. A scheduling environment will automatically become 'AVAILABLE' when all of its resources are in the required state (ON or OFF). Sample modify commands are shown in Figure 13.

When the WLM is ready to accept modify commands, it issues message "IWM041I WORKLOAD MANAGEMENT ADDRESS SPACE MODIFY COMMAND AVAILABLE". During IPL, COMMNDxx is processed long before the WLM will accept modify commands, so the modify commands cannot be issued from COMMNDxx. You can use an automation product to issue the commands, triggered by message IWM041I.

The difficulty in automating the modify commands is in predefining the state of resources. Most automation products can read a file and issue commands based on file contents. For each system in a Sysplex there will have to be such a file that reflects the correct state of the resources. The file will have to be kept updated to reflect the current resources. When a new policy is activated at a time other than during WLM startup, the WLM issues message "IWM001I WORKLOAD MANGEMENT POLICY XXXXXXXX NOW IN EFFECT". If the file is updated before activating the new policy, automation will be able to set the resource state. Otherwise, the modify commands will have to be entered manually.

**COMMUNICATING SCHEDULING ENVIRONMENTS**

Scheduling environments and their associated resources must be communicated

**Figure 1: Choose Service Definition Panel**

```
File Help
-----
Command ===> _____

                Choose Service Definition

                Select one of the following options.
                3  1. Read saved definition
                   2. Extract definition from WLM
                     couple data set
                   3. Create new definition

                ENTER to continue
```

**Figure 2: WLM Definition Menu**

```
File Utilities Notes Options Help
-----
Functionality LEVEL001          Definition Menu          WLM Appl LEVEL004
Command ===> _____

Definition data set . . . : none

Definition name . . . . . _____ (Required)
Description . . . . . _____

Select one of the
following options. . . . . 10_
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 3: Define a Scheduling Environment or Resource(s)**

```
                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 4: Create a Scheduling Environment**

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

Scheduling Environment Name  OSVS_COBOL          Required
Description . . . . . OSVS COBOL Compiler

Action Codes: A=Add D=Delete

                Required
Action Resource Name State Resource Description
a_ _____
```

to the user community. A company standards manual is the ideal place to describe these environments. Another possibility is a JESNEWS-like facility that lists the environments when requested. A third possibility is to allow users to display the scheduling environments through SDSF. Bear in mind that a scheduling environment becomes fairly permanent once users change their JCL to request it.

**ERROR MESSAGES**

WLM error messages are documented in “OS/390 MVS Programming: Workload Management Services” not in “Messages and Codes.” The messages are unfriendly to the point of hostility. While attempting to install a definition we received the following dialog message:

*Install failed, service definition is not valid. Validation reason code: 3D25, Validation offset: 1068. (IWMAM044)*

Seeking clarification, we read the explanation for message IWMAM044 and found that reason 3D25 indicates that “In the scheduling environment to resource connection, the scheduling environment name (SVSEA\_SR\_SCHENV\_NAME) is not in the list of defined scheduling environments.” Still clueless as to what was wrong, we ran the IEHIBALL utility several times seeking obvious errors. This too proved futile so we trashed everything and started over. Suffice it to say the error messages and online help need work.

**SDSF**

SDSF 1.8 provides support for Resource Affinity Scheduling. Figure 14 displays the status of the scheduling environments and shows state of resources. Since this snapshot was taken as resource states were being set, only the 3420\_TAPE scheduling environment is available to System P390.

**NEXT STEPS**

Resource Affinity Scheduling has potential, but some shortcomings need to be addressed:

1. The requirement to identify systems with resources through operator commands is unacceptable. Since the state of resources must be known in advance in order to automate the modify commands, it makes sense to define the state for each system in the Sysplex on the Resource Definition Panel.

**Figure 5: Add a Resource to a Scheduling Environment**

```

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

Selection For Scheduling Environment OSVS_COBOL

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

Action  Resource Name  In Use  Resource Description
_a_
    
```

**Figure 6: Define a Resource**

```

Define Resource

Command ==>
Resource name . . . . . OSVS_COBOL Required
Resource Description . . . . . Antique COBOL Compiler
    
```

**Figure 7: Select a Resource to be Added to the Scheduling Environment**

```

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

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

Action  Resource Name  In Use  Resource Description
_s_    OSVS_COBOL      Antiqu
    
```

**Figure 8: Indicate the Required State of a Resource**

```

Scheduling-Environments  Notes  Options  Help
-----
Command ==>

Scheduling Environment Name  OSVS_COBOL  Required
Description . . . . .

Action Codes: A=Add  D=Delete  State for a resource must be specified.
Specify either ON or OFF. (IWMAM676)

Action  Resource Name  State  Resource Description
_      OSVS_COBOL    ON
_      PROD_ONLINE   OFF
    
```

**Figure 9: Scheduling Environment List After Creating OSVS\_COBOL**

```

Scheduling-Environments  Notes  Options  Resources  Help
-----
Command ==>
Scheduling Environment Selection ListRow  1 to 1 of 1

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

Action  Scheduling Environment Name  Description
_      OSVS_COBOL

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

2. Scheduling scope is limited to a single Sysplex. Since most shops will have multiple Sysplexes, a scheduling environment should identify the Sysplex to which it pertains so that if a scheduling environment is moved to another Sysplex the requesting job will automatically be routed.

3. An improved help facility with improved error diagnoses is required.

**MAINTENANCE**

Review info APAR II10760 "JES2 Maintenance for OS/390 R4." We recommend PTFs be applied for the following APARs:

- OW26619 PER FINAL FUNCTION TEST CORRECTIONS FOR WLM AND SRM COMPONENTS IN OS/390 R4
- OW30336 PER WLM INITIATOR NOT STARTED FOR JOB THAT USES SYSAFF OR SCHENV
- OW28859 PER SDSF JOB INFO POPUP DOES NOT DISPLAY SYSTEMS WHERE SCHENV IS AVAILABLE

**RACF SUPPORT**

RACF APAR OW25727 (closed 97/05/19) provides the ability to restrict execution of a particular resource to a specific system or systems. Although this APAR is not required for Resource Affinity Scheduling, it can guarantee that products are only used on the systems for which they are licensed.

```
PERMIT program_name CLASS(PROGRAM)
ID(id) ACCESS(READ)
WHEN(SYSID(system_name))
```

**CONCLUSION**

Despite some shortcomings, Resource Affinity Scheduling is a welcome addition to OS/390. A number of shops that modified JES2 to achieve similar functionality can remove their modifications. Shops without these modifications now have functionality that was previously lacking. Historically, IBM products have improved significantly in response to user requirements. We expect the same for Resource Affinity Scheduling.

**BIBLIOGRAPHY**

- OS/390 MVS Planning: Workload Management*; GC28-1761-03; IBM Corporation
- OS/390 MVS JCL Reference*; GC28-1757-04; IBM Corporation
- OS/390 MVS Programming: Workload Management Services*; GC28-1773-02; IBM Corporation
- OS/390 MVS System Commands*; GC28-1781-03; IBM Corporation

**Figure 10: Scheduling Environment List After Multiple Environments Were Created**

```
Scheduling-Environments  Notes  Options  Resources  Help
-----
Scheduling Environment Selection ListRow      1 to 4 of 4
Command ==>

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

Action  Scheduling Environment Name  Description
--
      DEVELOPMENT                  Standard Development Env
      FIN_PROD                      Financial Production Env
      3420_TAPE                     Round Tape
      OSVS_COBOL                    Antique COBOL Compiler
```

**Figure 11: Resource Definition List With Multiple Resources Defined**

```
Resources  Notes  Options  XREF  Help
-----
Resource Definition List      Row 1 to 10 of 10
Command ==>

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

Action  Resource Name  In Use  Resource Description
--
      DB2_V4      YES      DB2 Version 4
      DB2_V5      YES      DB2 Version 5
      FIN_PROD_DB  YES      Financial Division Prod DB
      FIN_TEST_DB  YES      Financial Division Test Database
      OPTICAL      YES      3995 Optical Storage
      OSVS_COBOL   YES      Antique COBOL Compiler
      OS390_COBOL  YES      Current COBOL Compiler
      PROD_ONLINE  YES      Production Online Processing
      ROUND_TAPE   YES      3420 Tape Drives
      SAS          YES      Statistical Analysis System
```

**Figure 12: Required Resource States**

```
Scheduling Environment Name. . DEVELOPMENT
Description. . . . . Standard Development Env

Resource Name  State  Resource Description
-----
DB2_V5         ON     DB2 Version 5
FIN_TEST_DB    ON     Financial Division Test Database
OS390_COBOL    ON     Current COBOL Compiler
PROD_ONLINE    OFF    Production Online Processing
SAS            ON     Statistical Analysis System

Scheduling Environment Name. . FIN_PROD
Description. . . . . Financial Production Env

Resource Name  State  Resource Description
-----
DB2_V4         ON     DB2 Version 4
FIN_PROD_DB    ON     Financial Division Prod DB
OPTICAL        ON     3995 Optical Storage

Scheduling Environment Name. . 3420_TAPE
Description. . . . . 3420 Tape

Resource Name  State  Resource Description
-----
ROUND_TAPE     ON     3420 Tape Drives

Scheduling Environment Name. . OSVS_COBOL
Description. . . . . Antique COBOL Compiler

Resource Name  State  Resource Description
-----
OSVS_COBOL     ON     Antique COBOL Compiler
PROD_ONLINE    OFF    Production Online Processing
```

Figure 13: Sample Modify Commands to Set Resource State

```
f wlm,resource=3420tape,on
IWM039I RESOURCE 3420TAPE IS NOW IN THE ON STATE
f wlm,resource=PROD_ONLINE,ON
IWM039I RESOURCE PROD_ONLINE IS NOW IN THE ON STATE
F WLM,resource=optical,on
IWM039I RESOURCE OPTICAL IS NOW IN THE ON STATE
f wlm,resource=osvs_cobol,on
IWM039I RESOURCE OSVS_COBOL IS NOW IN THE ON STATE
f wlm,resource=os390_cobol,on
IWM039I RESOURCE OS390_COBOL IS NOW IN THE ON STATE
f wlm,resource=fin_test_db,on
IWM039I RESOURCE FIN_TEST_DB IS NOW IN THE ON STATE
```

Figure 14: SDSF Scheduling Environment Display

```
Display Filter View Print Options Help
-----
SDSF SCHEDULING ENVIRONMENT DISPLAY MAS SYSTEMS      LINE 1-4 (4)
COMMAND INPUT ==>                                  SCROLL ==> PAGE
PREFIX=RPS*  DEST=(ALL)  OWNER=*
NP   SCHEDULING-ENV      DESCRIPTION          SYSTEMS
     DEVELOPMENT        Standard Development Env
     FIN_PROD           Financial Production Env
     3420_TAPE          3420 Tape          P390
     OSVS_COBOL        Antique COBOL Compiler

SDSF RESOURCE DISPLAY MAS SYSTEMS
COMMAND INPUT ==>
PREFIX=RPS*  DEST=(ALL)  OWNER=*
NP   RESOURCE          P390
     DB2_V4             RESET
     DB2_V5             RESET
     FIN_PROD_DB        RESET
     FIN_TEST_DB        ON
     OPTICAL            ON
     OSVS_COBOL         ON
     OS390_COBOL        ON
     PROD_ONLINE        ON
     ROUND_TAPE         ON
     SAS                 RESET
```

*OS/390 JES2 Commands*; GC28-1790-02; IBM Corporation

*SDSF Guide and Reference*; SC28-1622-02; IBM Corporation

*Goal Based Initiator Management*, John Arwe, IBM Corporation\*

*WLM: Resource Affinity Scheduling*; Mike Cox, IBM Corporation\*

*WLM: Managed Initiators*; Mike Cox, IBM Corporation\* 

\* Available at [www.s390.ibm.com/wlm/documents/](http://www.s390.ibm.com/wlm/documents/)

*Opinions stated in this article are the author's and do not necessarily reflect those of Programart.*

**NaSPA member R.P. (Bob) Shannon is a software developer for Programart Corporation. He has more than 20 years of data processing experience, primarily in the areas of MVS systems programming and performance. He is an MVS Cluster Leader at SHARE, and can be reached at [shannon@programart.com](mailto:shannon@programart.com).**

©1998 Technical Enterprises, Inc. For reprints of this document contact [sales@naspa.net](mailto:sales@naspa.net).