

**Improve your Availability  
With  
Sysplex Failure Management**

**Riaz Ahmad  
IBM Washington Systems Center  
Gaithersburg, Maryland**

Customer experience has shown that the ability to detect and recover from component failures is critical for high availability in clustered systems like OS/390 Parallel Sysplex.

Several availability studies have demonstrated that a percentage of Parallel Sysplex outages can be avoided or reduced if sysplex failure management (SFM) is in effect. This, and the fact that SFM has been enhanced with APAR OW30926, should be important reasons why you need to implement SFM in your Parallel Sysplex. This paper describes what SFM is and why Parallel Sysplex customers should use it.

Nearly all users of computing systems are making availability a requirement. The strong demand for higher performance is being met with clustered architectures. The redundancy inherent in clustered systems offers the opportunity to provide increased levels of availability. However, the presence of more and larger components increases the chances that something can go wrong and tends to decrease availability.

Sysplex failure management (SFM) is an *availability function* that is applicable to both basic and Parallel Sysplex environments. SFM is integrated into the OS/390 base and it is recommended that all installations take advantage of SFM and implement it when they are configured in either a basic or a Parallel Sysplex. Although you can use SFM in a sysplex without a coupling facility, to take advantage of the full range of failure management capabilities that SFM offers, a coupling facility must be configured in the sysplex.

Sysplex failure management allows you to define a sysplex-wide policy that specifies the actions OS/390 is to take when certain failures occur in the sysplex. A number of situations might occur during the operation of a sysplex when one or more systems need to be removed so that the remaining systems can continue to do work. The goals of failure management in a sysplex are to minimize the impact that a failing system might have on the sysplex workload so that work can continue with little or no operator intervention. The type of failures that are handled by SFM are:

- XCF signaling connectivity failures in the sysplex
- Status update missing condition
- PR/SM Reconfiguration

### **XCF Signalling Connectivity Failures**

All systems in the sysplex must have full signaling connectivity at all times. Loss of signaling connectivity between sysplex members can result in one or more systems being removed from the sysplex so that the remaining members in the sysplex retain full signaling connectivity to one another.

### **Status Update Missing Condition**

Each system periodically updates its own status and monitors the status of other systems in the sysplex. The status of the systems is maintained in a Couple Data Set

(CDS) on DASD. A *status update missing condition* occurs when a system in the sysplex does not update its status information within the failure detection interval, specified on the INTERVAL keyword in COUPLExx, and appears dormant. SFM allows you to specify how a system is to respond to this condition. System isolation allows a system to be removed from the sysplex as a result of status update missing condition, without operator intervention, while ensuring that the data integrity in the sysplex is preserved. Specifically, system isolation terminates I/O and coupling facility accesses, resets the channel subsystem, and loads a non-restartable wait state on the failing system, thus ensuring that the system is unable to corrupt shared resources.

Additionally, with APAR OW30926, which is available for OS/390 Release 2 or higher, and SFM active with ISOLATE specified, the operator will be prompted if a *status update missing condition* occurs for a system, but XCF signalling is still active with that system. For instance, when temporary conditions occur in which the system status cannot be updated, messages IXC427A and IXC426D are issued if XCF signalling is active. The operator or system programmer can then evaluate the state of the system targeted in the messages (through display commands and performance monitors for example) and determine whether to remove the system from the Parallel Sysplex. If the operator or system programmer determines that the system is performing as expected, the operator or system programmer does not respond to IXC426D and the message is DOMed once the system starts updating status, thus avoiding the removal of the system. Otherwise, the operator or system programmer responds to message IXC426D, and removes the system from the Parallel Sysplex.

## **PR/SM Reconfiguration**

After a system running in PR/SM partition is removed from the sysplex, SFM allows a remaining system in the sysplex to reconfigure processor central and expanded storage for use by the remaining system.

## **The SFM Policy**

For a sysplex to take advantage of SFM, the policy must be active on all systems in the sysplex. That is:

- All systems must be at the SP Version 5 level or higher.
- An SFM policy must be started in the SFM couple data set.
- All systems must have connectivity to the SFM couple data set.

If any system loses access to the SFM couple data sets, the policy becomes inactive in the sysplex. If that system regains access to the SFM couple data set, SFM automatically becomes active again in the sysplex.

The ability of SFM to determine which system to remove from the sysplex in a connectivity loss situation is influenced by a **weight** that has been assigned to each system according to its importance in the sysplex.

SFM determines which systems to keep, which to remove from the sysplex and attempts to implement that decision by system isolation. If a system that is being removed is connected to a coupling facility that is also connected to an active system, SFM can reconfigure the sysplex by isolating and removing the system without operator intervention. Without a coupling facility, a system cannot be isolated automatically and a message prompts the operator to reset the system manually before it is removed from the sysplex.

### **Structure Rebuilds**

When there is a loss of connectivity between a coupling facility structure and its connector, it makes sense for the structure that has suffered a connectivity loss to rebuild in another available coupling facility, thereby allowing the user of the structure to connect to it in the new location and continue processing. However, not all structures support this rebuild processing, and it is also advantageous to allow an installation to control which structures will be rebuilt in a failure situation. The way an installation can influence which structures should be rebuilt is to use the REBUILDPERCENT keyword in the coupling facility resource management (CFRM) policy when you define the structures. The value specified for this keyword is used as part of a formula that also incorporates the system weights.

If the rebuild value calculated by the system is equal to or greater than the value specified by the installation on the REBUILDPERCENT keyword in the CFRM policy, then and only then will the structure rebuild operation begin.

Some structures, such as XCF structures, are considered so important to the integrity of the sysplex that they will always initiate a rebuild themselves as opposed to having the system initiate the action regardless of whether or not there is an SFM policy active. If an SFM policy is not active in the sysplex, a structure rebuild will not be initiated for any structure even if the structure supports REBUILD.

With OS/390 Release 3 and higher systems and fix for APAR OW30814 installed, the loss of coupling facility connectivity processing has changed. In the event there is no SFM policy active in the Parallel Sysplex and there is a loss of CF connectivity, a structure rebuild will be initiated by OS/390 R3 and higher systems as long as the structure users support the structure rebuild.

Enhancements have been made to improve the coupling facility selection process used during structure allocation by factoring in the SFM system weights and the connectivity that each system has to the coupling facility. . Also, structure rebuild processing has been improved by insuring that a rebuilt structure due to loss of connectivity has better

connectivity to the systems in the sysplex than the old structure. If no SFM policy is active, however, all systems are treated as having equal weights when determining the suitability of a coupling facility for initial structure allocation.

## **Benefits and Summary**

It is evident that if you are running in a Parallel Sysplex, you *should* have an active SFM policy, and you should specify appropriate weights for your systems according to their importance in the sysplex.

In a Parallel Sysplex, if you are using XCF signalling through multiple XCF signalling structures, it is unlikely that you will suffer a total signalling connectivity loss. If there is a loss of signalling connectivity, sysplex failure management will help you isolate the image and partition it out of the Parallel Sysplex according to your SFM policy and without the operator involvement.

In the case where the status update missing condition has been detected, SFM will help isolate the image using the coupling facility fencing services without operator intervention. This is even more critical when you are in a data sharing environment. It is very likely, that the image missing the status update is not in good health and is holding resources that are critical to the continuous operation of the Parallel Sysplex as well as data sharing environment. Steps must be taken to remove this image as soon as possible and release the resources that were held by this image. With the help of Automatic Restart Manager (ARM), you can restart the failing workload on other images within the sysplex for a swift recovery.

Implementing sysplex failure management will provide the ability to detect and recover from component failures in a timely fashion and therefore provide a sound foundation for a highly available Parallel Sysplex.