IBM @server zSeries

Value of Resource Sharing

October 2001

Table of Contents

Introduction — zSeries Parallel Sysplex-Systems Enablement zSeries Resource Sharing

A significant value proposition for zSeries Parallel Sysplex clustering has emerged as exploitation and hardware integration of coupling technology has evolved over time.

A number of base z/OS™ and OS/390® components have discovered that the IBM® S/390® Coupling Facility shared storage provides an excellent medium for sharing component information for the purpose of multi-system resource management. This exploitation called IBM @server zSeries™ Resource Sharing enables sharing of physical resources such as files, tape drives, consoles, catalogs, etc. with significant improvements in cost, performance and simplified systems management. This is NOT to be confused with Parallel Sysplex® data sharing by the database subsystems. zSeries Resource Sharing delivers immediate value even for customers who are not leveraging data sharing, through native system exploitation delivered with the base z/OS and OS/390 software stack.

zSeries Resource Sharing delivers immediate value even for customers who are not leveraging data sharing, through native system exploitation delivered with the base z/OS software stack.

In 1994, systems enablement (resource sharing) was just the first step to data sharing and may not have been considered relevant for those customers who did not have a requirement for database data sharing (continuous availability). With all the latest enhancements in hardware, software and microcode this zSeries Resource Sharing support allows midsize customers and those doing server consolidation, ERP or Parallel Web serving to embrace Parallel Sysplex technology without requiring them to move to a database data sharing environment to gain the benefits mentioned in this paper.

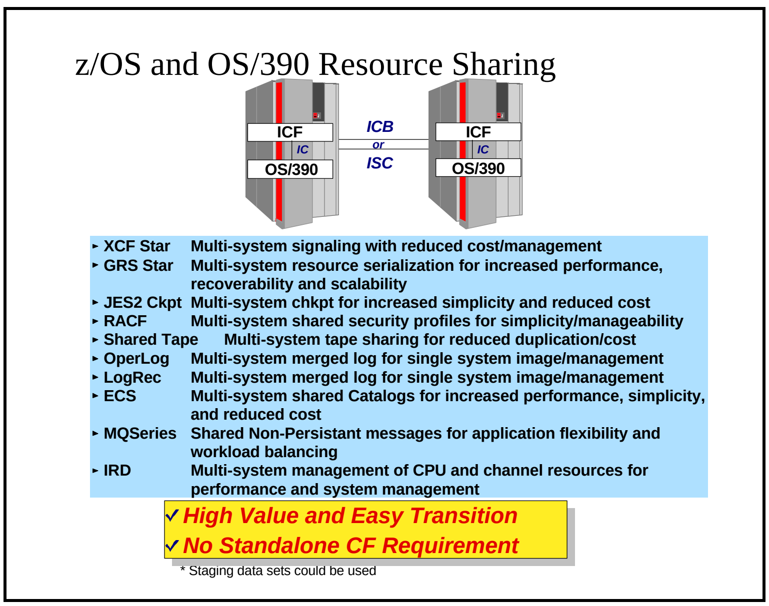

In addition to base z/OS and OS/390 software exploitation, a number of systems integration enhancements were made in the hardware and microcode, aimed at reducing the complexity and cost of Parallel Sysplex technology. This began with HiPerlinks with the Generation 3 (G3) CMOS machines in 1997, which enabled customers to replace their very complex CTC configurations with XCF Signaling structures. Next, Internal Coupling Facilities (ICF), were introduced on G4 which enables customers to replace their standalone CFs with ICFs in a resource sharing environment. Finally, Internal Cluster Busses (ICB) and Integrated Coupling Channels (IC) starting on the G5/G6, significantly optimize the performance and cost benefits one gets when sharing resources.

S/390 Resource Sharing is an attractive legitimate configuration ***target*** *objective for many customers TODAY.*

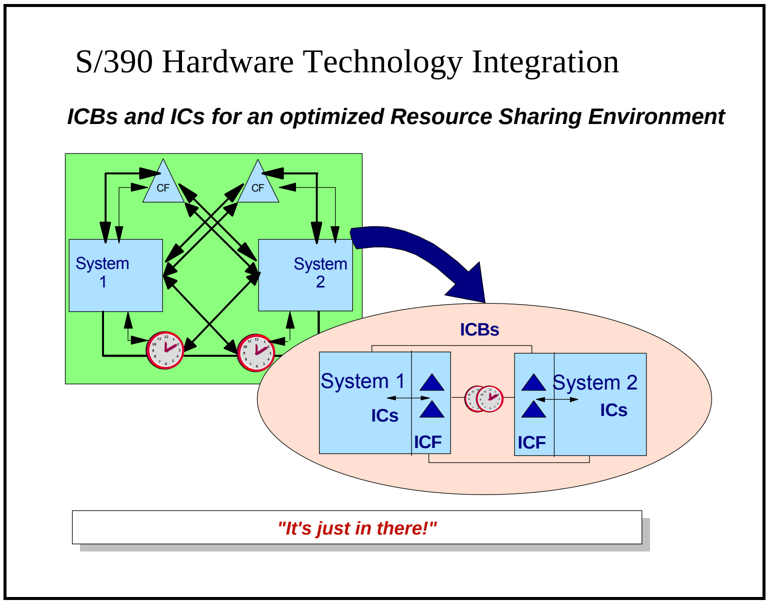

S/390 Resource Sharing is the superior way to run multiple z/OS and OS/390 systems.

No standalone CF required

These functions, along with numerous availability and performance enhancements for both single server and multi-server environments, make zSeries and S/390 resource sharing a valid end target environment for Parallel Sysplex customers today.

S/390 Resource Sharing — Value in Both Multiple Server and Single Server Environments

Combines state of the art Parallel Sysplex technology with LPAR PR/SM ™ hipervisor to provide:

- *Ease of use, simplified multi system management/ operations and reduced costs*
- *Single Logical Systems Resource Image through sharing of many of the base systems functions/ hardware*
- *Value in both a multi-server and single-server environment*

As businesses rapidly grow, so too do their Information Technology (IT) resources. This growth can be due to business volumes, data center consolidations, or business mergers/acquisitions which can lead to an increase in the number of applications and z/OS and OS/390 system images and/or LPARs (Logical Partitions) as a natural result. In many instances IT organizations will consolidate by just adding another system image to the configuration without change to the application workload or existing systems. An unfortunate side-effect of this approach is it can force the installation to duplicate many of the resources used by each physical system (or in each LPAR image) such as tape drives, consoles, printers, networks and common data sets. This often leads to an environment that is highly complex and expensive to maintain and operate.

In a non-resource sharing environment, each OS/390 system image has its own set of systems resources: point-to-point CTCs configured between every pair of systems for communication between systems; dedicated master console and 3174 control unit; dedicated tape drives; logs; redundant master catalogs and user catalogs.

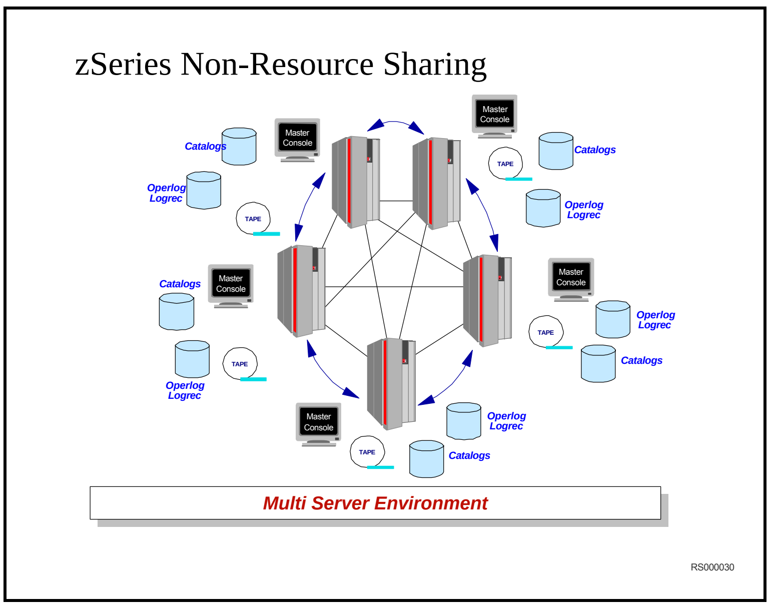

As this picture demonstrates, this environment is not very scalable or manageable, and can be very costly. As new systems are added into this configuration, the complexity worsens and is usually a disruptive process.

A resource sharing environment requires no failure independent CFs and therefore can run entirely with Internal Coupling Facilities. Best yet, the majority of resource sharing system exploitations are delivered in OS/390 itself. It's just in there!

zSeries Resource Sharing allows customers to share common resources across multiple z/OS and OS/390 images, regardless of whether they are running native, as multiple LPARs on a single server or across multiple zSeries or S/390 servers. Resource Sharing exploits S/390 Coupling Technology to allow multiple OS/390 images to manage and share a resource as if it were a single resource attached to a single image. Being able to efficiently share resources between multiple images enables IT organizations to utilize existing resources more efficiently and reduces or eliminates the need to purchase additional hardware and software to support multiple systems. In addition, z/OS Resource Sharing can reduce the overall cost and complexity in managing multiple system images.

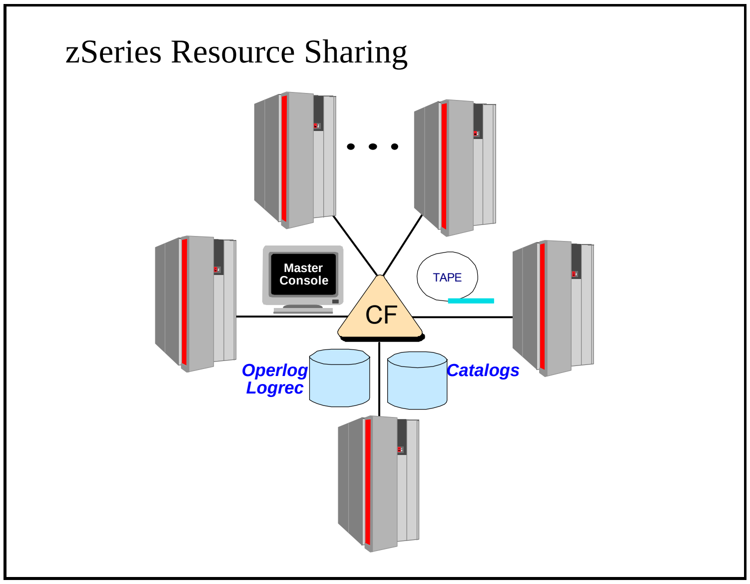

Inter-system communication improves in terms of reduced configuration complexity and improved performance, as systems are able to send messages through the Coupling Facility in "STAR" mode thus eliminating CTC point to point connections. Serialization of shared files and other resources across the multiple systems gains a substantial performance boost through GRS STAR mode, where lock requests are sent to the CF and passing locks around a "ring" is eliminated. Catalog sharing across systems is enabled through shared caching of index records for high performance in the CF. Tape drives can be managed as a shared resource pool across the multiple systems through the Coupling Facility. System logs for systems operations and diagnostics can be merged real-time to provide a single system image for all of the systems in the configuration. The Master Console can be shared across systems, eliminating the cost of redundant consoles and associated 3174 control units.

The same Resource Sharing benefits afforded in a multi-server environment are applicable to single server multiple z/OS and OS/390 image configurations as well.

While it may come as a surprise, the same Resource Sharing benefits afforded in a multi-server environment are applicable to single server multiple z/OS and OS/390 image (LPAR) configurations as well. This is due, in large part, to the effectiveness of PR/SM to make the single server appear as a separate system to each instance of OS/390 running on the server. While this has obvious benefits in terms of isolation and granular sharing of CPU, I/O and Memory resources, it also has software systems management drawbacks. In a multiple-LPAR single physical server, every operating system believes it is running on a totally separate server. Therefore, many of the management and cost issues associated managing a multi-server configuration apply.

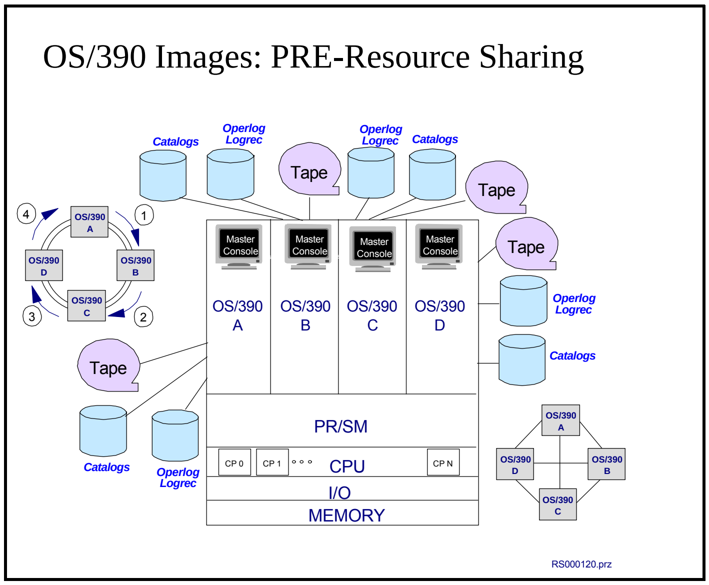

Without the benefits of zSeries Resource Sharing, each operating system image requires its own master console and a dedicated 3174 control unit. Each image requires point-to-point logical definitions for every other system image it must communicate with, and all images have their own Operations Log and Sys1.Logrec data sets. Each system image maintains a set of dedicated Tape Drives, and every z/OS and OS/390 image has its own Master and User Catalogs. Serialization of access to shared file and DASD resources is through some form of physical DISK serialization or GRS Ring serialization protocols.

In short, from a hardware perspective, a single server multiple image (LPAR) configuration is economical and simpler to manage than multiple physical boxes. However, from a software perspective, all of the same issues associated with resource replication, configuration complexity and performance bottlenecks in sharing physical resources apply.

Resource sharing cost/performance optimization is achieved starting on IBM 9672 G5 servers through the integration of Internal Coupling (IC) links and Internal Coupling Facilities (ICF).

S/390 Resource Sharing, leveraging Parallel Sysplex Coupling Facility technology, can significantly reduce the cost and complexity of managing a multiple image single server configuration. Further, resource sharing cost/performance optimization is achieved starting on IBM 9672 G5 servers through the integration of Internal Coupling (IC) links and Internal Coupling Facilities (ICF).

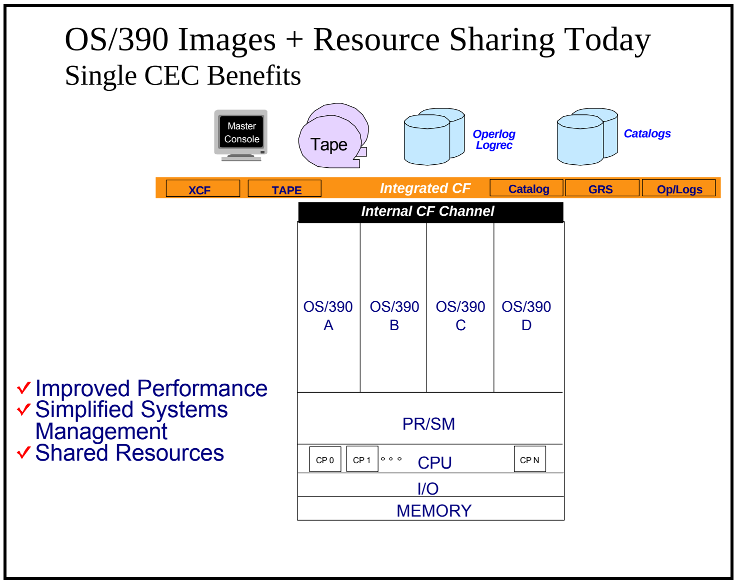

To better understand the value of resource sharing, let's examine some of the major components delivering value in the resource sharing environment. We will begin by discussing the benefits of using XCF Star (XCF signaling), available with OS/390 R1.

S/390 Resource Sharing – Using XCF Star (OS/390 Release 1)

The cross system coupling facility (XCF) component of OS/390 allows authorized applications to communicate (XCF signaling) with applications on other OS/390 images in a sysplex. In a Parallel Sysplex cluster, signaling can be achieved through ESCON® Channel to Channel connections (CTCs), Coupling Facility structures, or a combination of both CTCs and CF structures. CTC signaling paths are unidirectional and require at least four signaling paths between each system (two inbound and two outbound paths). CTC connections become more complex to manage as the number of systems in the configuration increases. CF signaling paths are bi-directional and much less complex to maintain. With XCF Star (XCF Signaling), each system in the sysplex can automatically discover its connectivity to other systems via the CF, without having to define point-to-point connections on every system in the configuration. Furthermore, in the event that a signaling path or signaling structure is lost, the recovery of CF signaling structures is automated and greatly simplified if more than one structure is allocated (IBM recommends that two structures be allocated, in different CFs). If using a combination of CTCs and CF signaling paths, XCF will always chose the fastest path for signaling.

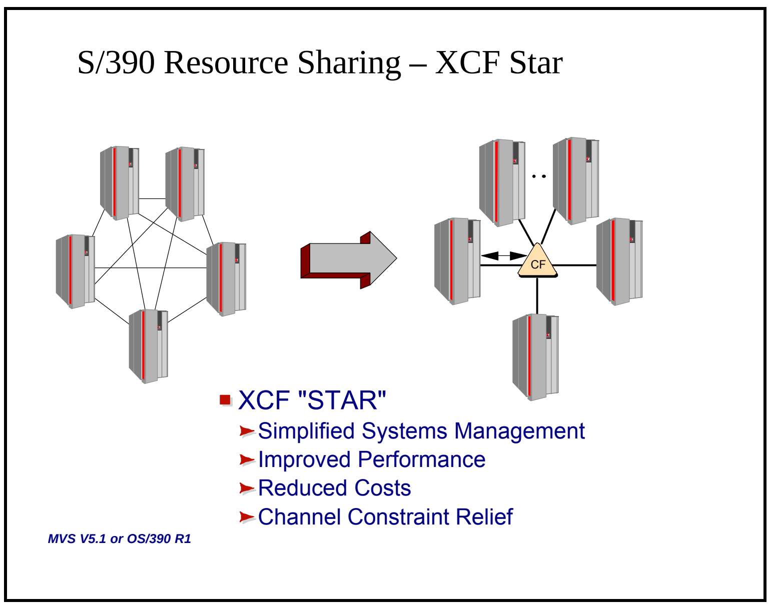

When XCF signaling is implemented using HiPerLink technology from IBM, performance is comparable to CTCs. With coupling technology improvements delivered with faster zSeries and 9672 engines, HiPerLinks, and the introduction of Integrated Cluster Bus (ICB) on the G5 servers, XCF signaling through CF structures outperforms CTC signaling. Further, XCF Star allows customers to consolidate CTC links used by VTAM®, JES, and GRS onto XCF Star, freeing up channel paths for other use in constrained configurations. These advancements in CF coupling technologies combined with simplified systems management, ease of recovery, and better cost efficiencies, make XCF Star the clear choice for configuring XCF signaling paths in a Parallel Sysplex cluster.

S/390 Resource Sharing – Shared Console

Every z/OS and OS/390 image is required to have a master console to operate and manage the subsystems and applications residing on that image. In a multi-system environment, operators must monitor and control multiple images through each image's individual console. This creates a difficult and costly situation to manage, often leading to increased staffing and operational errors that result in outages or interruptions in service. Cost and complexity is further aggravated by the fact that the master console is attached to a system via a dedicated 3174 control unit, which cannot be shared even across LPARs (images) on the same physical server.

With Parallel Sysplex resource sharing, the cost and operation of multiple systems is greatly optimized through the use of "shared consoles." Parallel Sysplex clustering exploits XCF signaling technology to enable messages and system commands to be seamlessly transported between systems in the sysplex. This eliminates the need to physically attach a console and 3174 controller to every image in the sysplex. Minimally, only one master console and one alternate console is required for the entire Parallel Sysplex cluster, regardless of the number of images in the sysplex. Additional consoles can be configured for greater availability/redundancy as desired.

Because of the console limitations and the changes to console management, it is recommended that you thoroughly understand the roles of master and alternate consoles in a Parallel Sysplex environment. There are several publications available to assist you with planning your console configuration, including: SG24-4626, *OS/390 MVS™ Multi-system Consoles Implementing Sysplex Operations* and GC28-1441, *MVS/ESA™ V5 Planning: Operations*.

S/390 Resource Sharing – Shared Console

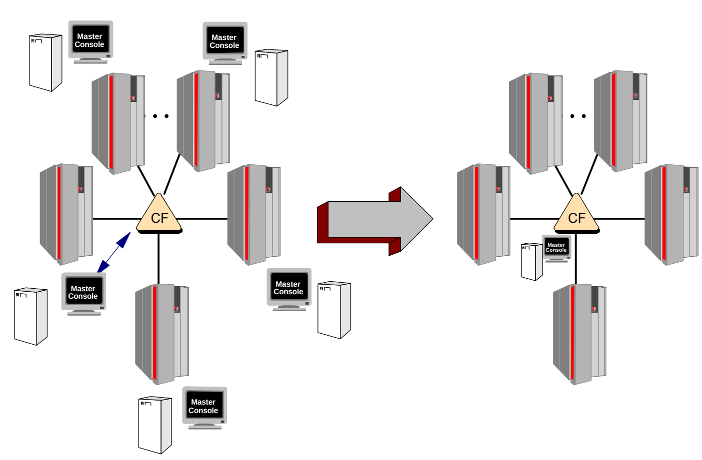

"GRS Star was like winning the lotto for our users. [Prior to using GRS Star we were] having serious problems keeping up with the number of enqueues being issued. This would frequently fail and the restart process would be extremely painful and unproductive for at least an hour while all the images resync'd. So we decided to move to GRS Star.

This was accomplished over one weekend.

I saw tso logon's go from minutes to seconds and a daily batch jobstream go from 14 hours to less than 2 hours.

We are now converting S/SYSZVVDS reserves to global enqueues and this has lessened the number of deadlocks between jobs by 50% and has enhanced the performance of the catalog address space."

Dean Tessar

Hewitt Assoc.

S/390 Parallel Sysplex – GRS Star (OS/390 Release 2)

Global Resource Serialization (GRS) is the operating system component that provides serialized access to shared or "global" resources in multitasking and multi-system environments. Systems that need to share resources are combined into a GRS complex to provide the serialization controls needed to ensure the integrity of global resources, such as files and disks across multiple OS/390 images. The traditional mechanism to provide serialization for shared resources was GRS Ring. With a ring complex, GRS used CTC links to pass a GRS RSA (Ring System Authority) between systems in a round-robin or ring fashion. Only one RSA can exist in a GRS ring complex. Each system had to wait for the arrival of the RSA to send serialization requests for resources to other systems. Only when the RSA was returned could the originating system begin to process that resource request. As the number of systems in the ring complex increased, the delays and overhead required to pass this RSA across all of the systems made sharing resources less efficient and more costly. Additionally, each system has to maintain an in-storage copy of all the global resources allocated for all systems in the sysplex.

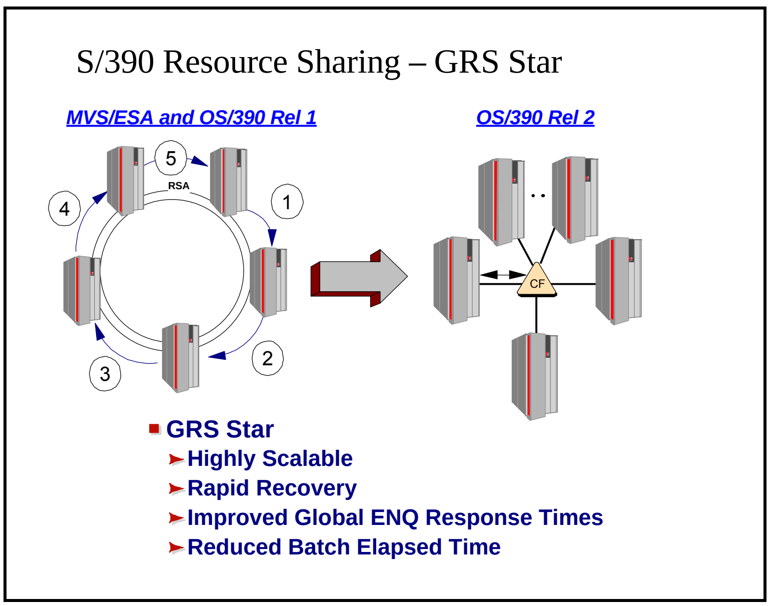

With the introduction of GRS Star in OS/390 Release 2, a new method of communicating global resource allocation requests was introduced. GRS Star places the Coupling Facility at the hub of the GRS complex and eliminates the delays and processing overhead inherent in the traditional ring complex.

Global resource access requests are now processed in microseconds (versus milliseconds), and the overhead associated with allocating a shared resource is limited to the requesting system only.

GRS Ring versus Star Throughput Example

DB2 Tablespace Cleanup Job
(elapsed time h:mm)

	Job 1	Job 2
RESMIL (10)	4:39	5:21
RESMIL (1)	1:11	1:26
Star Mode	0:09	0:12

Elapsed time went from hours to minutes

The performance data was determined in a controlled environment. The results obtained in other operating environments may vary.

GRS Star is designed to provide measurable improvements to overall resource sharing performance, delivering a technology that is cost efficient and highly scalable:

- The resource is registered synchronously in the Coupling Facility at the time the request is made; no need to wait for the RSA.
- Each system only needs to maintain a local view of the resources it allocated.

Furthermore, Star is designed to provide significant improvements to availability and recovery times, by eliminating the extensive processing needs to reinitialize or reform GRS rings and queues on each system when a system joins or leaves the complex.

S/390 Parallel Sysplex – Shared Tape Drives

Unlike DASD devices that can be online and accessed simultaneously by multiple systems, tape devices must be dedicated to one system at a time. As customers move to multi-system environments to increase processing capacity and improve availability they are faced with costly alternatives for handling tape processing. Some of these alternatives include; incrementally increasing tape capacity by adding more physical drives to support additional systems; or having operators manually manage tape devices by moving available devices between systems; or purchasing third party software solutions to manage tape devices.

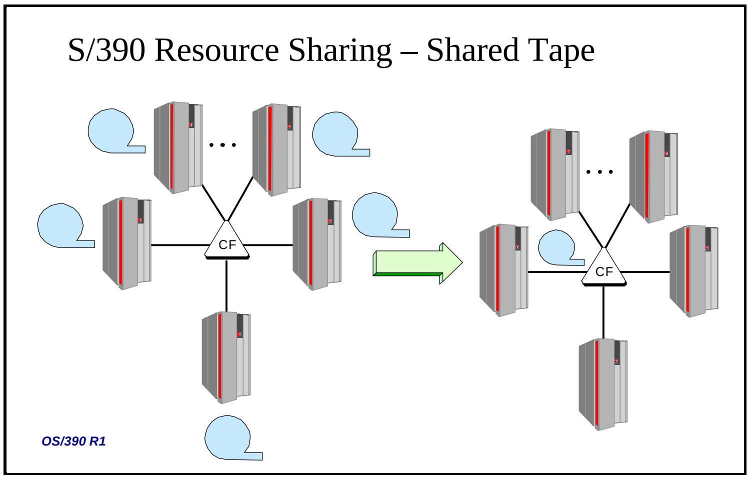

Beginning with MVS/ESA Version 5.2, and subsequent z/OS and OS/390 releases, tape units no longer need to be dedicated to a single system. Automatic Tape Switching manages the allocation of shared tape devices between two or more systems in a Parallel Sysplex cluster by automatically switching devices between systems. To enable dynamic tape sharing, devices that can be automatically switched between multiple systems are managed through a structure in the Coupling Facility. Systems communicate with each other through the CF structure to track and maintain the availability of shared tape devices, and to select and automatically allocate the appropriate device needed to satisfy requests. Because the same tape devices are available to multiple systems and can be automatically switched between systems, jobs that perform tape processing can be dynamically balanced across multiple systems automatically, allowing customers to utilize existing tape and CPU resources more efficiently.

In z/OS 1.2, the tape switching support has been rewritten to make it more reliable, robust, and serviceable. The new design makes use of GRS resources and XCF services to maintain serialization when allocating shared tape drives, dropping the use of the IEFAUTOS CF structure. In order to maximize the benefits of this function, GRS STAR is highly recommended.

S/390 Parallel Sysplex – Merged Logs

In z/OS and OS/390 environments applications and subsystems generate log data that is maintained in time sequence. In a multi-system environment, each system generates and writes logs independently of the others. In the event that the combined data from multiple systems needs to be accessed, the proper log data for each system must be identified and merged into time sequence order before it can be processed. The System Logger addresses the problem of log management in a multi-system environment by creating a single real-time merged log for log-streams such as Logrec and System console logs.

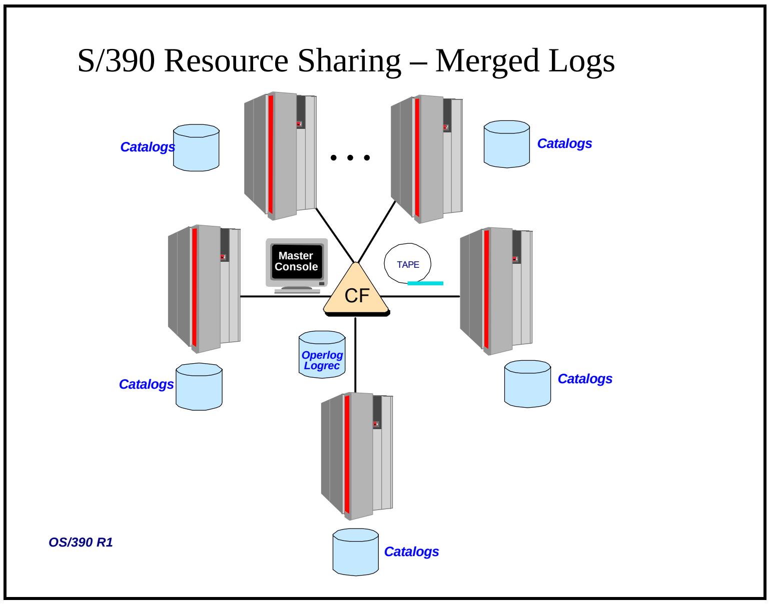

System Logger is a set of services that enables the merging of log data generated by several systems in a Parallel Sysplex cluster. It is implemented as a subsystem with its own address space. Applications that exploit system logger services (such as Operlog, Logrec, CICS® Transaction Server, and IMS™ for Shared Message Queues) write log data into a log stream, which is simply a collection of data. Data in a log stream is written to an interim storage area where log data can be accessed quickly without incurring I/O (Interim storage can be either a structure located in one or more Coupling Facilities, or in the case of DASD-Only log streams data is contained in local storage buffers and duplexed to a DASD staging area). When the interim storage area reaches a user defined threshold, the log data is automatically offloaded to DASD data sets where it is stored for longer term access. Once the data is offloaded to DASD the system logger component will keep track of the log data. Installations can define a retention period and have the system logger component automatically delete logs once the user-defined retention period has expired.

The OS/390 System Logger provides a Single Point of Control and Single System Image for common log data generated on multiple systems in a Parallel Sysplex configuration.

S/390 Parallel Sysplex – Enhanced Catalog Sharing

The Benefits of Enhanced Catalog Sharing are compounded when used in conjunction with GRS Star for optimal performance, availability, operations simplification and reduced DASD cost.

Each S/390 system requires access to a master catalog and user catalogs. In a multi-system environment catalog data sets can be shared across multiple z/OS and OS/390 images, which reduces the need for redundant hardware and eliminates configuration and operational complexity. However, each system image must have read/write access to the master and user catalogs. Catalog sharing requires that all changes to the catalog be communicated to all systems to ensure data integrity and data accessibility. To maintain the integrity of the catalog, each system uses a combination of device locking and multiple I/Os to access a "shared record" for each of the catalogs. Using this method to communicate changes to each z/OS or OS/390 image is costly and can have a negative impact on performance. As a result, many customers do not share catalogs across systems, even though they'd like to.

ECS support could decrease GRS CPU usage by as much as 33%, CAS CPU usage by 23%, and elapsed time for the job by 16% (results will vary depending on number of data sets being accessed, etc,...

Enhanced Catalog Sharing (ECS), available with OS/390 Release 7, allows you to "have your cake and eat it too." ECS enhances catalog sharing in a Parallel Sysplex cluster through the use of the Coupling Facility. ECS provides a new catalog sharing method that moves the contents of the "shared record" to a cache structure in the Coupling Facility. This allows OS/390 images to access this information at CF speeds, eliminating most of the GRS and I/O activity required to access shared catalogs. In a GRS Ring environment, this enhanced sharing capability is expected to decrease CPU usage for jobs accessing shared catalogs by as much as 40% and decrease elapsed time by as much as 50%. For GRS Star (recommendation is to run ECS with GRS Star) environments running older release of DFSMS™, ECS support could decrease GRS CPU usage by as much as 33%, CAS CPU usage by 23%, and elapsed time for the job by 16% (results will vary depending on number of data sets being accessed, etc.).

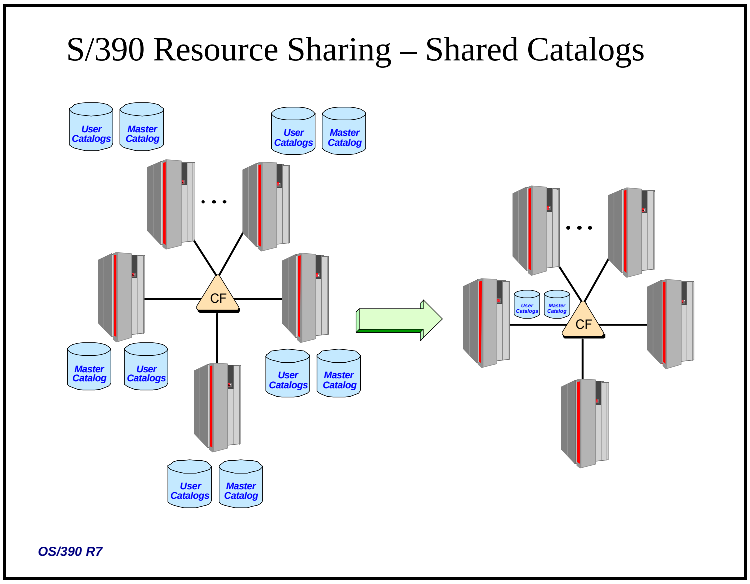

MQSeries Shared Message Queues

MQSeries® messaging enables applications to communicate across different operating system platforms by use of messages in a way that is straightforward and easy for programmers to implement. It takes care of network interfaces, assures delivery of messages, deals with communication protocols, and handles recovery after system problems.

MQSeries of OS/390 V5.2 has extended support to allow applications running on multiple queue managers in the same queue-sharing group anywhere in the Parallel Sysplex cluster to access shared queues for non-persistent messages up to 63KB long. This provides the advantages of:

- Application flexibility
- High availability
- High capacity
- Workload balancing

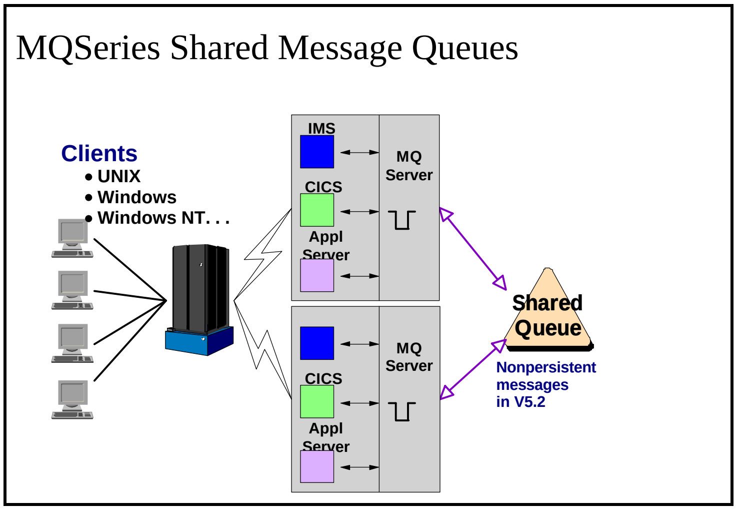

MQSeries applications running on these queue managers can MQPUT and MQGET non-persistent messages of up to 63 KB to the same shared queues as if were a local queue. Shared queue support for trigger first or trigger depth is based on committed messages only, and can start an application to run on any or every queue manager in the queue-sharing group for each trigger event. If a queue manager fails, all queue managers in the same queue-sharing group will cooperate to recover its shared queue-related work.

The shared queues can be accessed by applications running on any queue manager in the same queue-sharing group anywhere in the Parallel Sysplex cluster. This provides high availability, high capacity, and full workload balancing. Work does not stop if an individual application instance or queue manager should fail or be recycled, as other instances of the same application accessing the same shared queues continue to do the work. Throughput is no longer constrained by the capability of a single queue manager, as multiple queue managers can access the same shared queues. Automatic pull workload balancing is achieved as the least constrained application instance will process the most messages.

Intelligent Resource Director

IRD consists of three functions

- *LPAR CPU Management*
- *Channel Subsystem Priority Queuing*
- *Dynamic Channel Path Management*

An LPAR Cluster is the set of z/OS partitions on the same IBM zSeries @server z900 running in the same Parallel Sysplex cluster.

Without IRD, one must:

1.Perform manual reconfiguration changes to adjust partition weights and I/O configuration. This is difficult if the workload changes are frequent and unpredictable

2.Buy enough channel and CPU resources so all the workloads are happy. This is expensive and is limited by the maximum number of channels supported.

IRD automates and simplifies the process while doing so in an efficient and dynamic manner.

In many installations, there has been a proliferation of logical partitions. This is driven by many reasons, such as application isolation, or footprint consolidation. Intelligent Resource Director, or IRD, automates **dynamically** many of the manual planning and tuning efforts that go into maintaining these logical partitions. LPAR CPU Management fine tunes the LPAR weights and the number of "logical CPs." Channel Subsystem Priority Queuing prevents less important, but I/O intensive, workload from crowding out I/O access from a more important workload on this or another partition. Dynamic Channel Path Management, or DCM, enables the system to respond to channel requirements as they occur by moving additional channels to DASD logical control units when they are needed in response to changing demand for I/O channel bandwidth. This is now done dynamically. All this combines to:

- Give more efficient use of hardware resource
- Reduce channel requirements
- Simplify I/O configuration planning and definition
- Improve overall system performance
- Enhance RAS

The scope of IRD management is within the z900 processor. The LPAR CPU Management requires Workload Manager Goal mode. Although Dynamic Channel Path Management and Channel Subsystem Priority Queuing aspects will work with the Workload Manager in Compatibility mode, it provides its greatest value in Goal mode. A WLM Coupling Facility structure is required for LPAR CPU Management and DCM. z/OS 1.1 is required for IRD support, although z/OS 1.2 can dynamically manage CPU resources of non-z/OS partitions, in particular Linux® images.

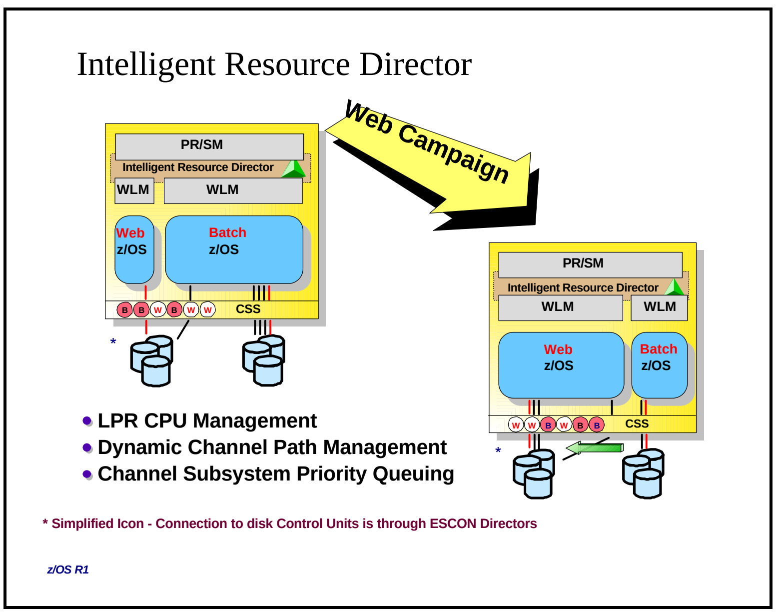

Additional Resource Sharing Exploitation Benefits

OS/390 Security Server (RACF), High-speed Access to Security Profiles

RACF® can use the Coupling Facility to improve performance by using a cache structure to keep frequently used information located in the RACF database. The cache structure is also used by RACF to perform very efficient, high-speed cross invalidates of changed pages within the RACF database. Doing so preserves the working set that RACF maintains in common storage.

RACF is another example of an exploiter that provides improved performance and scalability. RACF is capable of using the CF to read and register interest as the RACF database is referenced. When the CF is not used, updates to the RACF database will result in discarding the entire database cache working set that RACF has built in common storage within each system. If an installation enables RACF to use the CF, RACF can now selectively invalidate only changed entries in the database cache working set(s) thus improving efficiency. Further, RACF will locally cache the results of certain command operations. When administrative changes occur, such commands need to be executed on each individual system.

RACF databases can also be consolidated to obtain one security environment for the enterprise. Many installations today have a RACF database per each image or JES (Job Entry Subsystem) complex. Consolidating to one database (1) provides a consistent security environment to all users on all systems, (2) eliminates the need to synchronize RACF databases through vendor or local products and (3) enables the security administrator to define security profiles and classes in one place instead of in every RACF environment.

JES2, Checkpoint

Beginning with JES2 V5, JES2 supports placing its checkpoint in the Coupling Facility. When this option is selected a Coupling Facility list structure is used for the primary checkpoint data set. The alternate checkpoint data set can reside on DASD. The benefits of having the JES2 checkpoint in the Coupling Facility include equitable access to the checkpoint lock across all members within the MAS complex. In addition, members of the MAS are now capable of identifying who owns the lock in the event of a JES failure.

Cross System BatchPipes

BatchPipes® uses a list structure in a Coupling Facility to enable cross-system Pipes. Cross-system BatchPipes is also known as Pipeplex. BatchPipes allows for the parallelization of input/output processing across systems in a Parallel Sysplex cluster to minimize total elapsed time for job execution.

Conclusion

S/390 Resource Sharing has many advantages:

- Systems Management Simplification, Single System Image and Single Point of Control
- Configuration Management Simplification and Cost control
- Performance Optimization
- Elimination of redundant hardware and software components

From a systems management point of view, Resource Sharing facilitates accessing multi-system resources from any system in the sysplex. It eliminates the necessity to keep multiple copies of files across different systems. It reduces the costs associated with copy management as well as the redundant disk storage that results from multiple copies.

The ability to consolidate SMS, HSM, and Tape complexes into one environment also dramatically improves the systems management of these environments. With such multiple complexes, system programmers today have to maintain multiple policies, system libraries and definitions. A change to one complex may require a change to the other complexes; or, worse yet, a change isn't required of one or more of the others, so different environments must be maintained. Each environment needs to be managed and operated individually. Different procedures and tools may exist for each. These factors are why the problem becomes worse as the number of systems/complexes increases.

The consolidation and standardization of system data sets, such as PARMLIB, PROCLIB, Master Catalog and SYSRES volumes, can significantly ease the management of zSeries and S/390 environments. Today, multiple copies of these data sets exist. Each must be maintained, and they are likely to be very disparate from one another. System programmers must keep track of the differences, the reasons for those differences, and updates of the appropriate ones. When an update is introduced across systems, multiple copies of the same data sets must be updated. Much of this process can be simplified by consolidating onto one or two system data sets (e.g., PARMLIB) per Parallel Sysplex cluster. System programmers can then maintain these libraries much more easily. With the use of symbolics and cloning, one simple change to one member of a data set can be sufficient. This consolidation also leads to standardization of these system resources across the Parallel Sysplex cluster, which further makes it easier to implement changes.

S/390 Resource Sharing has substantial benefits for all OS/390 customers running multiple instances of the operating system, in both single server and multi-server configurations, and to top it all off, the benefits are shipped as part of the base OS/390 product set.

Is this you?

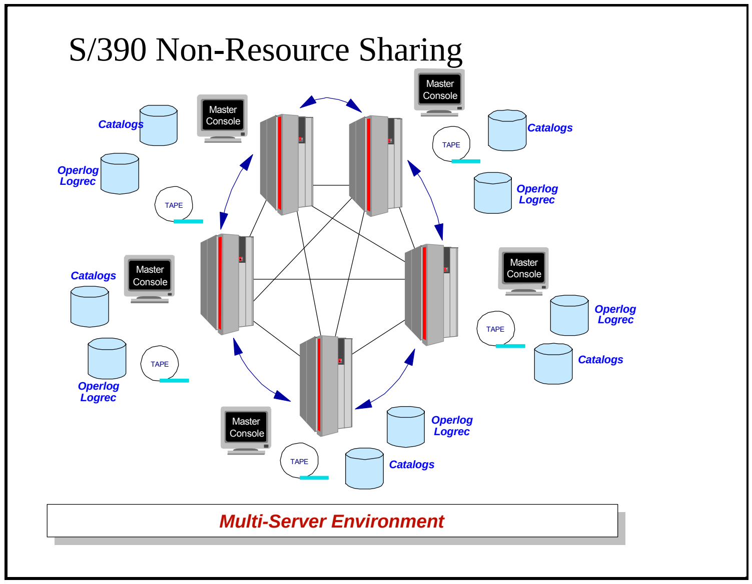

Then zSeries Resource Sharing is for you:

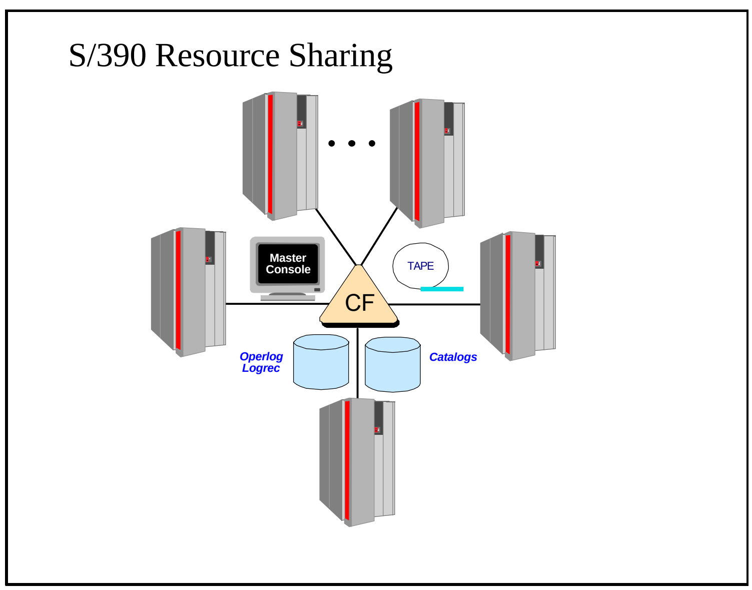

Is this you?

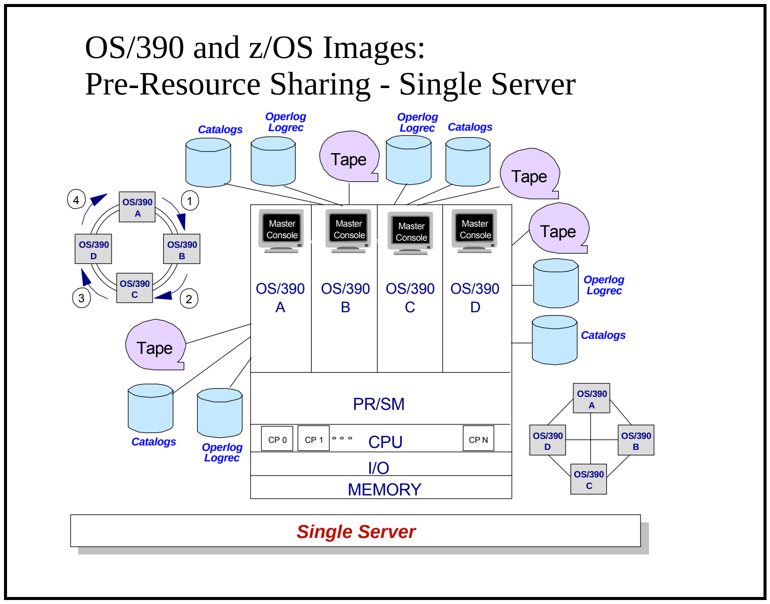

Then zSeries Resource Sharing is for you:

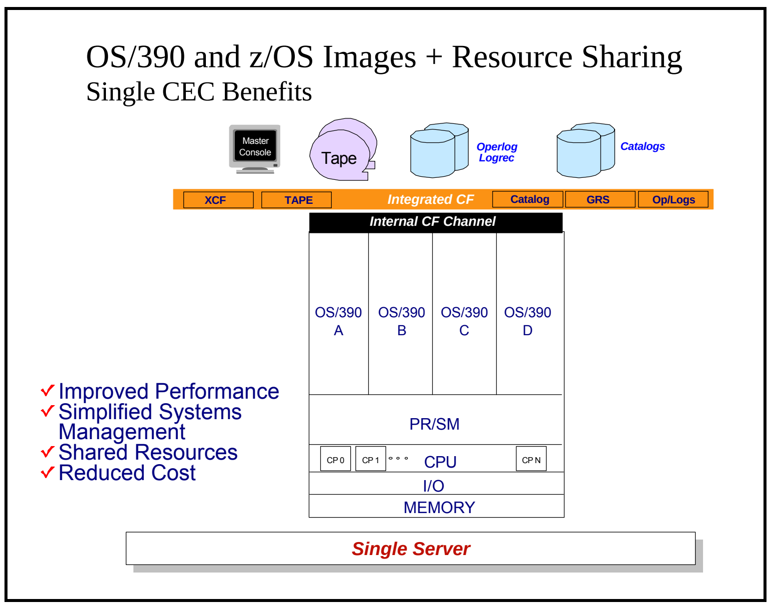

IBM Corporation
Marketing Communications, Server Group
Route 100
Somers, NY 10589

Printed in the United States of America, 10/01

This publication was produced in the United States.

IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

You can find additional information via IBM's World Wide Web server at http://www.ibm.com.

IBM Hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

GF22-5115-02

ZSW03333-USEN-00