

TOPICS

A z/OS Newsletter, Issue 4

February 2001

Talking with Tom Rosamilia z/OS: The path to the new world



Now that Y2K is behind us and we move forward into the next millennium, the IT industry is being challenged with the constantly changing and increasing demands of the e-business environment.

We are rapidly moving beyond the first generation of e-business, which was gaining accessibility to the Internet to exploit the Internet to create new business opportunities. This second generation of e-business requires an infrastructure that can handle the explosion of more users, more devices, more transactions anytime, anywhere.

On October 3, 2000, IBM launched the IBM @server™ product line, which includes the zSeries, iSeries, pSeries and xSeries platforms that can be combined in a heterogeneous mix of systems seamlessly integrated to meet the needs of businesses in the new world of e-commerce. The first zSeries server, the IBM @server zSeries 900, is an enterprise class e-business server optimized for the integration of business data and transactions. z/OS is the z900's new 64-bit operating system which exploits the z/Architecture, delivers high speed connectivity to the network and the storage systems, and provides scalability, reliability and security. This reliable, mission-critical data transaction server is capable of dynamically shifting workloads across the server and the network based on customer priorities.

Based on OS/390®, z/OS delivers innovative technology that is tightly integrated with the z900 hardware and licensed internal code. The core of this support is Intelligent Resource Director (IRD) formally known as LPAR clusters. IRD is Workload Manager (WLM) Goal mode resource management, Parallel Sysplex resource sharing and PR/SM all working cohesively together to deliver maximum value. There are three components of IRD - Dynamic CPU Management, Dynamic Channel Path Management and Channel Subsystem Priority Management – which enable the effective use of CPU and I/O resources across LPARs. The dynamic combination of these resources are selected based on your service level goals.

All of the new technology exploited by zSeries operating systems allows for better support of mixed workloads on the platform. z/OS provides an environment for classic mainframe workloads, and for Java and UNIX® workloads that need the highest quality of service. LINUX™ for zSeries offers a wealth of applications and development tools such that your business has an expanded set of application choices. Linux can be deployed in zSeries LPARs or with z/VM. With z/VM, it is possible to run hundreds of Linux systems as virtual servers in one LPAR, eliminating the need for hundreds of real servers. The combination of running Linux and

z/OS side by side on the same machine provides unique strengths that are not possible on other platforms. Virtual servers can dynamically expand and contract using IRD and applications can communicate through memory using HiperSockets. For example, combining a Web Application Server running in a Linux partition with legacy applications and mission-critical data in z/OS will improve performance, increase availability, reduce complexity and protect your investments.

Central to the zSeries offering is the IBM License Manager which enables a new software pricing strategy designed to support e-business. This workload pricing allows your business to define a maximum capacity level and pay only for the software capacity needed to meet the demands of current workloads. This pricing structure is designed to leverage the new technology so that the systems can allow peaks in transaction volumes, can provide the level of granularity needed to grow one workload at a time and can quickly adjust and enable additional capacity on demand. The new software licensing capabilities are receiving support from other industry software providers.

These are exciting times for the IT industry and IBM® and the zSeries team is committed to building support for an e-business infrastructure that will deliver world-class value, leverage current investments in the platform and define a path into the new world of e-commerce. •

Tom Rosamilia is Vice President, zSeries Software Development



GA22-7501-00



Contents

Talking with Tom Rosamilia	
z/OS: The path to the new world	1
In this issue	3
Letters to the editor	3
Introducing the Intelligent Resource Director	4
LPAR CPU Management —	
Making a good thing even better	6
Dynamic Channel Path Management	8
Channel Subsystem Priority Queueing	11
Linux meets the enterprise servers	15
Hot topic: cool product -	
an introduction to Linux on the enterprise servers	16
Getting personal with the penguin	18
64-bit real operation in OS/390 Release 10	22
Self-managing your software licenses and costs	24
A piece of the action	26
License Manager Standards	27
Exploiting z/Architecture with EXCP I/O	29
Ramp up to IBM License Manager	32
ISVs and the IBM License Manager	33
z/OS makes your system programmers more effective!	36
Parallel Sysplex® in z/OS made easier!	38
What happened to ServerPac for z/OS?	42
Parallel Sysplex update	45
IBM's strategy for software service	48
New courses!	49
Our contributors	50

TOPICS

A z/OS and OS/390 Newsletter Issue 4 February 2001

EXECUTIVE EDITOR BOB WARD
 MANAGING EDITOR DICK WAGENAAR
 CREATIVE DIRECTOR ... KERRY COLLIA

© International Business Machines
 Corporation, 1999, 2001

Produced in the United States of America
 1-2001
 All Rights Reserved

The z/OS and OS/390 Hot Topics Newsletter
 is published twice yearly as a service to
 customers, providing articles of interest to the
 z/OS and OS/390 community.

The information in this Newsletter is presented
 as is and is based on the best knowledge of
 the authors. No warranty is provided (neither
 expressed nor implied).

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

© CICS, CICSplex, ESCON, IBM, MQSeries,
 OS/390, Parallel Sysplex, RACF, S/390, Tivoli,
 VTAM are registered trademarks of International
 Business Machines Corporation in the United
 States, other countries, or both.

™ @server, DB2Connect, GDPS, FICON, MVS,
 MVS/XA are trademarks of International
 Business Machines Corporation.

z/OS, zSeries, z/VM, and z/Architecture are
 trademarks of IBM.

Java, and all Java-based trademarks are
 trademarks of Sun Microsystems, Inc., in the
 United States, or other countries, or both.

LINUX is a trademark of Linus Torvalds.

Unicode Standard and Unicode Consortium are
 trademarks of The Unicode Consortium.

Microsoft, Windows, Windows NT are
 trademarks of Microsoft Corporation in the
 United States, other countries, or both.

UNIX is a registered trademark in the United
 States and other countries licensed exclusively
 through The Open Group.

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

In this issue ...

Welcome to our fourth issue! As you've probably noticed by now, we've got two front covers this time. No, it's not some avant-garde concept in twenty-first century newsletter design. It's our way of showing that there's a new generation of enterprise-level operating system in your future, a rebranded OS/390, called z/OS. Be sure to check out Tom Rosamilia's article to read all about it. Tom is the Vice President of zSeries Software Development, and his article will help you understand why we rebranded and what value it brings to you. Additional articles in the z/OS side of Hot Topics explain in more detail the enhanced function it delivers. The z/OS side also includes a series of answers to questions you've asked about Linux, and how it plays in the OS/390 and z/OS worlds.



On the flipside ... we have an assortment of articles of general interest, including two articles on TCP/IP, and an article that explains that latest about ordering products.

Also, be sure to check out our new 'Letters to the Editor' column. We want to make this a regular feature of the newsletter, but, of course, that depends on you. Don't hesitate to let us know what you

think, or to ask us a question, or a bunch of questions. If we think others would be interested, we'll publish your letter and our response. The address is newsletr@us.ibm.com.

Our Supplement this issue, edited by our guest editor, Susan Chalenski, offers a wealth of information about how we package and deliver the information that helps you manage and use OS/390 and z/OS.

As always, we hope we've again provided information that you can use.

The Editors,
newsletr@us.ibm.com

All issues of HOT TOPICS are available online at the OS/390 and z/OS homepages.

Letters to the editor

To the editor at
HOT TOPICS Newsletter:

Currently we deliver all the doc (usually 5-25 books per product) on a CD-ROM as BOTH BookManager® .BOO softcopy with a BookShelf and BookIndex, AND as Adobe Acrobat .PDF. We have found that the Acrobat strength lies in printing while the BookManager lies in searching and research. This gives us (and our customers) the best of both worlds.

We have noted the drastic increase of PDF usage within IBM and are concerned with the future of BookManager. There are strong opinions (primarily from those with a PC only background) that we drop BookManager as "no one

uses it", and equally strong opinions from the development and maintenance groups (both within the company and from customers) that PDFs are wasted space on the CD-ROM. So, what is the future of documentation, at least as far as BookManager is concerned?

Jim Bartlett,
Senior Systems Architect at CSC
(Computer Sciences Corporation)

Dear Jim:

As you know, we actually followed up on this letter earlier with a conference call. But here, for those people who may also be interested, is the gist of what we talked about. You are obviously

very familiar with both of these softcopy formats (BookManager and PDF), and summarized the strengths of each one very well. Because BookManager and PDF are complementary formats, with some people having a strong preference for one or the other, we are continuing to produce both softcopy formats for our z/OS operating system books. IBM is continuing to make some enhancements to the BookManager set of products, and we intend to continue to deliver both BookManager and PDF formats for our z/OS operating system books.

John Sefcik,
Information Strategist,
IBM Poughkeepsie User
Interface Team

Introducing the Intelligent Resource Director

STEVE HAMILTON AND MADELINE NICK

The “Intelligent Resource Director” — does that sound like a fancy job title, somebody who’d get one of the prime parking spaces, maybe right next to the Chief Information Officer? The Intelligent Resource Director is actually an important new function, introduced as part of the IBM zSeries and z/OS. Using the strengths of logical partitions, Partition Resource/System Manager (PR/SM)[®], Workload Manager (WLM), I/O Supervisor (IOS), Channel Subsystem, and Parallel Sysplex, the Intelligent Resource Director gives your installation more power to dynamically move resources to your most important work.



Kerry Collia

The concept of logically partitioning your work is not new, of course. For years, S/390[®] architecture has allowed you to divide your Central Processing Complex into separate logical partitions, so that different types of work can run in their own unique environments. Each logical partition is assigned its own “LPAR weight,” meaning the percentage of overall processing power that’s guaranteed to the work in that partition. If the workload shifts to the extent that more processing power is needed in a particular partition, the operator would have to change it manually via (PR/SM). If the demand is unpredictable and irregular (as in a web server environment), you’ll need that

operator to monitor the system at all times, day and night, ready to change the weights.

In addition, each logical partition also has its own statically-defined set of channel paths to I/O control units. Again, in the event of a significant shift in workload, these channel paths could only be changed by an operator, via the HCD definitions.

Once an I/O request made it to the channel subsystem, it would be serviced on a first-in first-out (FIFO) basis. This could at times cause your highest priority work to be delayed because a lower priority batch request came into the queue just before it.

The Intelligent Resource Director extends the concept of goal-oriented resource management by allowing you to group logical partitions that are resident on the same physical server, and in the same sysplex, into an “LPAR cluster.” This gives the z/OS Workload Manager the ability to manage resources, both processor and I/O, not just in one single image but across the entire cluster of logical partitions.

The Intelligent Resource Director is made up of three parts: LPAR CPU Management, Dynamic Channel Path Management, and Channel Subsystem Priority Queuing.

LPAR CPU Management

LPAR CPU Management lets Workload Manager distribute processor resource across an LPAR cluster by dynamically adjusting the LPAR weights in response to changes in the workload requirements. When important work is not

meeting its goals, WLM will raise the weight of the partition where that work is running, thereby giving it more processing power.

As part of LPAR CPU Management, WLM will also optimize the number of online logical CPUs dedicated to each partition. As the LPAR weights change, the number of online logical CPUs are also changed to maintain the closest match between logical CPU speed and physical CPU speed. LPAR CPU Management requires a CF level 9 coupling facility structure and z/OS images running in goal mode on a zSeries machine in z/Architecture mode.

For more, see Gail Whistance’s article on LPAR CPU Management!

Dynamic Channel Path Management

Dynamic Channel Path Management lets Workload Manager dynamically move channel paths through the ESCON[®] Director from one I/O control unit to another, in response to changes in the workload requirements. By defining a number of channel paths as “managed,” they become eligible for this dynamic assignment. By moving more bandwidth to the important work that needs it, your DASD I/O resources are used much more efficiently. This may decrease the number of channel paths you need in the first place, and could improve availability — in the event of a hardware failure, another channel could be dynamically moved over to handle the work requests.

For more, see William Rooney’s article on Dynamic Channel Path Management!

Dynamic Channel Path Management requires z/OS and a zSeries machine in z/Architecture mode, which can be running in either Basic or LPAR mode. It supports DASD control units which are non synchronous and are connected via an ESCON Director. (Initially, it supports ESCON channels only.) It's optimized when WLM is running in goal mode.

Channel Subsystem Priority Queuing

Channel Subsystem Priority Queuing is an extension of I/O priority queuing, a concept that has been evolving in MVS™ over the past few years. In an LPAR cluster, if important work is missing its goals due to I/O contention on channels shared with other work, it will be given a higher channel subsystem I/O priority than the less important work. This function goes hand in hand with the Dynamic Channel Path Management described above — as additional channels are moved to the partition running the important work, Channel Subsystem Priority Queuing ensures that the important work that *really* needs it receives the additional I/O resource, not necessarily the other work that just so happens to be running in the same LPAR cluster.

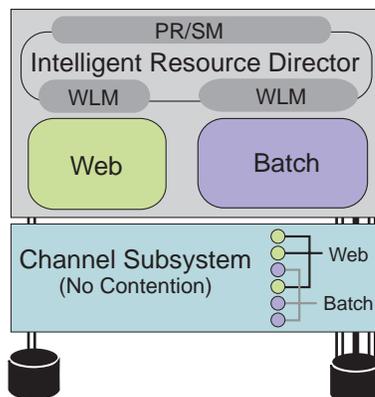
For more, see Jürgen Märgner and William Rooney's article on Channel Subsystem Priority Queuing!

Channel Subsystem Priority Queuing requires a zSeries machine in z/Architecture mode. It initially supports ESCON and OEMI channels only. It's optimized when WLM is running in goal mode.

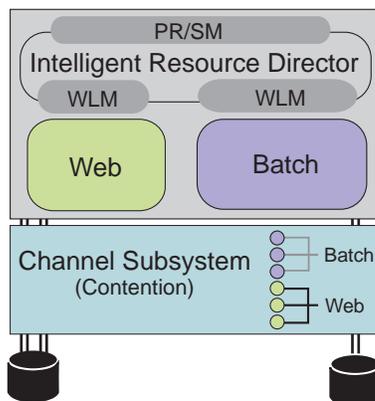
Here's an example!

To illustrate what the Intelligent Resource Director can do for a

rapidly changing, mixed workload environment, consider this scenario:



In this first figure, you're serving your web site with one logical partition, and running batch work in another. Suddenly, your web site is overrun with hits! Because you've defined the web work to have a higher importance than the batch work, the Intelligent Resource Director will dynamically adjust the weight of the logical partition serving the web work. Also, a number of channel paths will be dynamically reassigned to the DASD used by the web work. This additional I/O bandwidth will ensure that the web work does not miss its goals. Finally, because there is now contention in the channel subsystem, the Channel Subsystem I/O priority will ensure that the web I/O requests are given priority. With no human intervention whatsoever, the Intelligent Resource Director responds to the sudden shift in workload, and your most important work gets the resources it needs!



So what will the Intelligent Resource Director do for me?

If you're migrating into the e-business environment with Websphere application servers, you should consider z/OS and the benefits of consolidating those applications onto one machine. With each in its own separate logical partition, the Intelligent Resource Director will make sure that each gets the CPU and I/O resources it needs, while you reap the benefits of workload isolation

A Redbook called z/OS Intelligent Resource Director (SG24-5952) will be available first quarter of 2001. The book consists of three parts - one part for each of the constituent parts of IRD - LPAR CPU Management, Dynamic Channel Path Management, and Channel Subsystem Priority Queuing. Each part provides an introduction to that component, and helps the reader determine the benefits that IRD will provide in their environment. After the introduction, there are sections that discuss how the component works, how to plan for it, how to implement, and operational performance, and tuning information. The book can be read sequentially, or each part can be read independently of the rest of the book.

and the ability to rapidly deploy new applications with minimal impact to existing workloads.

If you're already enjoying the benefits of S/390 Resource Sharing, the Intelligent Resource Director is a natural addition to your environment. It extends all of the strengths of Resource Sharing and lets you more easily and efficiently run many diverse workloads on one machine.

Finally, Parallel Sysplex data sharing customers will find significant value when combining the power of the Intelligent Resource Director on one server with data sharing across multiple servers in a sysplex. With the combined ability to transparently route transactions across the multiple servers in your datasharing group while dynamically moving the required resources across non-datasharing partitions within one server machine, all your business goals are met. What more could you ask for?

Where do I go for more information?

For more information on the Intelligence Resource Director, see *z/OS Planning: Workload Management, SA22-7602* and the IBM Redbook, *z/OS Intelligent Resource Director, SG24-5952*.

You can go to the Parallel Sysplex website: <http://www.ibm.com/s390/ps> for additional information on IRD, Parallel Sysplex and the new z/OS msys for Setup support to easily migrate into an IRD environment. •

LPAR CPU Management — Making a good thing even better

GAIL WHISTANCE

LPAR CPU management is an exciting new element of the z/OS's Intelligent Resource Director. CPU management has for years been the centerpiece of the Workload Manager's dynamic resource management. It's well known that on a single z/OS image, you can consolidate workloads for efficient use of resources and at the same time differentiate those workloads through an WLM policy according to their importance to the business. As the workload mix changes, WLM quickly responds and ensures the "most deserving" work gets the resources it needs.

Now with IRD, dynamic CPU management is extended to include the automatic management of CPU resource across z/OS images running as logical partitions (LPARs) on the same CPC. This was formerly a manual task of adjusting LPAR weights and logical CPU assignments. Now the power of WLM, Parallel Sysplex, and PR/SM combine to allow you to make the most efficient use of

your CPC's CPU power while honoring your business goals — all automatically.

Why Logical Partitions?

Even with z/OS's strength as a workload consolidation platform, there are still many reasons why customers use the PR/SM facility to separate workloads onto different LPARs on a physical machine rather than running them all on the same z/OS image. These include:

- Workload isolation for availability — It is less likely for a healthy workload to suffer an outage when another workload fails.
- Server consolidation — There is a trend to combine workloads from several smaller servers onto one large CPC. A convenient way to do this is to keep those workloads on separate z/OS images but run them in LPARs on the same CPC.

- Different quality of service — Different responsiveness requirements for production, test, and application development workloads motivate many customers to keep these types of work separated.

Logical partitioning will continue to grow in value as the size of a single CPC grows and as more customers consolidate servers. Throughout the server industry, most major platforms are pursuing some form of partitioning.

So what's the problem?

One of the resources that can be allocated to a logical partition is processing weight (LPAR weight). Weight determines the portion of the CEC's shared CPU resource allocated to an LPAR when there is competition for CPU among two or more partitions. It is the job of the system programmer to set LPAR weights so that the workloads running in the logical partitions get the CPU they need. If the workload mix changes, then the weights have

to be reevaluated and manually changed. To properly adjust LPAR weights requires a human to constantly monitor workload performance and to understand the business goals and tradeoffs for these workloads. Not an easy task! An alternative is to configure enough CPU capacity to allow for peak workload demands so that no critical work is never delayed. But when the workload is not at peak demand, the excess capacity goes unused, and as one customer put it, “A CPU cycle unused is lost forever.” It is an expensive alternative to overconfigure.

So we have WLM managing CPU resources across workloads within a single image, and the PR/SM facility managing images but not workloads. What is needed is workload management across images. This is where the Intelligent Resource Director comes into the picture.

Let LPAR CPU Management do it for you!

IRD’s LPAR CPU management offers a way to get the most out of your CPU resource by dynamically adjusting LPAR weights based on the WLM policy. WLM moves CPU resource to the partition with the most “deserving” workload and allows the system to adapt to changes in the workload mix. The weight exchange is done among LPARs in a cluster, which is a set of z/OS images belonging to the same sysplex and running on the same CPC. The total weight of the cluster remains constant, so LPARs outside the cluster are unaffected. Weight management kicks in when all LPARs in the cluster are running CPU-constrained and the weights are being enforced by PR/SM. When just one LPAR is using all or most of the CPC

because other LPARs are idling, then weight management will have no effect on resource distribution.

To illustrate the value of LPAR CPU Management in a mixed workload environment, here’s an example:

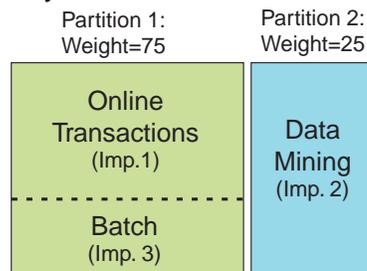
You have three workloads running on one server:

- **Online Transactions**, your most important workload. This runs only during the day shift.
- **Data Mining**, which has a medium importance. This is always running, and will consume as much resource as you give it.
- **Batch**, which is your lowest importance work. Like data mining, it is always running, and will consume as much resource as you give it.

In this example, the server is divided into two partitions:

- **Partition 1** runs both the online transactions and the batch work, as they happen to share the same database.
- **Partition 2** runs the data mining work.

Day Shift



This first figure shows the day shift. As the online transaction workload is the most important, Partition 1 is given a high enough weight to ensure that the online transaction work does not miss its goals due to CPU delay. Within the partition, Workload Manager is making sure that the online transaction work is meeting its

goals before giving any CPU resource to the batch work. (That piece of the puzzle isn’t new, of course — WLM has been managing workloads in this way for years.)

Night Shift



This figure shows the night shift. Everybody has gone home, so there are no more online transactions. If the partition weights had remained the same, then the batch work would be consuming most of the CPU resource. LPAR CPU management automatically adjusts to this change in workload, adjusting the partition weights accordingly. Now the more important data mining work will receive most of the CPU resource.

LPAR CPU Management takes advantage of the combined power of Parallel Sysplex, PR/SM, and WLM. A WLM coupling facility structure is used to maintain cluster performance data, PR/SM provides the mechanism to change weights, and WLM runs the whole show with its goal-based policy.

The following figure shows how weights are adjusted among three LPARs in a laboratory test cluster, where partition A is running the most important work, partition B is running medium importance work and partition C the least important work. When LPAR A’s workload increases at time 00:07 and starts

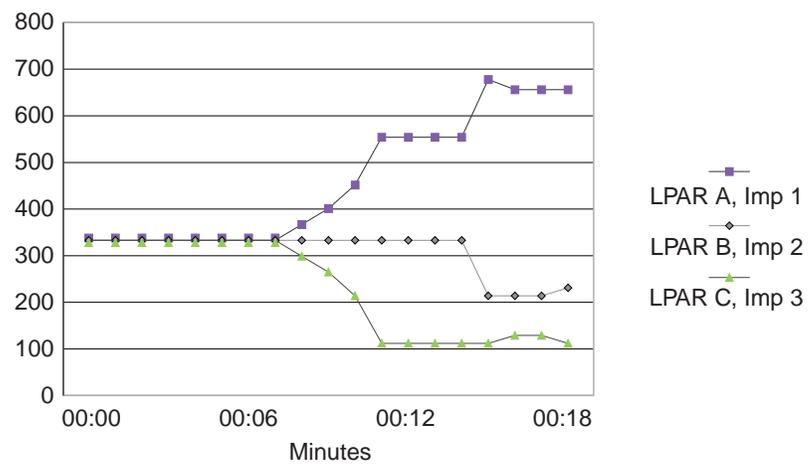
missing its goal, you can see weight start to move to that LPAR with first LPAR C as donor, then LPAR B as donor.

Logical CPU Management

A companion function to LPAR weight management is logical CPU management. This function makes sure that the right number of logical CPUs are online based on the partition's current weight and CPU consumption. The idea is to set the number of online CPUs so the "logical CPU speed" is close to the physical CPU speed. For example, if a partition is consuming 1000 SUs a second and the physical CPU speed is 200 SUs a second, you want 5 logical CPUs online. Optimizing the number of logical CPUs benefits workloads that have large amounts of work done under single tasks, and minimizes LPAR overhead for all workloads.

LPAR Weight Management Experiment

LPAR A's importance 1 workload increased at time 00:07 and started missing its goals. Weight is first taken from LPAR C, which is running importance 3 work, then from LPAR B, which is running importance 2 work.



Summary

IRD's LPAR CPU Management, consisting of both weight management and logical CPU management, takes WLM another step forward in managing the resources of your enterprise. It takes WLM's

ability to run diverse competing workloads in a single system and extends it to multiple LPARs in a cluster. It allows you to take advantage of both logical partitioning and goal-based resource management. •

Dynamic Channel Path Management

WILLIAM J. ROONEY

Introduction

Dynamic Channel Path Management (DCM) is a feature of the new IBM zSeries and z/OS that further extends zSeries' lead in managing multiple heterogeneous workloads with various business priorities to achieve their goals.

It is one of two important new features of the I/O portion of Intelligent Resource Director (IRD). Dynamic Channel Path Management allows z/OS to dynamically change channel path definitions to ESCON Director attached DASD control units in response to changing workloads, and in doing so, move channel resources to the control units where they are

required. When combined with Workload Manager (WLM) running in goal mode, DCM moves the channel resources to control units that are being used by business critical workloads to help them meet their goals.

When further combined with Control Unit Priority Queueing in the Enterprise Storage Server (ESS or Shark), Channel Subsystem Priority Queueing (see accompanying article) and Parallel Access Volume (PAV) Alias Tuning, the z/OS system becomes even more self tuning and self defining to meet workload goals.

Background

Today, the I/O configuration definition process is complex, and requires significant skills. The process normally involves determining how many channels are required by a control unit, and how many other control units, if any, can share that same set of channels.

There are availability considerations. For example, even if only a single channel is ever required by a control unit, two or more are normally defined to it in case of a failure somewhere along the path.

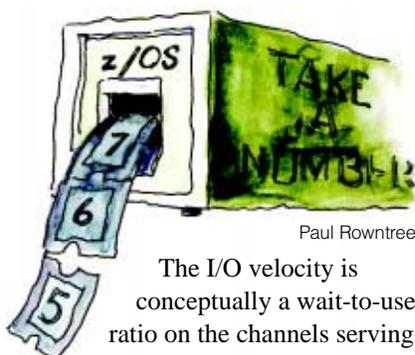
The configuration process tends to be iterative. A configuration is defined, measured and analyzed.

This is repeated until the configuration has the desired response time and availability characteristics.

Even if the configuration is defined perfectly the first time, the characteristics of the workloads on systems change. New workloads are added, some workloads grow more quickly than others, and new I/O devices are added to the configuration. Some workloads are less predictive than others. Some may vary slightly from week to week, and others may have sudden and unpredictable spikes in demand. These factors combine to create a situation where an I/O configuration that allowed you to meet your response time goals last week is inadequate to meet your goals this week. To make matters worse, you may actually have sufficient I/O resources, just not where you currently need them.

Subsystem I/O velocity

Once every interval, currently defined to be ten seconds, DCM goes through a process known as Data Gathering. During this process DCM collects several measurements on DASD subsystems¹, and uses that information to calculate a new metric called the subsystem I/O velocity (IOV).



The I/O velocity is conceptually a wait-to-use ratio on the channels serving a particular subsystem. It tells us how long we must wait for a channel to a subsystem compared to how long we actually use those channels.

For example, assume that we have a subsystem which has accumulated eight seconds of connect time in the last ten seconds, and one second of pending time. Also assume that switch busy, control unit busy and device busy are all zero. Then we would have an I/O velocity of $\frac{8}{8+1} = 0.89$. If the pending time grew to three seconds and the connect time stayed the same, the I/O velocity would drop to 0.72. So the closer the I/O velocity gets to 1.00, the less I/O requests are waiting for channels to a subsystem. The closer it gets to 0.00, the more constrained the subsystem is for channel bandwidth.

DCM calculates this I/O velocity for all DASD subsystems in the system. It then calculates a Default Target I/O Velocity. DCM then looks at the I/O velocity actually being achieved for each of the managed subsystems, and attempts to get all of the subsystems close to the default target.

A list of all subsystems whose I/O velocity falls below the target range, is then passed to another phase of DCM called Imbalance Correction. Imbalance correction processes subsystems one by one, starting with the one furthest from target, and looks for changes that would help the subsystem reach its target by searching tables which represent the physical I/O topology built by DCM during system initialization and maintained since then. There are two types of changes that DCM looks for; paths that can be added to this subsystem and paths that can be removed from subsystems that share the path with the subsystem that requires more bandwidth.

Each of the possible changes may also impact the availability characteristics of the I/O configuration. If a path needs to be removed from a subsystem, DCM

favors choices which do not leave all the remaining paths with a single point of failure. When adding a path, DCM tries not to choose one that will have a single point of failure in common with all the paths already defined for the subsystem.

A measure of configuration complexity is calculated for each of the possible changes. This measure of complexity is referred to an I/O entropy index. For example, consider a configuration where Subsystem 1 is connected to CHPIDs 23 and 24, Subsystem 2 is connected to CHPID 24 and Subsystem 3 is connected to CHPIDs 25 and 26. Assume that DCM has determined that subsystem 2 is below target and that it could add bandwidth either by adding CHPID 23 or 25. If it uses CHPID 23, then the number of interconnected subsystems and channels will not change, but if it adds CHPID 25, then Subsystems 1, 2 and 3 will be interconnected with CHPIDs 23, 24, 25 and 26. This is more complicated, and all other things being equal, DCM would prefer the less complicated configuration.

Next, DCM projects the I/O velocity achieved by each of the possible changes. DCM also projects I/O velocities for any subsystems that are affected by the change. (An affected subsystem is any subsystem which shares a channel, either directly or indirectly, with the subsystem being helped.)

Finally, with each of the possible decisions evaluated, DCM considers all of the attributes and selects one to implement.

Workload Manager

Dynamic Channel Path Management operates in two modes, balance mode and goal mode. In balance mode, DCM operates as

described above. WLM has no influence on DCM's decisions.

In balance mode DCM does not know which subsystems are important to critical work, so as work builds on a system, and the channel resources become over-committed, the default target I/O velocity will begin to drop. Under these conditions, the only option DCM has is to try keep each subsystem's I/O Velocity within tolerance of an ever decreasing target.

If WLM is operating in goal mode on all systems in the Logical Partition (LP) Cluster², DCM is considered to be operating in goal mode. In goal mode, WLM uses its knowledge of what the service class goal are, what subsystems are being used by the service classes and whether the IOV of those subsystems are low. If the channel resources become over-committed and are impacting the ability of service classes to achieve their goals, WLM then has the ability to set an explicit IOV target for subsystems. DCM will then attempt to meet the explicit goals for these subsystems, while continuing to drive the other managed subsystems toward the default target.

Activating Dynamic Channel Path Management

There are some new terms introduced with DCM. They are static and managed channels, and managed control units.

Static channels are channels that are defined to control units in the traditional way. This is typically done by adding the CHPID number of the channel on the "Select Processor / Control Unit" panel of HCD. Managed channels are channels that are owned by DCM. These channels never appear on HCD's control unit definition panel.

Managed control units are control units that DCM may dynamically add channels to.

In order to activate DCM, you must first define a pool of managed channels and a set of managed control units that may receive them. To define managed channels, under the "Add (or Change) Channel Path Definition" panel you need to indicate that the channel is managed, and if the definition is to be used for an LPAR mode configuration, you must include an I/O Cluster name. The I/O cluster name is the name of the sysplex that will own and manage this channel.

Once you have defined a set of channels that DCM may manage, you need to tell DCM which control units may receive them, and the maximum number of managed channels that may be added. To do this, under the "Select Processor / Control Unit" panel, enter an asterisk (*) where you would normally enter the "Channel Path ID. Link Address". For each asterisk that is entered, DCM will add that many managed channels if it needs to. So if you enter two asterisks, DCM will add as many as two managed channel, but no more than that. You may add as many asterisks to the control unit definition as you wish, as long as the total number of statically defined channels and asterisks does not exceed eight.

Logical Partitions

DCM will share (or EMIF³) channels across logical partitions which are all members of the same sysplex. Since DCM removes channels from control units, it must assure that it is not removing the last channel from a control unit in another partition. To do this, it first communicates with other members of the sysplex on the same Central Processing Complex (CPC), and request that they vary the path

offline. When all systems have responded that they have successfully done so, the channel is removed from the control unit. Since partitions which are not members of the sysplex cannot be notified of the change via XCF, DCM does not allow channels to be shared outside the sysplex.

This restriction is enforced by LPAR. When a partition is IPLed, all managed channels are configured offline to that partition. During IPL, z/OS will identify the sysplex it is joining to LPAR, and will be returned a list of channels which had the same name specified in the "I/O Cluster" name under HCD, which DCM will then configure online.

Channels which have the "managed" attribute cannot be configured online, even via operator command, to a partition that has not identified itself as belonging to the sysplex.

Frequency of change

Dynamic Channel Path Management will consider changes to control unit definitions once every ten seconds. If there are multiple partitions which are members of the same sysplex, each partition will consider changes every ten seconds. However, if one partition finds that another partition is already implementing changes, it will skip this processing for this interval.

During this interval, DCM will attempt to help subsystems which are below their target velocity. Subsystems which have had channels added or removed from them are not eligible to have additional changes for another two intervals. This means that the most frequent that changes will occur to a single subsystem are once every 30 seconds.

The I/O velocity that is used to determine if a subsystem is below target is actually a moving average of up to eleven intervals. This is done to dampen the subsystems I/O velocity and prevent frequent and unnecessary changes.

In addition to the subsystem that is being helped, any subsystem that is affected by the change will also be ineligible for changes for two intervals. An affected subsystem is any subsystem which directly or indirectly shares a channel with the subsystem being helped.

When to use Dynamic Channel Path Management

Dynamic Channel Path Management works best when there are variations over time between the channel requirements of different DASD subsystems. If the nature of the workload is such that most

subsystem have their peak channel requirements at the same time, DCM will be of little help, since it can only reassign existing channels.

Consider an environment where a set of subsystems normally needs four channels, but occasionally needs as many as six. Consider also that it is rare for all of the subsystems to require more than four channels at the same time. In this case the user may define four static channels and the additional two channels as managed. By doing this, the total channel requirements of the system may be reduced, since DCM can move the required channels to the subsystems when they are required. In addition, the user may decide to allow an additional four managed channels, giving the subsystem even greater access to channel bandwidth in the case of unforeseen peaks.

Availability

DCM also enhances availability. In cases where hardware failures cause channel paths to be lost, the effect can be that the responsiveness of the workload can drop so low that it is no longer achieving its service class goals. When DCM detects that these subsystems are no longer achieving their IOV targets, it will respond by attempting to add additional channel paths. This in turn may get the workloads back to where they are again achieving their goals. •

1. A subsystem is equivalent to a Logical Control Unit (LCU). DCM works on a subsystem basis, however the more familiar term "control unit" is used throughout the article.
2. A Logical Partition Cluster is the set of logical partitions on one Central Processing Complex (CPC) which are all members of the same sysplex.

Channel Subsystem Priority Queueing

JÜRGEN MÄRGNER AND WILLIAM J. ROONEY

Introduction

Channel Subsystem Priority Queueing (CSSPQ) is a feature of the new IBM® zSeries and z/OS that further extends zSeries' lead in managing multiple heterogeneous workloads with various business priorities to achieve their goals.

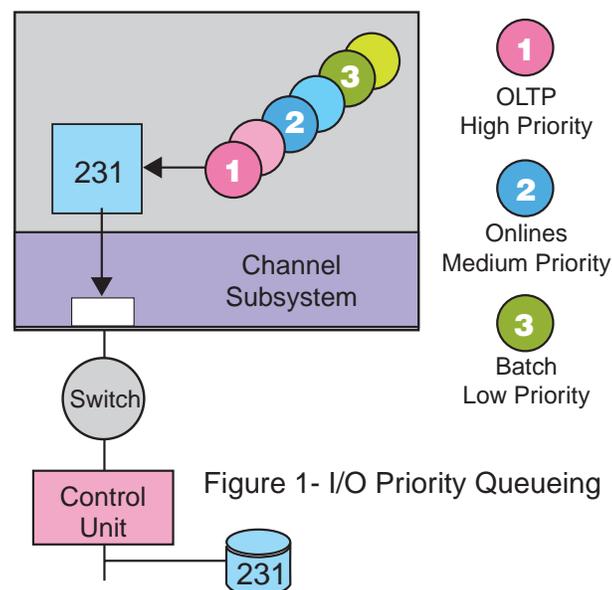
It is one of two important new features of the I/O portion of Intelligent Resource Director (IRD).

Channel subsystem priority queueing allows z/OS to set priorities on I/O requests, which are then used by the z900 to schedule the work to resources within the channel subsystem. In this way customers may identify workloads that are the most

critical, and z/OS will work with the zSeries processor to allow the critical work greater access to channel subsystem resources.

How is it different from today's I/O Priority Queueing

Prioritizing I/O requests is not new for z/OS. Back when the system was known as MVS, I/O requests were queued within MVS on a device basis (see Figure 1). Each device had its own queue of I/O requests



waiting to be served. Different types of work within the system could then be assigned different

priorities so that I/O associated with higher priority work would be scheduled ahead of less important work. These priorities were relatively static. A system programmer would define a relationship between different types of work and their priorities and leave them alone.

This solution had its limitations. Although it was better than having all work within the system of the same priority, consider an environment where you have two business critical workloads. You assign the highest priority to the first workload, a slightly lower one for the second workload. All other workloads in the system receive yet a lower one.

Normally this works for you. But one day, the highest priority workload is achieving its goals, but the second workload is not. You may realize that the first workload is actually significantly exceeding its goals, and if you were to set the priorities equal, or even reversed them, both workloads would be able to achieve their goals (see Figure 2).

This is what Workload Manager added to I/O priority queueing in OS/390 Release 3. Instead of setting fixed I/O priorities for workloads, WLM is able to monitor the work toward goals, and if I/O delay is the reason work is not achieving its goals, it can change the work's I/O priority to help it reach its goal. WLM does this while considering the impact of such a change on other work in the system, and then further monitoring the effect after the change.

The next step in the evolution of I/O priority queueing was introduced with the Enterprise Storage Server (ESS or Shark) and is available back to OS/390 Release 3. Although MVS prioritized I/O requests scheduled to a single device, this priority was not honored at all places where queueing occurs. The control unit for the device was such a place. For one, a single control unit normally supports several devices. If all of the I/O requests going to a database on one device have a high priority, and all of the I/O requests going to a set of files on another device have low priority, then the device level priority queueing will not help in this situation. Both the high priority I/O requests and the

low priority I/O requests will be treated equally when sent to the control unit. I/O requests from the lower priority work could dominate the control unit at the expense of the business critical work.

In addition, the work going to the control unit could be coming from other partitions in the same Central Processing Complex (CPC) or from other CPCs (see Figure 3). In these cases device level priority queueing will not help.

With the introduction of the Enterprise Storage Server, WLM is now able to set a priority that is honored by the control unit. In this way, I/O requests that are destined to different devices, or come from other partitions or other CPCs can be given preferential treatment at the control unit.

Channel subsystem priority queueing addresses yet one other potential queueing delay for I/O; Channel subsystem resources. In the discussion above we mentioned two situations in which device level priority queueing will not help. In the case of I/O requests going to different devices (see Figure 4) or when the I/O requests are being issued in different

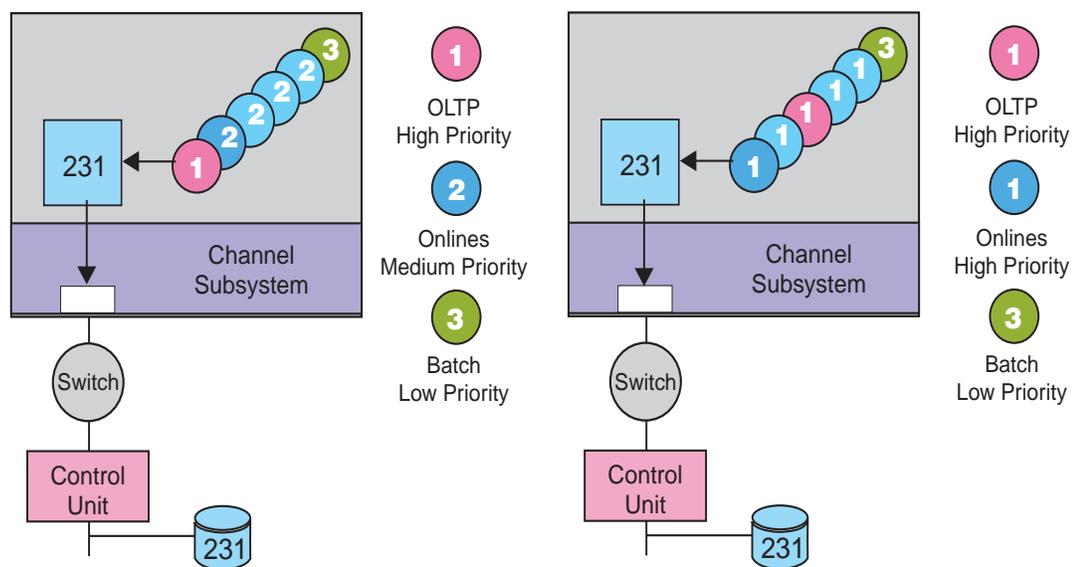


Figure 2- Priority Queueing with WLM

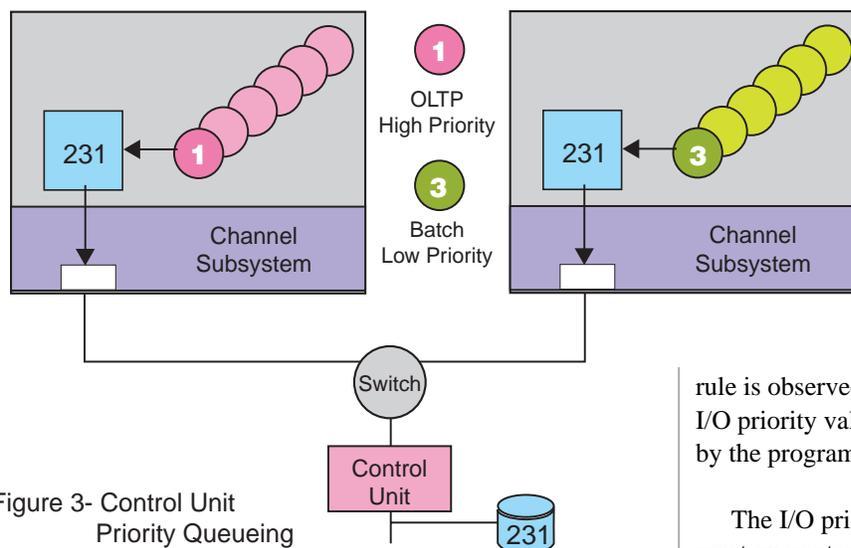


Figure 3- Control Unit Priority Queueing

partitions on the same CPC, the control unit priority queueing provided by the ESS will be of some help, but this presumes that the channel subsystem is able to get the I/O requests to the control unit in the first place. It also presumes that only one control unit is attached to the channels.

This is where CSSPQ helps. When I/O queues form within the channel subsystem, due to heavy I/O contention to several devices served by the same channels, (or the same device, but from different logical partitions), the channel subsystem will give preference to the requests with the highest priority.

How it works

The I/O part of the z/Architecture specifies that I/O requests are passed from Central Processing Units (CPUs) to the Channel Subsystem (CSS) by means of Start Subchannel (SSCH) instructions. These instructions are issued by the z/OS I/O Supervisor (IOS) on behalf of applications or other system functions. These requests are copied into Subchannel Control Blocks (SCBs) which are added to the tail of a particular Work Queue (WQ). Each SCB enqueued at the bottom of a work queue contains a field designating the I/O priority associated with this request.

As the CSS in zSeries 900 systems is implemented by one or more so-called System Assist Processors (SAPs), one WQ is associated with each SAP. The SAP is responsible for initiating I/O requests by dequeuing SCBs from the WQ and passing the request on to one or more channel paths designated to handle this request. In case all of the channel paths controlled by a particular SAP are busy, the request is re-queued at the bottom of the same or another WQ.

Therefore, work queues do not only contain requests (SCBs) just appended by any of the CPUs by means of SSCH instruction, but in addition SCBs which have been appended to the WQ due an “all channel paths busy” condition.

All other busy conditions (i.e. control unit busy, device busy and switch busy) are not handled via WQs. It should be noted here that I/O priorities are observed only when an SCB gets dequeued from a WQ.

Basic algorithm

The I/O priority mechanism is active when the corresponding global I/O priority switch is enabled, via the system console. Otherwise, a strict first-in/first-out

rule is observed which ignores the I/O priority value(s) specified by the program.

The I/O priority mechanism must guarantee a reasonable fairness algorithm so that on one hand I/O requests are handled according to their priority values but on the other hand no I/O request will starve due to low(er) priority values assigned to them. Furthermore, the algorithm should be fairly simple in order not to impose too much latency during priority evaluation.

The algorithm works by first taking any (either initial or redrive) I/O request off the top of the work queue, then taking the oldest highest priority initial request off the work queue, then the oldest highest priority redrive request off the work queue, and lastly the oldest initial request off the work queue, regardless of its priority.

In doing so, this algorithm makes sure I/O priority values as specified by software are well observed by the channel subsystem but starvation is totally avoided. The algorithm makes sure that

- (High frequency) high priority requests do not monopolize the channels at the expense of others with lower priorities.
- High frequent busy/redrive conditions do not block out initial requests (and vice versa). The alternate handling of initial or redrive requests is independent of any priority values.

Activating Channel Subsystem Priority Queueing

In order to activate Channel Subsystem Priority Queueing, you must first activate it at the service element. For a logically partitioned machine, you need also need to specify a range of values for each partition that the operating system image within the partition is limited to. The maximum range that WLM will use is eight, but you may choose a narrower range.

The machine will enforce that the I/O priority of any I/O requests from partitions are within the specified range. If it is too high, it will be set to the highest value allowed and if it is too low, it will be set to the lowest value allowed.

In order to have WLM manage CSS priorities, it must be running in goal mode and the I/O Priority Management option in the WLM service policy must be turned on.

How are the priorities set

The I/O priority that is passed on the I/O request to the channel subsystem is set by WLM. There are no externals that a user can use to specify the priorities on their I/O requests. Instead, WLM sets the priorities for I/O requests issued by a workload using this basic approach:

1. System related work is given the highest priority
2. High importance work missing goals is given a higher priority than other work
3. Work meeting goals is managed so that light weight I/O users will have a higher priority than heavy weight I/O users
4. Discretionary work is given the lowest priority in the system

In this way, Channel Subsystem Priority Queueing as just another tool that WLM may use to help it

achieve the goals for service classes as specified by the user.

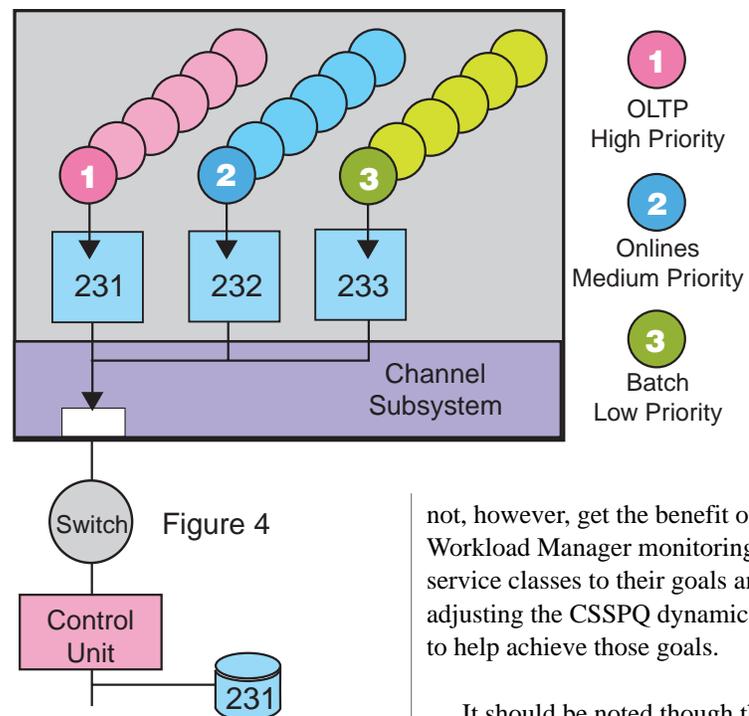
Multiple logical partitions

By activating CSSPQ and then specifying an allowable range, you can control how different partitions interact with it. If you are running other operating systems on the same CPC that do not support CSSPQ (e.g. VM, TPF, VSE and OS/390), you may specify a default priority for these partitions.

For example, if you have a VM partition that has work that is as important as the discretionary work

preference over the z/OS partition, but the z/OS partition would still be able to make tradeoffs across the work loads it was managing.

One does not even need to be running z/OS with WLM in goal mode to get some benefit from CSSPQ. If you have three workloads, one of high importance, one of medium importance and another of low importance; you may set the priority ranges to - for example (8-8), (7-7) and (6-6) respectively, and they will be managed accordingly. They will



in the z/OS partition, you might set the VM partition to a priority range of (8-8), and on the z/OS partition set the priority range to (8-15). This way I/O requests from the VM partition will be treated equal to that of the discretionary work from the z/OS partition (as opposed to even lower if you let the priority default to zero). Alternatively, you might set the priority range for an OS/390 partition, running your business critical OLTP applications to (16-16), while the z/OS partition was set to (8-15). In this way, the OLTP work would always have

not, however, get the benefit of Workload Manager monitoring service classes to their goals and adjusting the CSSPQ dynamically to help achieve those goals.

It should be noted though that partitions which are members of the same z/OS Logical Partition (LP) Cluster¹ should have all their ranges the same. In this way WLM will be able to manage service classes across the cluster in a consistent way.

Dynamic Channel Path Management

Channel subsystem priority queueing compliments the other I/O portion of IRD, Dynamic Channel Path Management (DCM) (see accompanying article). DCM allows the system to dynamically

redefine channel paths to control units, to more efficiently use the channel resources that are available to the processor.

The assignment of channels to control units by DCM is done with the cooperation of the Workload Manager, when it is running in goal mode. This allows DCM to give additional channel paths to control units that are being used by service classes that require that additional bandwidth to achieve their goals. These additional channel paths, however, are available to all work in the LP Cluster, even the low priority work. So adding channel paths may not get the desired results, if low priority work just consumes the added bandwidth.

This is where CSSPQ helps. By assigning the more important work a higher priority, we are assured that it will get greater access to the available channels.

How do I know if I will benefit from it

There are several situations where CSSPQ would be of benefit. One would be in environments where ESCON Multiple Image Facility (EMIF) is not currently exploited due to a concern that shared channels might be consumed by discretionary work in one partition at the expense of business critical work in another. CSSPQ may allow such environments to begin exploiting EMIF, which in turn may reduce total channel requirements for the processor.

In addition, in situations where unexpected demand on the system requires business critical workloads get preferential access to shared channel subsystem resources in order to continue achieving its goals, CSSPQ may help.

What prerequisites do I need?

In order to exploit channel subsystem priority queueing, you need to have an zSeries. In addition, in order to be able to have these priorities dynamically managed in order to help your business critical work achieve its goals, you will need z/OS with Workload Manager in goal mode. •

1. A Logical Partition Cluster is the set of logical partitions on one Central Processing Complex (CPC) which are all members of the same sysplex.

LINUX meets the enterprise servers

SUSAN BULEY, JOHN DIMARTINO, AND J.D. ROSS

Information for these articles was gathered by Susan, John, and J.D. during interviews with the following:

- Greg Burke, Program Director for Linux for zSeries
- Chris Rohrbach, S/390 Senior Software Designer for Linux for zSeries
- Boas Betzler, S/390 Senior Software Designer for Linux for zSeries

You have probably heard all the buzz about the introduction of the Linux operating system on to the S/390 hardware as well as the recent announcements about Linux on the new zSeries architecture. But what is all this buzz about running Linux on your enterprise server?

IBM has worked to bring Linux to both of the enterprise servers, zSeries and S/390, for several reasons:

- There is a large and rapidly growing number of programmers that are familiar with Linux. It is the most popular platform in universities and colleges today.
- There are a large number of applications available to run on Linux that can be quickly and easily ported to either Linux for S/390 or Linux for zSeries.
- Linux is an open and flexible system that when combined with the robustness of the enterprise servers, provides an industrial strength Linux environment.

- Using enterprise server architecture, customers can easily consolidate large numbers of Linux and UNIX servers onto a single platform.
- For many of IBM's customers, the majority of their data currently resides on S/390 hardware. Combining that data with the Web capabilities of Linux and a high-speed interconnection delivers a highly integrated e-business solution.

In the next articles we will introduce you to the Linux operating system, explain why the combination of Linux and an enterprise server, whether it is Linux for zSeries or Linux for S/390, is so beneficial for you, and we will also

try to answer the questions you have on this topic.

The Linux for S/390 software is already available for you to install (see the “download” section of the Linux for S/390 web site at <http://www.ibm.com/s390/linux>).

The Linux for zSeries software will be available soon. Since Linux for S/390 and Linux for zSeries are the same on most of the topics we will cover in this article, we use the term Linux for enterprise servers to refer to both versions. •

Hot topic: cool product- an introduction to Linux on the enterprise server

What is this Linux thing everybody keeps talking about? Chances are you have heard of this new wonder operating system. But what exactly is Linux? What can it do for you? More importantly, what can it do for your enterprise server?

Linus Torvalds announced the release of the original Linux kernel (base code) in August of 1991 in a news group of Minix (a clone of UNIX) users. When he released Linux 0.01 over the news group he asked for help to make Linux more workable. In under four months the users of Linux, those who became known as the Open Source Community, produced Linux 0.1. By 1992, the Linux operating system already had 1000 users and developers. By 1999, the number of faithful leapt to 9 million. The wave continues to this day. Industry estimates show that Linux was shipped on at least 25% of new servers last year. But this wave has just begun. IBM, SuSE®, TurboLinux®, and Red Hat® joined together and made Linux for S/390 possible, thus propelling Linux into the world of mainframes. But why has Linux captured the hearts of IS professionals so quickly? It's not just Tux, that cute little penguin used as a mascot.

The reason behind the strong support for the Linux operating system is its base in the Open Source Community. This group of Linux users reports bugs and fixes to other members of the Community. The Community then implements those fixes at the speed of the Internet. So running Linux on enterprise servers gives you the power of 9 million programmers and the strength and backing of IBM. The best part is that even with all these programmers involved there is still only one Linux. This ensures that Linux will remain pure.



There are different flavors or distributions of Linux but all share the base Linux code. The different flavors represent applications, wizards, and possibly support added to the Linux operating system by companies like SuSe, TurboLinux and Red Hat. Each

flavor is hardware platform neutral. A program you can run on an Intel platform you can run on a zSeries machine with a simple recompile in most cases. This makes porting programs fast, cheap, and easy. It also allows wide spread availability of commercially made applications.

The key e-business uses of the Linux operating system are in the server environment, providing core Internet based services such as mail, World Wide Web service, FTP file servers, and a growing role in the areas of LAN file and printer service. Industry estimates are that 60% of the production servers on the World Wide Web are Linux based technology

Now you ask, “What is the benefit of Linux on an enterprise server?” You receive many benefits from using Linux on a zSeries or on a S/390. Among them is power. The power to be able to consolidate an entire server farm to one enterprise server. The power of the enterprise server’s legendary reputation of availability (99.999 - a proven figure) and performance.

By running Linux on your enterprise server, the Linux operating system inherits the qualities of service that your S/390 hardware has traditionally provided and that the zSeries

will continue to provide, such as availability, manageability, and scalability. These attributes become vital as, with e-business, Information Technology changes from being a component of your business to being the core of your business.

The addition of Linux to your current environment is highly beneficial. Linux for enterprise servers has the ability to access data on a z/OS or an OS/390 image. The need for redundant data is minimized. Hey, penguins are friendly, remember? Network traffic between LPARS will be reduced by the introduction of HiperSockets. This break through technology will allow Linux to communicate with partitions containing Linux, z/OS, or OS/390 at memory speed rather than network speeds. The data will be transferred between partitions within the same server. This “network in a box” will reduce network latency and increases overall performance of the machine. The HiperSockets Interface is TCP/IP socket compatible and will transfer data transparently to the application.

Linux will also inherit enterprise server qualities such as self-healing and self-configuring attributes. The self healing attribute allows Linux for enterprise servers to detect and correct errors without effecting the application. This is done on the hardware level with no input from any application or person. There is very little costly down time for failure or replacement of hardware.

Self configuration is a valuable tool that Linux also inherits transparently from the enterprise server. It is not only beneficial because of the increased scalability but also for its increase in capacity. Your enterprise server is able to deliver resources such as processors and I/O paths to the applications that require it, in real time. This improves cost effectiveness by avoiding over configuration for complementary peaks in a web server for example. This will make consolidation into a single zSeries or a single S/390 cheap and easy.

Besides scalability, effectiveness and consolidation, long time non-failure is a key to a successful server. Nothing beats a mainframe’s availability. Even if there is an incident of a hardware failure the enterprise server will transparently swap out faulty hardware. Linux and any other applications will not be affected. Linux inherits all of these qualities from your enterprise server. No more restarts and no costly down time.

Walter G. Weir, Chief Information Officer of the University of Nebraska at Omaha agrees, stating, “For Linux to come and be on this environment opens up whole new vistas of possibilities in terms of what applications I’m going to run on S/390. This will enable me to get rid of some of my midrange systems that I don’t need anymore. It will help me consolidate support and so I will have less support costs to run a machine that I know is going to run well. But we now have an environment that I never have to worry about failing, it is so reliable, it is like clockwork, its like dial tone on the phone.”

And now that we have your attention, here are other abilities Linux will inherit from your enterprise server.

Both the hardware, S/390 and zSeries, and the Linux operating system support multitasking. The big iron can accommodate tens to hundreds of Linux images on a single machine. Actual figures depend on applications and workload and can be determined by a sales consultant. This replaces the server farms with low availability numbers. Not only do you get centralized management back at the mainframe level but VM can manage the workload of each Linux image at real-time. This means no more Internet servers running at 10% or lower capacity waiting for the next peak. Plus, on each of these hundreds of Linux images, you can run multiple users from a single Linux install. Linux has the ability to support more than one user at a time, typically multitasking so well that the other user is invisible to the first, second, third and so on. What you get is teamwork. Linux manages the hardware it shares with multiple users and VM manages hardware between the different Linux images. This is all going on transparently to the actual users and applications.

Linux is application-rich and it is powerful. Linux is the future and the future of Linux at IBM is clear, we are committed to Linux on all platforms. Consider how Linux for zSeries or Linux for S/390 could enhance your e-business. •

Getting personal with the penguin

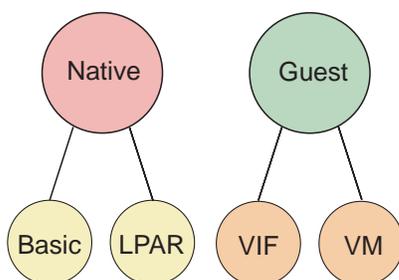
Linux plays an important role in the future of e-business, and it is important for you to understand how its introduction affects you and your installation. Just as Linux benefits from running on your enterprise server, your z/OS or OS/390 shop can benefit by adding Linux for zSeries or Linux for S/390 to your installations.

As an OS/390 programmer, the idea of adding the Linux operating system into the mix at your installation probably brings up a lot of questions in your mind. The Linux operating system is obviously powerful and popular, but you are used to the robust qualities of service that IBM's S/390 hardware and OS/390 software have always delivered. The same robust qualities that you know will continue and grow with the zSeries architecture and z/OS operating system.

In this article, we answer the questions that we have been hearing from OS/390 programmers as they learn about and plan for implementing Linux for zSeries or Linux for S/390.

Question: How do you install Linux on a S/390?

Answer: Linux for enterprise servers can be installed a variety of different ways. There are 2 modes that can be installed on an S/390: native and guest. These modes have several possibilities of their own.



A native install occurs when the Linux for enterprise server operating system is installed directly onto the hardware. There are two ways of installing Linux for enterprise servers in native mode. These two installs are basic and LPAR. In basic, there is a single Linux image running on the hardware. LPAR mode allows you to create up to 15 logical partitions on a single machine. Each LPAR is completely invisible to the other. This means that on one enterprise server you could run Linux, z/OS or OS/390, VM and any other S/390 compatible operating system side by side without affecting the other. This provides various combinations of configurations where you can share data between the logical partitions.

A guest install occurs when the Linux for enterprise server operating system is installed on top of another operating system that has already been installed on the system, such as VM. The two ways of installing Linux in guest mode are VM and Virtual Image Facility. Installing on VM allows you to define hundreds of virtual machines on a single enterprise server. You can use VM to run multiple Linux web servers or to test new Linux environments without duplicating hardware or disrupting production workloads. VM allows consolidation of a multitude of servers and allows Linux to inherit operational attributes. With Virtual Image Facility you can not only create but also manage multiple Linux images on your enterprise server. Depending on each images demands for resources, 10's to hundreds of Linux images can be run efficiently on one S/390 or one zSeries.

VM and Virtual Image Facility can be run in any of the logical partitions created by using LPAR mode. This unique combination of native and guest mode is the key to creating hundreds to thousands of Linux images on a single machine.

Once you have chosen the method that is correct for your installation, see the library section of the Linux for S/390 Web Site at <http://www.ibm.com/s390/linux/> for guidance information on using these methods.

Question: I have never heard of the Virtual Image Facility before. Can you tell me more about it?

Answer: Virtual Image Facility, or VIF, is a high performance environment that operates within a logical partition (LPAR). This management system enables 10 to hundreds of Linux images on a single enterprise server. The administrator can access, create, and manage these images quickly and easily through one console. VIF also contains commands to delete, start, stop, and increase DASD for the image or for paging. It also allows you to increase or decrease memory available to the image.

VIF can be run anywhere that Linux for enterprise servers can be run. The images can be given dedicated processors or engines or share those same resources with other software, such as z/OS or OS/390. The costs on resources of VIF can be as little as 5% of the system in support of up to hundreds of Linux instances.

Question: I currently have several UNIX applications ported to and running on OS/390 UNIX System Services and I was planning to port additional applications.

Should I consider porting all of my applications to Linux for enterprise servers instead?

Answer: The addition of Linux for S/390 was not meant to be a replacement for OS/390 UNIX, but rather as a complement. The same applies in that Linux for zSeries is not meant to replace z/OS UNIX. There is still a strong business need to continue running applications on OS/390 UNIX and to run applications on z/OS UNIX in the future. Use Linux for enterprise servers as a vehicle to greatly expand the number of commercial applications that you run at your site.

So now that there is a choice, how do you decide which operating system is the correct one for running your applications?

The decision to run your applications on z/OS UNIX (or OS/390 UNIX) versus or Linux for enterprise servers is determined by the business demands and workload requirements of each application. You should run applications on z/OS UNIX or OS/390 UNIX when you specifically need function or qualities of service contained in the z/OS or OS/390 operating systems.

A benefit of porting applications to Linux for enterprise servers rather than OS/390 UNIX or z/OS UNIX is that since 98% of the Linux code is identical across hardware platforms, applications are easily and quickly ported. In contrast, UNIX code has variations across hardware platforms and much more customization is involved when porting applications.

Although almost all applications can be ported to Linux, IBM recommends that you deploy those applications needing robust qualities of service and scalability provided through features such as Parallel Sysplex and Workload

Manager on z/OS UNIX or OS/390 UNIX. Developers who have made this investment will continue to see benefits from z/OS UNIX and OS/390 UNIX and there should be no reason to migrate these applications. Linux for enterprise servers is a new environment with rapidly growing acceptance in the marketplace. It brings the promise of a large number of applications as well as providing for faster porting of existing UNIX applications to the enterprise servers that do not need to exploit z/OS or OS/390 features or qualities of service.

Currently, several applications and packages that run on Linux for S/390 are available on the Internet. One source is the Thinking Objects Web Site at <http://linux.s390.org>. The number of applications and packages is constantly growing, so check the web site periodically for new additions.

Question: This has been an OS/390 shop for a long time now. We are used to the performance, support, and reliability that running these systems on the S/390 hardware brings us. Will Linux for enterprise servers provide us with the same high quality I am used to in these attributes?

Answer: The z/OS and OS/390 operating systems have an architecture that is designed to meet the world's most critical computing needs. The attributes that offer continuous availability of z/OS and OS/390 are built right into the system. Linux inherits reliability characteristics, such as CPU sparing, LPAR, EMIF, and error correction and detection just from running on the enterprise server architecture. However, the more advanced functions that are implemented architecturally in today's S/390 hardware and software and that of the new

Linux on the big server: an interactive seminar

Come to <http://ibm.com/s390/linux/seminar> (launched February 25, 2001), and explore the concepts behind Linux for S/390 and Linux for zSeries. The seminar's animated graphics demonstrate how the work of an underutilized middle-tier server farm can be consolidated at tremendous savings into the logical partitions of a single server. The seminar dramatically reveals possibilities: You can now create applications at your desktop and easily port them to the "big iron." Who has 10,000 spare machines and 10,000 extra LAN ports to set up as a test? You can now run such a test all inside a single box. See the installation possibilities — single-image, LPAR, guest and nonguest modes— graphically illustrated. How can Virtual Image Facility save you effort? The seminar clarifies this and several other features, too.

zSeries architecture, for example hardware-cryptography and data compression, would require corresponding support in Linux.

The marriage of the two, Linux and enterprise server hardware, offers you the best of both worlds. Linux for enterprise server combines the hardware benefits of decades of enterprise server architectural design, the design that Fortune 500 companies are running their businesses on with the fast moving, application-rich Linux operating system.

By running Linux on the enterprise server hardware, Linux inherits the superior self-configuring and self-healing attributes of the hardware. This means that you can immediately improve the reliability, availability, and scalability of a Linux application by your choice of enterprise server hardware.

Open Source code has a reputation of high quality. Members of the Open Source community that contribute code take great pride in their contribution and work hard to ensure that their code is defect free and will uphold the standards that open source code has earned.

To ensure the quality of service that our customers expect, IBM Global Services provides fee based consulting, planning, implementation services, defect and remote technical support as you have come to expect. They will help you answer how-to questions, help define problems and determine their source. IBM consultants are available to help you evaluate your Linux for S/390 requirements and to assist in implementing and optimizing your Linux for S/390 solutions. These services will also be available for Linux for zSeries.

So bottom line, what does this all mean? It means that you can

have hundreds to even thousands, of powerful Linux images running applications with basically no down time - the most reliable Linux available.

Question: How do I back-up and restore my data on Linux for S/390?

Answer: With Linux for S/390, there are several methods available for backing-up and restoring your data. The option that is best suited for you depends upon the specifics of your installation and your familiarity with UNIX tools. Using the table on the facing page, choose the scenario that matches your situation to find the method of back-up and restore that IBM recommends.

When you have selected the method of backing-up and restoring data that is correct for you, consult the Linux for S/390 Web site at <http://www.ibm.com/s390/linux/> for additional information to help you implement your method.

If you plan to use Linux for zSeries, check our Web site at <http://www.ibm.com/s390/linux> for up-to-date information on choosing the best back-up and restore method for you.

Question: Will my middleware work with Linux for S/390?

Answer: IBM is committed to enable all of its key middle ware to run on the Linux environment. In the fourth quarter of 2000, the following middle ware was made available: CICS® Transaction Gateway V3.1, MQSeries® Client, WebSphere® Advanced Edition V3.5 with Java JDK V1.2.2, DB2 UDB Enterprise Edition V7, DB2 Connect™ V7, IMS™ Connect, and Tivoli® Storage Manager Client. In addition, a beta version of Tivoli Management Framework with selected applications was released

in forth quarter of 2000, with General Availability slated for the first half of 2001.

For additional information on middle ware and availability dates, please see the Linux for S/390 web page at <http://www.ibm.com/s390/linux/>.

Question: Is there a difference between Linux for S/390 and Linux for zSeries?

Answer: Yes. The new Linux for zSeries will run on the IBM eserver zSeries architecture. Linux for zSeries plays a vital role in the next generation of e-business.

Linux for zSeries is a 64-bit version of Linux that runs on the new IBM eserver zSeries 900 architecture and exploits the benefits of this new architecture. It contains a 64-bit version of the Linux kernel, the key device drivers, the compiler and the Runtime library. Running the Linux for zSeries 64-bit version gives you several advantages, such as support for 64-bit real and 64-bit virtual modes, and being able to use main storage above the 2 GB line.

Linux for S/390 is a 31-bit version that runs on the current S/390 hardware machines. Linux for S/390 can be run on the new zSeries architecture, but it will only work in 31-bit mode. You need the 64-bit Linux for zSeries to take full advantage of the zSeries architecture.

Question: Now you really have my interest. Where can I get more information on Linux for enterprise servers?

Answer: If you want to learn more about Linux for S/390 or Linux for zSeries, you should visit the IBM Linux for S/390 web site at: <http://www.ibm.com/s390/linux/>. This site contains the latest

news and technical information on both Linux for S/390 and Linux for zSeries.

The Linux for S/390 Web Site also contains a library of documentation for your use in installing and preparing to install Linux on S/390. The manuals are available

in PDF format for you to view or download and print. The library will grow to include similar documentation for Linux for zSeries soon. Watch the Web site for details.

The redbook, *Linux for S/390 (SG24-4987)*, helps you install the Linux for S/390 Marist College distribution in different environments, and documents basic system administration tasks that help you manage your Linux for S/390 system. This book is

	Scenario 1	Scenario 2	Scenario 3
Type of Customer	Are you currently using a client/server method for backing-up and restoring data, such as OS/390, Tivoli Storage Manager from IBM, or ARCserve from Computer Associates?	Are you familiar with UNIX tools for backing-up and restoring data?	Do you run Linux along side a VM system and need a solution that provides disaster recovery?
Required knowledge of UNIX or Linux	Very little knowledge of UNIX or Linux required. Knowledge of your current client/server back-up and restore method is important.	Experience using UNIX tools is useful.	Very little knowledge of UNIX or Linux required. Knowledge of VM is important.
Recommended back-up and restore method	If you already have an existing client/server method for backing-up and restoring your data, just connect your Linux for S/390 with your existing back-up and restore method and use it for your Linux data.	Use existing Open Source options for backing-up and restoring data. Good examples are: Amanda, or the tar, pax, or cpio command.	Use your VM system for disaster recovery of the data on your Linux for S/390.
Benefits	You already know how to use your current method.	Open source options are readily available for download from the web. The tools tar, pax, and cpio are available with your Linux for S/390 system.	If you are already running VM, you know how to establish disaster recovery there.
Where to get more information on this method	Review the documentation that came with your current method of backing-up and restoring data. ARCserve is available at: http://www.ca.com/ Tivoli Storage Manager is available at: http://www.tivoli.com	Amanda is available for download at: http://www.amanda.org A description of how to use this tool is provided in the redbook, Linux for S/390 (SG24-4987), available at http://www.ibm.com/redbooks .	Review the documentation for disaster recovery that came with your VM system.

available on the Internet at
<http://www.ibm.com/redbooks/>.

If you would like to take a class, IBM Learning Services provides basic Linux education as well as a 3 day workshop, *EW04 - Linux for S/390 Implementation*. This workshop covers such topics as: distributions, installation, administration, security, and S/390 Virtual Image Facility. Contact IBM Learning Services at 1-800-IBM-TEACH (1-800-426-8322) for details on when this class is offered.

To subscribe to the Linux for S/390 discussion group with other customers and IBMers that share your interest in the topic of Linux for S/390, send a note to listserv@marist.edu. The body of your note should contain:

SUBSCRIBE LINUX-390 your-name, your-organization
 Once you are subscribed, you will receive further instructions on how to use the mailing list.

If there are other topics concerning Linux for zSeries or Linux for S/390 that you would like to see discussed, please send a note to rossjd@us.ibm.com.

64-bit real operation in OS/390 Release 10

BOB ROGERS AND JIM GROUND

Great news: OS/390 Release 10 supports more than 2 gigabytes of central storage when operating under the new z/Architecture on a zSeries processor.



Kerry Collia

OS/390 Release 10 can operate either in the older ESA/390 architecture or in z/Architecture. In ESA/390 architecture (on an S/390 G5 processor, for example), it continues to support expanded storage and a maximum of 2 gigabytes of real storage like prior releases of OS/390. Expanded storage is not supported when operating under z/Architecture on a zSeries processor; however, interfaces that previously used expanded storage (for example, hiperspaces) work compatibly using real storage, rather than expanded storage. To understand the impact of z/Architecture, consider the basic memory management concepts introduced by 370-XA 20 years ago that are still a part of OS/390 today.

System overhead managing memory

A system image, in native mode or logical partition mode, is configured with some amount of central storage. The CPUs of the system image can access data directly only when that data is in central storage. Data not in central storage must be read in from I/O devices or paged in from paging devices before the CPUs can operate on that data.

Each user or job is assigned a 2-gigabyte virtual address space. However, the total amount of virtual storage allocated for users and jobs typically exceeds the total amount of central storage available. Virtual pages that are not currently being accessed are paged out to paging devices and then paged back into central storage when they are needed. Two types of paging devices can be used for this: direct access storage devices (asynchronous), and expanded storage (synchronous). Many systems today are configured with enough central and expanded storage to almost eliminate that asynchronous paging to external

devices. Because there is typically more virtual storage than can be held completely within 2 gigabytes of central storage, the virtual pages must time-share that available central storage. For every page brought into central storage, other pages must be pushed out. To avoid delays when central storage frames are needed, the operating system continuously pages in some of the virtual pages — an action called “page stealing”. Anticipating future demand for central storage, the operating system selects which pages to page out based on a Least Recently Used (LRU) algorithm, assuming that the most recently referenced pages will soon be referenced again. The operating system determines which are the least recently referenced pages by maintaining an Unreferenced Interval Count (UIC) for each page in central storage. Moving and tracking these pages is a lot of work for the operating system. This overhead affects system performance.

Note:

Occasionally, a page must be “fixed” in central storage to prevent it from being paged out. The page is fixed to allow I/O channels address central storage directly. I/O channels and devices are directed by channel programs. Channel programs contain only real addresses and all of the data areas they refer to must be fixed in central storage.

Enhanced addressing architectures and new channel programming formats have been introduced to deal with the larger real addresses. This creates an additional complexity: when pages are fixed for I/O, they must be placed in frames with real addresses that can be accommodated by the channel programming format being used. The location in real storage where a page will be fixed is controlled by the program that allocates the virtual storage — either by default or by explicit specification. If a page resides in the “wrong type” frame when it is requested to be fixed, the operating system moves the data from its original frame to a frame of the “right type”. This process potentially increases overhead, but it also provides a strong level of compatibility for old programs.

Introducing 64-bit real addressing

64-bit real addressing can theoretically allow central storage to support up to 2^{64} bytes (16 exabytes) of memory, which is 8 billion times 2 gigabytes or a trillion times 16 megabytes. It is not expected that systems will actually have this incomprehensibly large amount of real memory for decades. The first IBM z900 processors can be configured with up to 64 gigabytes of real central storage, which is only four-billionths of the memory that could eventually be accommodated by a 64-bit architecture.

Two major aspects to 64-bit support in OS/390 R10 when operating under z/Architecture are:

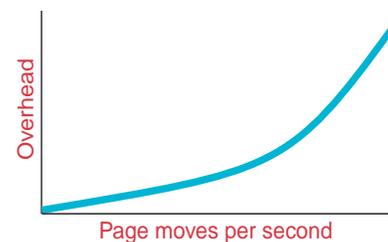
- Support for greater than 2 gigabytes of central storage
- Elimination of support for expanded storage

The value of 64-bit real support to you

64-bit real addressing in z/Architecture eliminates data movement between central storage and expanded storage. This elimination improves the performance of the system and enables customization that can improve response time and performance without incurring unnecessary CPU overhead. It is not expected to reduce external paging to direct access devices unless, for a given workload, the amount of configured real storage increases beyond the previous total of central and expanded storage.

The use of 64-bit real addressing reduces system overhead; but how much overhead is a current system experiencing with central plus expanded storage? There are a number of different overheads associated with expanded storage paging. Some are subtle and not reported by typical performance monitoring tools. For example, overhead rises as the computing power of the system increases. Overhead also increases with the number of CPUs, due to both hardware and software serialization. Even the count of pages moved per second cannot be interpreted in a straightforward way; demand paging driven by application page faults is much more expensive than movement due to page stealing or swapping. Page fault processing spends much more time in other storage manager code than it does actually moving data from expanded to real storage.

Determining the exact overhead experienced by a system due to expanded storage paging is complex due to the non-linear behavior of the overhead. For relatively low page move rates, the overhead is negligible; but, for very high movement rates, the overhead is punishing. The following schematic graph depicts the approximate shape of the cost-of-page-movement curve.



Many customers will find that in their current configuration, expanded storage paging overhead is minimal. For them, the value of 64-bit real addressing comes as added confidence that their systems can grow without real storage constraint and the potential to consolidate multiple system images into one image without experiencing real storage limitations.

What do you have to do to use 64-bit real addressing?

To use 64-bit real addressing support in OS/390 R10:

1. You must have a zSeries system. (z/Architecture is not available on any S/390 system)
2. Apply all IBM and ISV maintenance to enable middleware products and applications and to tolerate the changes. For the IBM maintenance, see the PSP buckets; for ISV maintenance, see <http://www.ibm.com/s390/s390da/osnp.html>.
3. Configure all processor memory as real.

4. Add the statement ARCHLVL 2 to LOADxx.
5. IPL, and the system does the rest.

You may require additional real memory depending on your system configuration when running in z/Architecture mode due to the increase in the size of some internal control structures and code.

It's that easy!

64-bit real addressing provides immediate real storage constraint relief for systems that have done a lot of paging under ESA/390 architecture. Furthermore, it enables continued growth and possible system image consolidation without constraining real storage. And, it is easy to enable. •

Self-managing your software licenses and costs

VICKY POLLLEN

Today, when running our state-of-the-art operating system, OS/390 V2R10, on our state-of-the-art machine, Generation 6 of the Enterprise Server, the system can tell you which IBM software products have been installed. It cannot tell you, however, if those are the same products for which you are licensed, and thus paying for each month. You must track that information on your own. The challenge is to keep the following two lists in synch:

- products you are licensed for, by machine, and
- products you are using on each machine.

To the extent that these lists differ, you may be paying for products you no longer use, or running products out of compliance with the Terms and Conditions you have agreed to follow. In addition to keeping your own records current, each time you want to make a simple change, such as enabling an optional priced feature, discontinuing a product, or moving a product to another machine, you must contact IBM 30 days in advance. The larger the number of software products and servers you

run in your shop, the more difficult these management tasks become.

With z/OS running on an IBM eserver zSeries 900 server, we are bringing the task of managing software product licenses into the electronic age, reducing the burden on you to manage them on your



Kerry Collia

own. IBM License Manager (ILM), a new function provided with z/OS, product license certificates, files that represent the terms and conditions contained in your product license agreement, and new hardware features unique to the z900 server, form the basis for this new support.

ILM was developed in partnership with ISOGON[®] Corporation, a market leader providing solutions for enterprise software asset management. It is based on the X-Open Software License Manager standard from The Open Group and is active only when z/OS is running in z/Architecture (64-bit) mode on a zSeries 900 or equivalent machine. ILM is a generic tool, with an API which any software publisher can use to help you manage that publisher's software products' licenses. IBM is using ILM to implement the new zSeries pricing model, called Workload License Charges (WLC).

When you order any Variable Workload License Charge (VWLC) product or one of a select set of Flat Workload License Charge (FWLC) products for a z900 server running z/OS, a product license certificate (or certificate) is sent along with the product code. The certificate is an electronic representation of the terms and conditions of the product license. VWLC product certificates contain a defined capacity value, which indicates the

portion of the machine on which the product is licensed to run and is the basis for the product charges. You install the product certificate at the same time you install the product code, using the ILM Management Tool.

At any time, you can use the ILM Management Tool to display the certificate, and thus license, data for products installed on your machine. You can also use the Management Tool to update certificates directly. These updates correspond to changes to your product license agreements, take effect immediately on your machine, and are recorded in a log file. You must use the Management Tool at least once per month to generate a reconciliation message file from the log file, and forward to IBM. Each certificate change will then be processed as an order, used to update IBM's inventory records, and included with the data used to generate your bill for the following month. Changes that can be made using the Management Tool include enabling or disabling optional priced features, discontinuing products, moving a license from one machine to another, and, for VWLC products, increasing or decreasing the defined capacity (i.e. the portion of the machine on which the product is entitled to run).

This use of ILM is called passive license management. Products come with certificates, which are installed on your system, and used to display and update product license data.

All VWLC products and some FWLC products take use of ILM one step further, using it for active license management. These products modify their code to call ILM each time the product is started, asking if a license can be granted. The product then continues to come up or will not start, based on the answer returned.

When an ILM-enabled product is invoked, ILM is called to determine whether or not the product is allowed to start on that server. ILM checks to ensure that the values of an installed product certificate match the current running environment (for example, machine serial number). For a VWLC product, the check includes ensuring that running the product in the designated LPAR will not exceed the capacity defined for that product on the product certificate.

If the conditions are met, the product is allowed to continue. Under normal circumstances, if any of the conditions is not met, the product will not be allowed to start. However, to give you a startup period and an ongoing way to recover from error situations or disaster conditions with minimal disruption to your key workloads, several operational modes are supported including normal mode, exception mode, observation

mode, and emergency mode. When products are allowed to run even though there is a mismatch between the product certificate and the system environment, exception data is written to the log file. Exception records are included in the reconciliation message file sent monthly to IBM, and used to build exception reports used by IBM and customers to determine what changes must be made to return the customer to full license compliance. This can include updates to product certificates (e.g. increasing a product's defined capacity) or the system environment (decreasing the size of an LPAR, or starting a workload in a smaller LPAR).

ILM gives you greater control and awareness of your IBM licensed software. It help you to manage your licensed software, allows you to make changes to your software licenses that take effect immediately, and gives you more control over your total software costs.

For more information on ILM and WLC, see *z/OS Planning for Workload License Charges* on the Web at <http://www-1.ibm.com/servers/eserver/zseries/zos/bkserv/>, or the information at <http://www.ibm.com/eserver/zseries>.

A piece of the action

JOHN URBANIC

You've spent all day at the county fair, and now you can't wait to get a piece of delicious home-made pie, so you head over to the pie booth. You see a man with a sign that says, "Apple Pies, \$8.00."



Kerry Collia

"How much for a slice?" you ask. "We're not selling individual slices," the man says. "I have to pay for the whole pie, even if I can only eat a slice?" "I'm afraid so."

You continue on to the next pie seller, a woman with blueberry pies.

"What are you charging?" you ask. "It's \$8.00 for pie," she says, "And \$1.00 for a small slice, 3.00 for a medium slice, and 5.00 for a large slice." "Great — I'll take one large slice."

Believe it or not, the man and the woman at the pie booth represent the OS/390 and z/OS pricing models. Prior to z/OS, IBM software products running on OS/390 were typically priced based on the computing capacity of the central processor complex (CPC) on which the software was running. This pricing method has some limitations, including the following:

- If you are running multiple workloads (using IBM licensed software) on the same CPC, you are charged for those products as

though each IBM license is using 100 percent of the CPC all of the time. Even if you want just a piece, you pay for the whole pie.

- If you purchase more system capacity (hardware) for an existing CPC, your software charges will increase, even if the added capacity is intended for future growth or for new workloads with different software products.

Software has become a large part of most I/S budgets. And with the explosive growth of new workloads, particularly in the area of e-business applications, IBM wants to make it more affordable for you to deploy these new workloads — and grow existing workloads — on a zSeries 900 server. This is the goal of Workload License Charges (WLC).

With z/OS running on a zSeries 900 you can use a combination of hardware and software to manage your costs and your system capacity with WLC. Under WLC, you tell IBM how big a piece of the action you want. That is, determine the portion of processor capacity you need to run your workloads. For selected IBM products running in these workloads, you are charged for the capacity you define rather than for the entire CPC. In this way you are charged only for the resources you need. Other advantages to this pricing model are:

- Purchasing additional hardware for future growth does not increase your software bill.
- You have the flexibility to manage e-business workload spikes.
- You can benefit from improved price performance as your workloads grow.

How are products priced?

There are two categories of products under WLC:

- Variable Workload License Charge (VWLC) products — you can specify a capacity less than the total capacity of the CPC on which the product runs. Your cost is based on the capacity you define. The IBM License Manager manages your use of the product to that capacity. Note that variable WLC is not available for all products supported by z/OS.
- Flat Workload License Charge (FWLC) products — you are charged a flat rate for these products, independent of the size of the CPC on which they run and how you are using them. That is, you pay for the use of the product on a CPC. You can run one or more instances of the product on this CPC.

How does WLC work?

For VWLC products, WLC allows you to specify that a product runs on less than the entire capacity of a CPC. A CPC's capacity is represented in terms of millions of service units per hour (MSUs).

For example, you can specify that a logical partition (LPAR) in which z/OS runs uses no more than 50 MSUs on a 200-MSU CPC. The Workload Manager makes sure that the LPAR averages no more than 50 MSUs over any four-hour period. This method is referred to as a "soft cap," in that it allows for spikes in your workload by allowing the usage to rise above 50 MSUs at times, while maintaining a 4-hour rolling average. Note that you can change these capacities dynamically to fit the needs of your installation.

Suppose your system has a capacity of 200 MSUs and is divided into three logical partitions (LPARs) as shown in the figure above. To see an example of what your cost is for a product, consider the following CICS scenario.

CICS runs in LPAR1 (25 MSUs) and LPAR2 (100 MSUs) for a total of 125 MSUs. For this example, we assume the same version of CICS is running in both LPARs. With WLC you pay only for the 125 MSUs you request for CICS.

200 MSUs		
LPAR 1	LPAR 2	LPAR 3
25 MSUs	100 MSUs	75 MSUs
CICS	CICS	DB2

Also in this figure:

- z/OS is priced based on a capacity of 200 MSUs and DB2 is priced based on 75 MSUs.
- LPAR1 is considered to have a defined capacity of 25 MSUs; LPAR2 has a defined capacity of 100 MSUs. A defined capacity (soft cap) refers to the LPAR capacity you define for a workload. A defined capacity for an LPAR allows for spikes

in your workload while maintaining an overall 4-hour average, and is defined in MSUs. A defined capacity applies to all the products running in an LPAR, regardless of the number of products the LPAR contains.

- CICS has a license certificate capacity of 125 MSUs. Under WLC, each product you install includes a license certificate. The IBM License Manager manages your use of the product according to the terms and conditions on the license certificate.

For more information on WLC, see *z/OS Planning for Workload License Charges* on the Web at <http://www-1.ibm.com/servers/eserver/zseries/zos/bkserv/> •

License Manager standards

TOM CIERECH

The X/Open Software License Management (XSLM) Standard is the end product of GUIDE's License Management (LM) Project. The GUIDE Project consisted of a core team of volunteers who collectively defined a license management standard based on the licensing requirements which had been submitted by both customers and vendors during a previous phase of the LM Project.

Customer requirements centered on controlling and managing growing software costs while remaining compliant with the vendor's software terms and conditions. Customers wanted the ability to monitor their software usage so that they could reduce their costs by consolidating and/or eliminating the need for specific software on a subset of the ma-

chines within their data centers. They were also interested in a consistent way of managing licenses across all vendors and platforms and wanted more flexible and granular licensing/pricing models.

Vendor requirements centered on ways of supporting new more flexible licensing models and on a compliance tool which would make such new granular license models feasible.

The core team for the GUIDE LM project consisted of customer representatives from British Columbia Government and Burlington Industries. The GUIDE Project's vendor reps were from Boole & Babbage (BMC)[®], Compuware[®], Gradient Technologies[®], IBM, Isogon[®] and Novell[®].

Note that the LM Project has been adopted by and is now a SHARE Project.

Creating a License Management Standard which satisfied all of the customer and vendor requirements was not the only challenge. In some respects the greater challenge was to create a standard which addressed license management in a consistent way across the entire enterprise, could be reasonably implemented on every major hardware/software platform and was easily extendable in the future. Creating a standard which could be implemented on every major platform was an especially difficult task, since the infrastructure and technologies commonly used and taken for granted on some platforms are not available on all platforms.

The XSLM core team eventually addressed the challenge by carefully choosing where to place the protocol boundaries which determine what to make a part of the standard and what to leave as implementation dependent. The XSLM Standard defines:

1. the application and management APIs, which provide a consistent way of making license requests and managing the licensing system.
2. the format, content and rules for license certificates, such that license certificates created with any compliant license certificate generation tool could be used by any compliant license system implementation. The license certificates and log records consist of self defining data elements, so that the standard can be extended to include additional capabilities by defining new self defining data elements in future XSLM releases.
3. the format, content and rules for logging license certificate and license granting related events, such that logs can be processed by any set of compliant tools which can process XSLM log files.
4. licensing system behaviors, such that customer and vendor can be assured of a consistent set of behaviors from all compliant licensing systems when the same set of license certificates are installed.

5. function towers as the means of specifying optional XSLM defined capabilities, so platforms can have licensing system implementations which are still compliant with the base XSLM standard without requiring the implementation of all XSLM defined functions. Each function tower can contain one or more of the following elements: a list of APIs, license certificate data elements and log data elements each of which is uniquely supported by that tower.
6. the Licensing System Section of the license certificate, which allows each licensing system implementation to support additional capabilities not defined in the standard. Such a section allows the licensing system publisher to specify those additional capabilities by adding the appropriate licensing system specific data elements to the license certificate.
7. the Publisher Specific Section of the license certificate, which allows each software application publisher to include any publisher unique information on the license certificate for retrieval by the software application via the application API. Except for passing this information to the application, the licensing system performs no other actions on the data elements in the Publisher Specific Section.

The XSLM core team chose to make the following areas implementation dependent so that licensing system implementations are free to choose the most appropriate technologies for each platform supported by that implementation:

1. the communication protocols between the licensing system components
2. the file system for storing license certificates
3. the file system for storing log data

4. the mechanism to sent log data from the customer to the publisher
5. the externals for all of the execution related licensing system components: the license certificate generation tool, the license management tool, the application agent, the management agent and the license server.

The specification created by the LM Project was adopted by the SYSMAN Project of The Open Group (TOG) as Version 1 of the X/Open Software License Management (XSLM) Standard in February of 1999. Note, however, that the IBM License Manager (ILM) being shipped with z/OS R1 is based on Version 2 of the XSLM specification, which is still being documented. The XSLM core team expects to submit Version 2 of the XSLM specification to The Open Group for their review and formal adoption as a standard in 1Q01. Also note that z/OS R1 only supports a subset of what will be the XSLM Version 2 Standard and that future z/OS releases will support additional XSLM functions.

Additional information about Version 1 of the standard, as well as XSLM newsletters, presentation schedules by the core team and XSLM links can be found at the XSLM core team's Web site:

<http://www.xslm.org> •

Exploiting z/Architecture with EXCP I/O

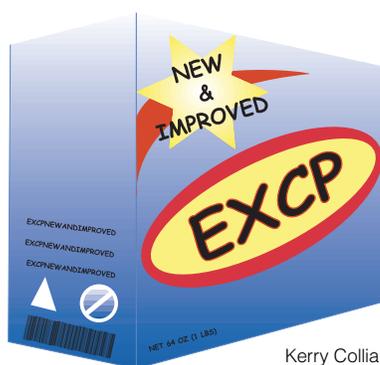
GREG DYCK

This article covers recent (and some not so recent) improvements to the EXCP interface for those assembler programmers who enjoy dealing with the bits and bytes at a hardware/assembler level. The primary focus is on changes made in OS/390 V2R10 to exploit z/Architecture mode. It assumes the programmers already have a basic understanding of CCW chains and EXCP/EXCPVR programming. It does not cover the entire programming interface,¹ nor is it a tutorial on how to write an EXCP application.

The words “Execute Channel Program” strike terror in the heart of many a systems programmer. The Execute Channel Program, or EXCP, interface in its full blown glory can be overwhelming. A “don’t touch” attitude often exists for applications that use the interface. And, for the most part there is no need to because the basic interface has remained unchanged. Yet, while the EXCP interface has remained upward compatible since the inception of MVS, new options have crept in that are not well known. These options exploit new hardware features on the processors; selective use of them can improve I/O performance. The changes started with MVS/XA™ 2.2 APAR OY31351 and have continued quietly through OS/390 V2R10, when you have very good reasons to look at these applications. Why? Because small changes to existing applications that use EXCP or EXCPVR can get great performance benefits from the new z/Architecture.

EXCP 64-bit IDAW support

Why are 64-bit IDAWs important to you? Because they allow I/O to be done using real storage above 2G in z/Architecture mode, which eliminates the need for the system to move your data into storage below 2G before doing the I/O. The movement of data can be very costly in terms of CPU overhead.



Kerry Colliia

Some IBM access methods, such as QSAM, have taken advantage of 64-bit IDAWs without requiring effort from you. To support z/Architecture, EXCP changed in OS/390 V2R10 to transparently build 64-bit IDAWs for devices that support them. For OS/390 V2R10, the support is limited to DASD and tape devices.² Minor changes must be made to EXCP applications to get the benefit and avoid movement of data for EXCP I/O.

You don’t need to change the EXCP request in the application to exploit 64-bit IDAWs; however, to request storage for I/O buffers that can be fixed in real above 2G, you must change the application. For DASD and tape applications, a simple change to the LOC specification when the I/O buffer is obtained using the GETMAIN or STORAGE OBTAIN services does this. These services have been extended to allow you to specify

that storage can be fixed in real above 2G, in addition to the current options — a specification of LOC=(xxx,64) gives this attribute to the storage. You can specify this in both z/Architecture and ESA/390 modes. In ESA/390 mode, it is treated as if you specified LOC=(xxx,31).³ Applications that are assembled using LOC=(xxx,64) run compatibly on all supported releases of OS/390 and z/OS.⁴

A program can determine if 64-bit IDAWs are supported by looking at bit UCBEIDAW in the UCB for a device. UCBEIDAW is set only when the system is running in z/Architecture mode and the device support code supports them. An example of checking UCBEIDAW and getting appropriate storage is:

```
LR R1,UCBPTR
Copy common UCB address to R1

IOSCMXR MF=S
Locate common extension

USING R1,UCBCMEXT
Establish addressability

TM UCBFL7,UCBEIDAW
Does device support 64-bit IDAWs?

JZ B31REAL
No, back I/O buffer below 2G
Yes, allow backing above 2G

DROP R1
STORAGE
OBTAIN,LENGTH=BUFLen,LOC=(RES,64)
J BDONE
Join common code

B31REAL STORAGE
OBTAIN,LENGTH=BUFLen,LOC=(RES,31)
BDONE ST R1,BUFADDR
Save the I/O buffer address
```

Using an IOBE

Why is the IOB Extension, or IOBE (mapped by IOSDIOBE), important? Because all new EXCP/EXCPVR facilities that are explicitly enabled are done so using it. In 1991, APAR OY31351 introduced the IOBE to EXCP because there was insufficient reserved space left in the IOB to introduce new features.⁵ Support for hardware, such as for FICON (including the FICON bridge) channels or Enterprise Storage Subsystem (ESS, aka “Shark”) DASD, has been added to the IOBE over the years.

An IOBE becomes part of the EXCP/EXCPVR request by setting flag IOBCEF to ‘1’b in the traditional IOB, and then loading the address of the IOBE into R0 before invoking the EXCP/EXCPVR service. An example of building and passing an IOBE is:

```
OI IOBFLG4,IOBCEF
An IOB Extension exists
```

```
XC IOBE,IOBE
Clear the IOBE
```

```
MVC IOBEID,=C'IOBE'
Set the eyecatcher
```

```
MVI IOBEVERS,IOBEVRSC
Set the version code
```

```
LA 0,IOBE
Set R0 for the EXCP request
EXCP IOB
Issue the EXCP request
```

Using an IOBE, an EXCP application can enable the following facilities.⁶ The OS/390 V2R10 DFSMSdfp Advanced Services book provides additional information in the use of these facilities.

- Request EXCP return up to 32 bytes of sense data following a unit check.
- Control the issuance of ERP error messages requiring operator intervention for tape devices.

- Request EXCP return a second CSW address if an error occurs when the control unit is executing ahead of the channel due to prefetching. (FICON™)
- Indicate that channel synchronization is not required on a read/write CCW transition because different storage areas are used for the read and write CCWs. (FICON Bridge, only)
- Indicate that Streaming Mode Control is disabled for the request. (FICON)
- Indicates that the device can bypass the channel program extent collision checking. Extent range enforcement remains active. (ESS DASD only)
- Indicates that the device is to bypass the validation checking of the parameters on Define Extent and Locate Record commands. Extent enforcement remains active. (ESS DASD only)

An EXCPVR application can request the same facilities as an EXCP application, but additional facilities exist for use only with EXCPVR:

- Indicate that prefetching of CCWs and data can be done by the channel. (FICON)
- Indicate that the channel program will not be modified during execution, other than to add CCWs at the end. (FICON)
- Indicate that channel synchronization is required after the CCW following a PCI. (FICON)

Use these facilities with care because incorrect specifications can cause unexpected results.

EXCPVR with 64-bit IDAWs

Unlike EXCP, applications that use EXCPVR in z/Architecture mode must change the CCW chains that they build if they wish to use I/O buffers backed in real storage >2G. In addition to changing the GETMAIN or STORAGE OBTAIN request, they must build a

64-bit IDAW to replace the 31-bit IDAW used in ESA/390 mode and build an IOBE to specify that it is being passed.

All CCWs in a CCW chain that specify use of an IDAW must use the same format IDAW; you cannot intermix 31 and 64-bit IDAWs in a single CCW chain. Flag IOBEEIDA in the IOBE indicates if a CCW chain uses 31 or 64-bit IDAWs. When IOBEEIDA is set to ‘1’b an application has built 64-bit IDAWs. 64-bit IDAWs are valid only when running in z/Architecture mode with devices that support 64-bit IDAWs.

Each 64-bit IDAW must begin on a double word boundary, and processes up to 4K of real storage. This is a change from 31-bit IDAWs, which each processed up to 2K of real storage. The 64-bit real address of the storage can be stored into the IDAW with the Store-Real-Address, or STRAG,⁷ instruction. An example of building a list of 64-bit IDAWs using STRAG is:

```
OI IOBEFLG2,IOBEEIDA
Using 64-bit IDAWs
```

```
LA 2,IDAWAREA
Set start of IDAWs
```

```
L 3,IOBUFFER
Get the I/O buffer start
```

```
LH 4,IOLLENGTH
Get the I/O length
```

```
AR 4,3
Calculate the ending address
```

```
NILL 4,X'F000'
Truncate end address to
4K boundary
```

```
IDAWLOOP STRAG 0(2),0(3)
Store a 64-bit real address
```

```
LA 2,8(,2)
Bump to the next IDAW location
```

```
NILL 3,X'F000'
Truncate address to a 4K boundary
```

CR 3,4

Are we done?

BE IDAWDONE

Yes, exit the loop

AHI 3,4096

No, bump to next 4K boundary

J IDAWLOOP

Continue building IDAWs

IDAWDONE DS 0H

EXCPVR with Format 1 CCWs

OS/390 V2R10 allows an EXCPVR application to build and execute format 1 CCW chains. Format 1 CCWs can be used in both ESA/390 and z/Architecture modes.

Format 1 CCWs can be backed anywhere in real storage below 2G. QSAM uses this support to move some of its CCW chains into storage above 16M real. Any EXCPVR application can do the same with the storage used for its CCWs. APAR OW46454 extends the format 1 CCW support to allow the CCW chains to be located anywhere in virtual storage.

Format 1 CCW chains are requested by setting IOBEFMT1 in the IOBE to '1'b. When this flag is set, the IOBSTART field is treated as a 31-bit address,^{8 10} the CSW address in the IOB will be 31 bits starting at field IOBFLAG3 and continuing through IOBCMDA, and the IOBSIOCC data will be returned in field IOBESIOC.

Conclusion

This article describes improvements to the EXCP interface, recently and over the years. This includes support for real storage >2G in z/Architecture mode and how you can take advantage of the support, how to use an IOBE to specify new options to EXCP, including requesting return of 32 bytes of sense data, and specification of FICON channel control flags to optimize I/O. The advanced assembler programmer can make additional changes to EXCPVR, including using support for format 1 CCW chains, 64 bit IDAWs in z/Architecture mode, control flags to optimize FICON channels, and control flags to optimize IBM ESS DASD devices. Now, it is up to you to examine your EXCP/EXCPVR applications and make changes to improve their performance, especially in z/Architecture mode. •

1. See the DFSMSdfp Advanced Services book for additional information on using EXCP and EXCPVR.
2. 64-bit IDAW support requires device recovery code changes. Only DASD and tape have been updated at this time.
3. LOC=(xxx,31) is treated the same as LOC=(xxx,ANY). LOC=(xxx,31) is the new and preferred format to use.
4. APAR OW46410 must be installed for downward compatibility.
5. Not all fields defined in the IOBE mapping are part of the intended programming interface.
6. APAR OW43557 must be installed for all of the following options to be available.
7. HLASM V1R4 provides support for z/Architecture instructions.
8. The processing described is new with APAR OW46454, and is not reflected in the currently published documentation.

Ramp up to IBM License Manager

JOHN CZUKKERMANN

On October 3, 2000, IBM announced a new software pricing model for the eserver zSeries 900 platform. The new model, called Workload License Charges (WLC)¹, is designed to support next generation e-business workloads. With WLC you can define the capacity upon which your software charges will be based for key IBM software products. To implement this new pricing model, z/OS introduces a new function: IBM License Manager (ILM).

As stated in the October announcements, ILM will be delivered initially on March 30, 2001, as part of z/OS Version 1 Release 1, and coincident with the availability of z/OS and WLC at full machine capacity. IBM testing of ILM will continue after March 30, 2001, and ILM will be made available for full production use when this testing is completed. Testing is planned to complete between the end of 2Q01 and the end of 3Q01. The ILM Pre-Production Period (PPP) is what we are calling the period of time between z/OS availability and availability of the service required to make ILM production-ready. WLC at less than full machine capacity will be available once you have applied this service to your z/OS system, and meet all other eligibility criteria. The purpose of this article is to help you understand how you can use the ILM PPP as your ILM ramp up time and to let you know where to find more information as we make it available.

We are creating a z/OS element level website for ILM that will be the primary communication vehicle to share information about the latest ILM developments. We will inform you about the availability of this

website via the Workload License Charges website². We plan to include information such as:

- Hints and Tips for License Administrators and System Programers
- General Q's and A's
- Links to IBM publications
- Announcements pertinent to the ILM PPP



Tracy Bondi

Other communication vehicles such as a listserver discussion list, Washington System Center Flashes and a PSP bucket are being investigated as well. Our goal is to give you the necessary information to enable you to leverage the PPP to gain first hand experience with ILM.

Rather than wait until we have fully completed our testing of ILM, we feel it is important to enable you to start experimenting with ILM as early as possible. This gives you the chance to check out ILM and experiment with various system configurations to understand ILM's operational characteristics without fear of impacting your current software bill. During the PPP, you will be charged at full machine capacity for all Variable Workload License Charge (VWLC) products. We will announce when we have completed our testing of ILM, and the service required to make ILM production-ready is available. This

also signals the availability of WLC at less than full machine capacity. At that time, you will be required to apply the recommended z/OS service to make your system ILM production-ready as one of the WLC Terms and Conditions.

What's the plan?

As you can imagine, we are tracking the progress of ILM development and testing very closely. Our current plan is to make ILM function available incrementally through the PPP, allowing you to begin experimenting with ILM prior to production use. The latest information about what function will be available and when will be shared through the ILM website and other communication vehicles as appropriate.

As functions become available, you will be able to perform activities such as installing product certificates, and changing certificates and LPAR definitions to understand how the license system reacts. We are planning to support performing reconciliation with IBM and viewing the resulting web-based reports, allowing you to become more comfortable with this process. The data submitted for reconciliation during the PPP will not affect your software bill in any way.

We want to ensure that you have the ability to experiment with ILM so you have a better understanding of how it works and how it will affect your day to day operations, before you are required to use it. The role of License Administrator may be new to some of you, so you will need time to learn the ins and outs of license management using ILM. This role is especially important because changes to the capacity values on

variable WLC product certificates can directly affect your software bill. Becoming proficient at performing the reconciliation process and understanding how ILM helps you stay within your license Terms and Conditions will be important as well. When ILM is production-ready and use of ILM becomes mandatory, the knowledge and experience you gained during the PPP will be invaluable.

Until our testing is complete and we have declared ILM to be

production-ready, you will have the option to not use ILM at all. To that end, we are taking steps so ServerPac and SystemPac systems are delivered so ILM either starts or doesn't start during IPL, as appropriate. We are looking into the possibility of including a memo about ILM in CBPDO orders, so those customers will have a reminder to check on the status of ILM and configure your system accordingly. All customers should monitor the ILM website for the latest ILM status.

We are working hard on ILM to optimize the benefit you can gain from PPP, and to ensure a smooth transition to your use of ILM in a production environment. As details regarding PPP and the transition to production become available, we will publish them on the ILM website. •

1. Announcement 200-354, *IBM Announces Workload License Charges*, can be found on IBMLink at <http://www.ibm.link.ibm.com> in the Announcements section under InfoLink.
2. http://www.ibm.com/servers/eserver/zseries/wlc_lm/

ISVs and the IBM License Manager

PAUL CARROLL

WARNING!!!

This is NOT a technical article. If you want bits and bytes or hints and tips, you won't find it here. However, if you are interested in what Independent Software Vendors (ISVs) think about IBM's License Manager and what their plans are for exploiting it, read on!

z/OS is the first platform, IBM or otherwise, to deliver the Open XSLM technology. This will be done with the IBM License Manager implementation on z/OS. As you have already seen, IBM's software products such as z/OS, DB2, and CICS to name a few, are exploiting this new technology. However, IBM is not the only company excited about these new standards. Over a year ago, IBM began working with over one hundred ISVs on this technology. Through presentations given at meetings of ISVs, vendors were made aware of the technology and how it would be implemented. Additionally, several large vendors took part in joint work groups with IBM to better understand this new technology and how IBM planned to implement it.

Many of the ISV's who will have products on z/OS have told us they to plan to exploit the License Manager technology. Some of the larger ISV who have expressed an interest in this technology are Allen Systems Group (ASG)[®], BMC Software[®], Candle Corp.[®], Compuware[®], Landmark Systems[®], and Software AG[®], to name but a few. (A website with a list of ISVs who plan to exploit z/Architecture can be found at <http://www.ibm.com/s390/s390da/osnp.html>) In particular, this article will focus on two companies, BMC Software and Allen Systems Group, and their plans to leverage the value of the XSLM standard as implemented by the IBM License Manager.

Ron Higgin, Director, Technology Planning for BMC Software, was a founding member of the License Use Manager (LUM) task force initiated by the GUIDE user group back in 1995. (Ownership of the LUM effort was transferred to the SHARE user group in the summer of 1999.) Ron was instrumental in making sure the ISV voice was heard when the

open XSLM standards were being formulated. Ron will provide some perspective on the XSLM standard and BMC's support of the standard's stated objectives. Ron will also offer some insight into BMC's early research activities focused on determining the most effective means of leveraging the inherent value of the License Manager technology.

Jim Bladich, Senior Pricing Analyst at Allen Systems Group, was exposed to the new technology through meetings with IBM. Although ASG's implementation of License Manager will take place over time, Jim will talk about how ASG is being aggressive in implementing a new product licensing methodology, called ESL, for ASG's product suite. The ESL will immediately introduce a workload capacity based pricing methodology to ASG's z/OS product line.

Here's what Ron Higgin from BMC Software has to say:

XSLM Background

The XSLM standard is the result

of over 5 years of work by a dedicated team of individuals representing the customer, application ISV, and software license management solution developer communities. The XSLM project began as a customer-driven initiative of the former GUIDE International user group back in 1995. The Open Group — a merger of the former X/Open and OSF standards groups — published the first version of the XSLM (standard) specification in early 1999. Ownership of the XSLM project was transferred to the SHARE user group in the summer of that same year.

The XSLM core team worked closely with IBM during both the design and implementation phases of the z/OS License Manager. This allowed the License Manager designers to insure their final product would comply with the XSLM specification. Conversely the core team was able to leverage the experience gained from this implementation activity to refine and clarify the specification. We expect The Open Group will publish the resulting (revised) specification sometime during the first half of 2001.

BMC Software Involvement with XSLM

The XSLM development team was pleased with IBM's decision to deliver the industry's first XSLM standard based software license manager as part of their recently announced z/OS operating system product.

BMC Software has actively participated in the development of the XSLM standard from the very beginning, and continues to do so. We strongly support the standard's two primary objectives:

1. To provide (through XSLM-compliant licensing system implementations) customers with a standardized, cost effective, and platform agnostic means of managing their software license resources — and —
2. To facilitate the implementation of new software pricing models that more closely relate the business

value of a product to the license fee charged for the use of that product. Such models are often referred to as use-based pricing models.



Paul Rowntree

BMC Software Use of the IBM License Manager

BMC Software is evaluating the z/OS (IBM) License Manager – ILM — as a potential replacement for its current proprietary license management technology. Early results have confirmed the feasibility of doing this for some subset of the z/OS-based BMC product portfolio. However, based on these early studies, the subset of XSLM-defined function implemented in the first License Manager release will not satisfy the technical licensing requirements for many BMC products. Thus while customers may see some BMC products enabled for use with the production-certified ILM relatively soon after its availability, the decision as to whether or not to ILM-enable many products must necessarily be deferred until a fully XSLM-compliant ILM is made available.

BMC supports the initiative to exploit industry standard (XSLM) licensing technology such as that implemented by the ILM. We believe such technology will ultimately provide significant value to our customers by:

- A. Simplifying the technical management of the customer's product license inventory.
- B. Facilitating the creation of innovative licensing models that provide more direct correlation between the business value

received from the use of a product and the license fee paid for that product.

Finally, it worth noting that BMC has chosen *not* to use the new WLM “Virtual Server ID and Capacity Query Service” function as a use-based license model implementation vehicle. There are a number of reasons underlying this decision. The most compelling reason is the fact that the License Manager — unlike the “Virtual Server ID and Capacity Query Service” function — offers the capability to aggregate product usage (in terms of CPU capacity) across multiple zSeries 900 partitions (LPARs) running the z/OS operating system.

About BMC Software

One of the world's largest independent software vendors, BMC Software delivers the most comprehensive e-business systems management software with the fastest guaranteed implementation. This Service Assurance™ strategy enhances the availability, performance and recoverability of companies' business-critical applications. Companies can use this management methodology to demonstrate their ability to deliver optimal service to their customers and partners by joining BMC Software OnSite™, a certification program that includes solution implementation and regular HealthChecks performed by BMC Software Professional Services.

BMC Software is a Forbes 500 company and a member of the S&P 500, with fiscal year 2000 revenues exceeding \$1.7 billion. The company is headquartered in Houston, Texas, with offices worldwide. For more information please visit BMC Software's Website at www.bmc.com, or call 800-841-2031 or 713-918-8800.

As you can see, Ron and BMC have been part of this for a long time and are aggressively working to implement the IBM License Manager. Now here are the thoughts of Jim Bladich from Allen Systems Group:

ASG recognizes that flexibility and budget predictability are critical to the measured success of an IT organization. In support of its commitment to its clients, ASG will develop its products to take full advantage of the 64-bit Z-series architecture and to facilitate workload pricing through both IBM's License Manager and the z/OS virtual server.

ASG values its relationships with its clients and is committed to providing ease-of-administration when it comes to contracting.

IBM's License Manager is a valuable tool for streamlining the vendor/client administrative relationship, said Bladich. This process assists our clients in maintaining their competitive advantage in the rapidly evolving IT marketplace.

ASG is a responsive, full-service provider of enterprise software solutions, including world-class operations, applications and business intelligence solutions. To best combine ASG's extensive product portfolio with the measurable price and performance advantages of workload pricing, ASG introduced its unique Enterprise Software License (ESL) program. ESL provides access to ASG's 120-product portfolio for a flat, value-based fee - where per-product, annual usage fees are based only on the LPAR's (or total MIPS) on which the product runs.

Emphasizing partnership with its customers, ASG's ESL will provide instant evaluation and installation of all ASG products in a real-time scenario, requiring no additional contract negotiation. Utilizing the "pay-for-what-you-use" methodology

of the IBM License Manager and z/OS Virtual Server, ESL delivers value to the customer by offering tremendous savings in administration costs, reductions in lengthy negotiations, predictability in budgeting and the undeniable competitive advantage of being able to instantly react in a competitive environment.

You can see that not only IBM, but the major ISVs are also planning to exploit the new technology of IBM's License Manager.

Although the ISVs are excited about the License Manager technology, don't expect immediate implementation by them. Full adaptation of the technology includes more than just changing some product code and issuing digital certificates. It also involves a company changing its back end business process and accounting programs as well. These are items that take much time and planning. IBM will continue to work with the ISVs to assist them in exploiting the License Manager technology. •

About ASG

Established in 1986, ASG provides software and services to the Global-5000 in the areas of Enterprise Operation Management, Enterprise Application Management and Business Intelligence. ASG helps these organizations improve the performance of their data center operations and facilitate their transition to e-business. Heralded as the recognized price/performance alternative in enterprise systems management, ASG is dedicated to providing integrated, world-class solutions to data centers where consolidating costs is imperative. Flexible in contracting and unparalleled in customer support, ASG is committed to consistent growth, stability and developing best-of-breed products in emerging technology fields.

ASG employs more than 800 associates worldwide. Headquartered in Naples, Florida, ASG has development, support, service and sales professionals on five continents and a network of value-added partners.

z/OS makes your system programmers more effective!

FRANZ-PETER BOLEY

Do your OS/390 system programmers complain about their workload? Are you under pressure to reduce your OS/390 system programming costs? Do you have trouble finding good OS/390 system programmers and getting them trained? Well, z/OS now exploits a new technology to greatly facilitate the configuration of z/OS systems. This article shows what this new technology, called *z/OS Managed System Infrastructure for Setup (msys for Setup)*, is all about.

It all started with Web-based assistants.

For about a year now, OS/390 has offered a number of Web-based assistants, or “wizards”. These wizards provide a convenient way for going through planning and ordering tasks, as well as for defining the configuration parameters of some OS/390 components.

Configuration wizards came about because configuring OS/390 products and OS/390 components is manual, documentation-driven process. Products usually come with Program Directories, Installation Guides and/or Configuration Guides that describe instructions that you have to carry out manually. Sometimes this means setting hundreds or even thousands of configuration parameters. In the process, you are often referred to detailed descriptions in one set of manuals, while usage information or recommendations – if available at all – are often in yet another set of books.

Wizards provide web-based, interactive dialogs for configuration that guide the user through a set of high-level questions. Wizards use defaults and best practises values wherever this is possible to drastically cut down on the number of decisions that a user has to make. Since the configuration process is handled by a program (instead of being a manual process), input is immediately checked for syntactical and semantic correctness. The right level of context-sensitive online help also is available, eliminating the need to dig through piles of books to find information. At the end of the configuration definition, the wizards generate job skeletons that apply the defined values to the system.

- Automatically updating the configuration resources in the system
- Understanding what the current system configuration is about
- Providing a central place, the *z/OS Management Directory*, for storing configuration definitions.

Automatic update of system configuration

Wizards require you to download the generated job skeletons to your host, adapt them to your needs, submit them and handle possible error situations. *msys for Setup* automates this task. The product-specific customization dialogs store the configuration parameters in the z/OS Management

Directory, along with so-called “update requests”. Update requests describe how the configuration is to be applied to the system (for example, set a field in a specific parmlib

member to a defined value or create a new user ID with a set of defined attributes).

After the configuration is completed, the user starts the process of updating the system configuration, which automatically carries out the update requests. Update requests are processed by support services on the host. A support service is available for each configuration resource, like parmlib members or the security databases.

Although this is an automated process, the user is still in control of what’s happening. The update



Extending wizard technology

msys for Setup builds upon this wizard technology. It employs the same easy “interview style” as wizards for defining configuration values. As with the wizards, *msys for Setup* does not ask users to specify hundreds of configuration parameters. Instead, *msys* uses defaults and best practises values, has built-in sizing algorithms and derives the low-level values from answers to high-level questions. *msys for Setup* extends this technology by:

requests can be viewed before the task is started. The update process also can be simulated in a “dry run” which exercises all the code for the update, without actually updating the system resources. An extensive log informs the user about what has been changed, or when in simulate mode, would have been changed in the system.

In case anything else goes wrong during the update, you can “undo” the processed update requests. msys for Setup also allows you to back up an msys configuration that you know is working. You can restore it at a later time when you encounter any problems with a new configuration.

Understanding the current configuration

In contrast to the wizards, msys for Setup runs directly on your system. Thus it can understand the current system configuration and make it available for the customization dialogs. msys for Setup enables the system to tell you what it already knows, so you don’t have to re-type information that you provided earlier.

This refresh management directory capability also allows to bring the content of the management directory back in sync with the actual configuration in the system. This can be necessary, for example, after emergency changes have been done directly to the system without working through msys.

z/OS Management Directory

As indicated, msys for Setup introduces a z/OS Management Directory. Today, configuration definitions are scattered across the entire system. For example, they exist in parmlib, in configuration data sets outside of parmlib, and in assembler macros which need to be assembled and link-edited.

With msys, the management directory will become the central repository for all configuration data. Thus configuration dialogs and management applications will have a consistent interface for getting at the current configuration of z/OS components and products. The z/OS Management Directory is an LDAP-based directory. We chose LDAP for the implementation, since it is a de-facto industry standard for directory services, and it’s widely used on a variety of platforms. We tried to introduce not again new, OS/390 or z/OS proprietary techniques with msys for Setup, but rather rely on technology which is available on other platforms as well (and for which skill is available from other platforms). This is also the reason, why msys for Setup internally uses XML extensively for the communication exchange between its components.

Want to know more about LDAP?

LDAP stands for “Light-weight directory access protocol”! It was derived from the X.500 standard for directory services. While X.500 was quite complex because it could handle all kinds of “exotic” situations, LDAP took the more commonly used elements from X.500. In addition it changed the underlying communication protocol from the OSI stack to TCP/IP. So, LDAP can be viewed as a light-weight X.500! LDAP is becoming increasingly popular. LDAP server implementations are available on virtually all platforms. Since V2R5, there is also an LDAP server implementation available on OS/390.

Driving the configuration of the target

Configuring system components may require to be able to

configure a system which is not IPLed or even IPLable. To solve this, msys for Setup allows to configure a system different from the one that runs msys. The system on which the msys for Setup code runs is called the “driving system”, the one that is being configured the “target system”. Both can be the same image, or they can be different, “connected” via shared DASD and a shared RACF database. To be able to configure components of a new level of z/OS, msys for Setup allows to run the driving system on a couple of release “below” the later z/OS release. You will then have to mount the msys HFS of the latest z/OS on your driving system. Speaking of the system environment, the LDAP server could be on the driving system or the target system, or on a completely different system.

What’s the value of all this?

The main advantage to using msys for Setup is reducing your total IT costs. For example, msys for Setup will help make experienced system programmers much more efficient when configuring z/OS components or products on z/OS. Smaller z/OS shops also will benefit, because msys for Setup can reduce the skill level for configuring a z/OS product.

msys for Setup provides a consistent, easy-to-use interface that can be used by all products running on z/OS. And it can ensure that a configuration is correct before it is applied to the system. The matrix below summarizes your benefits when comparing the old, manual configuration world with the new msys world.

Sounds too good to be true?

In z/OS R1, this new configuration process will be available for setting up a Parallel Sysplex

resource sharing environment, including support for the Intelligent Resource Director, XCF signaling and Global Resource Serialization Star. msys for Setup comes as a base element of z/OS. It is an integral part of z/OS, and no extra charge is associated with it.

msys for z/OS R1 allows to use OS/390 releases down to V2R8 as driving system. You then have to mount the msys HFS from a z/OS R1 system on the driving system. Since it wouldn't be too attractive to allow to configure a Parallel Sysplex consisting of only z/OS R1 systems, the Parallel Sysplex exploitation allows to configure OS/390 V2R8 (or higher) systems as well. However, at least one system, the driving system or one of the target systems must be z/OS R1 (you have to get the msys code somehow, don't you?). Note that this z/OS R1 system must be installed via CustomPac (i.e., ServerPac or SystemPac®), or if

you upgraded from OS/390 V2R10, the V2R10 system must be a CustomPac system. For a detailed description of the pre-requisite hardware and software, have a look into *z/OS Planning for Installation*, or the *z/OS Managed System Infrastructure for Setup: User's Guide* which will become available at GA of z/OS R1.

Watch this space! The infrastructure that is provided in z/OS R1 is open for further msys exploitation!

Additional information
For more information on z/OS msys for Setup, go to:
<http://www.ibm.com/eserver/zseries/msys> •

Benefits to using msys for Setup

In the past, you needed to....

- manually update a number of system resources.
- worry about multiple different interfaces and their syntax.
- calculate formulas.
- decide how to set each value.
- check validity after the configuration was applied to the system.
- read several feet of documentation.

With msys....

- updates to the system are done automatically.
- the system takes care of the syntax and interfaces.
- the system does the calculations.
- only a small number of values is shown when a decision is required.
- the validity of a configuration is checked by the system.
- online help is just a mouse click away.

Parallel Sysplex in z/OS Made Easier!

JOSE CASTANO

Users have long complained about the many planning and customization manuals that must be read and understood to correctly configure a Parallel Sysplex. This complexity often results in errors in the configuration and setup of a Sysplex. Some configuration errors are easily corrected and result in little or no impact to the customer. However, others may result in Parallel Sysplex outages - exactly counter to the availability objective of establishing a Parallel Sysplex! In recent Parallel Sysplex outage analysis studies, approximately 15% of multi-system outages were the result of configuration errors. In order to address these potentially disruptive errors, IBM initially

delivered the Parallel Sysplex Configuration Assistant to make configuring a Parallel Sysplex easier.

Now, with z/OS, IBM is extending the Configuration Assistant concept with significant new function to make Parallel Sysplex set up as easy as possible. System Programmers will be able to use msys for Setup to configure and administer a "basic" Parallel Sysplex. This includes configuring XCF signaling, multi-system serialization, and Coupling Facility resources. Further, this support will help configure the Coupling Facility resources required for the newly introduced Intelligent Resource Director. This article

describes the new Parallel Sysplex set up based on the z/OS msys for setup infrastructure.

Using msys for Setup allows Parallel Sysplex setup to go beyond the simple set up provided in the Parallel Sysplex Configuration Assistant. In particular, set up activities and decisions are aided with context sensitive help, detailed syntax checking, cross dependency management and incorporation of "best practices" to assist a system programmer in correctly configuring a Parallel Sysplex - all without referring to a stack of manuals. Parallel Sysplex setup with msys applies to both new and existing Parallel Sysplex

implementors. For customers with standalone systems or basic sysplexes that want to create a Parallel Sysplex, it is a quick and easy way to establish a foundation that can be extended to exploit Parallel Sysplex resource sharing functions such as shared tape, system logger, and Enhanced Catalog Sharing (ECS). For customers that already have configured a Parallel Sysplex, the definitions that presently exist for



their environment are read into msys and its z/OS management directory and then reflected in the customization dialogs. This delivers the complexity reduction benefits when the customer wants to make changes or add another system to the existing Parallel Sysplex. Further, as IBM extends the exploitation of msys for setup function to include other exploiters, customers will have have experience with msys for Setup, and will be well positioned to take advantage of the new capabilities more quickly.

msys for Setup

Before reviewing the details of the Parallel Sysplex exploitation of msys for Setup, let's first review msys for Setup itself. msys for Setup is made up of two principal parts. The first part is the infrastructure component. This component offers the services and functions to provide query and update support of the z/OS environment as well as a data store to save customization data. Before completing this article, you may want to read "z/OS makes your system programmers more effective" in this newsletter, which

provides the details of the concepts and technology provided by the msys for Setup infrastructure. The second part of msys for Setup is the exploitation of the msys for Setup infrastructure that "plugs into" the msys workplace. The Parallel Sysplex plugin uses the msys services to provide customization support. The services provide methods to create customization dialogs, learn about the environment that is being customized, and also update that environment. Each product or solution that wants to offer customization via msys must use the msys services provided by the infrastructure component. z/OS 1.1 introduces both the infrastructure and the first exploiter, Parallel Sysplex configuration.

Parallel Sysplex Configuration using msys for Setup makes it easier!

When configuring a Parallel Sysplex prior to z/OS 1.1, system programmers are very dependent on documentation and its correct interpretation. For example, to configure and customize the environment that msys supports in z/OS 1.1 would require the system programmer to read and understand the intricacies in the *MVS Setting up a Sysplex, MVS Planning: Global Resource Serialization, MVS Planning: Workload Management*, as well as the *MVS Initialization and Tuning Reference* manuals. One of the goals of the z/OS 1.1 Parallel Sysplex support is to reduce the amount of research and skill required to customize a z/OS Parallel Sysplex. One of the ways this is accomplished is to provide context sensitive online help in all the customization dialogs. The help is relevant to the task at hand and is detailed enough to enable a decision for the dialog or input field without referring to documentation.

Context sensitive help can address the selection and creation of a single parameter, but it doesn't begin to address more complicated situations where a system programmer needs to understand information from several books and multiple disciplines. For instance, when setting up a Sysplex, a system programmer requires knowledge spanning several areas: hardware (for the coupling facilities and their links), operating systems (for the setup of sysplex signaling and serialization), and security (for providing appropriate access to the environment). Also complicating matters are parameters or values that affect other parameters when they are changed. These dependencies and interrelationships are not always obvious and can negatively impact availability if not well understood. A few examples of Parallel Sysplex interrelationships are:

- **SYSPLEx** parameter in **LOADxx** and **COUPLExx** must be the same. This is also an example in which a value must be specified in multiple places. The value is inputted only once by the user and the plugin propagates the value as necessary.
- Changes to the number of systems in a sysplex may require the **MAXSYSTEMS** keyword to be changed in the Couple Dataset and a new CDS to be formatted. Also, the structure sizes of the signaling structures may have to be changed.
- Addition of new Structure Paths for XCF signaling in the **COUPLExx** member has to be reflected in the active CFRM polity.

msys for Setup manages these dependencies and ensures that either the appropriate related change is made or notifies the msys administrator of the depen-

dency which may require an action on the installation's part. For instance, to relieve the system programmer of the burden of remembering to format a new CDS (bullet two above), if a new Couple Dataset must be formatted, msys will automatically allocate and format the dataset, and then notify the msys administrator to issue a SETXCF command to activate the new CDS. The administrative complexity is managed by the msys exploitation code owned by the product/component development team at IBM.

Another advantage of customizing with msys is syntax checking. For example, today system programmers can optionally run the RNL syntax checker (ISGRNLCK) to verify that GRSRNL updates are syntactically correct. This is an extra step that unfortunately is often ignored. The Parallel Sysplex msys exploitation plugin, as well as future msys exploiters, includes syntax checking as an integral component of customization. As the administrator provides input, it is validated before the user moves to the next customization dialog. The actual syntax of each customization parameter or library member is masked from the end user who must only know the values related to the customization dialog input fields. The plugin then translates the values into keywords which are sent to the host to generate the proper syntax. This ensures that updates are correct and will not result in an error which may cause or extend an outage. Furthermore, because the Parallel Sysplex plugin constructs and executes the IXCMIAPU and IXLIDSU utilities on behalf of the enduser based on dialog input fields, the syntax and complexity of these utilities are masked.

One of the features that sets the Parallel Sysplex msys plugin apart from common systems management tools is the incorporation of "best practices". Based on field experience and detailed outage analysis related to Parallel Sysplexes, the plugin includes extensive field "experience" to help prevent common pitfalls and also to promote "best practices". For instance:

- Placing primary and alternate on the same volume can cause sysplex-wide outages. The plugin issues a warning message when such a condition is detected. The user can choose to ignore it (if for instance the Parallel Sysplex is a test sysplex with limited resources) if appropriate.
- Too many consoles with MSCOPE(*ALL) can expose a customer to WTO buffer shortages. The plugin issues a warning message when this condition is detected.
- The plugin prompts for naming conventions so that all the couple datasets can be named in a consistent way which simplifies operations.
- The plugin provides the option to allocate and use shared SYS1.DAE which is preferred in a Parallel sysplex environment.

Another aspect of best practices is to use functions that have been shown to provide value in many customer situations. The msys Parallel Sysplex plugin implements an "XCF Star". This means that XCF signaling is configured to use Coupling Facility Structures and not CTCs. It is easier to manage and add to a Star configuration than it is to an XCF CTC configuration. Further, for the star configuration, the plugin allocates two signaling structures on different Coupling Facilities to avoid a single point of failure (another example of best practices). If a customer has a Parallel Sysplex configured using

CTCs, the Parallel Sysplex plugin will maintain those definitions, but also add new structure definitions. A GRS Star is also established as part of the of the configuration. A GRS Star is superior in terms of recoverability and performance compared to a GRS Ring and so is the default environment configured by the plugin. If a customer has an existing GRS Ring, it will be converted to a GRS Star.

msys for Setup - some additional thoughts

In the "z/OS makes your system programmers more effective" article, the concept of driving and target system is introduced. Because the driving system executes the updates for the target system, the driving and target system must be in the same Parallel Sysplex. First, the RACF® database must be shared between the driving and target systems so that it can be queried and update requests can be driven against RACF. This usually means one Parallel Sysplex. Second, we may need to update an active Parallel Sysplex. An example is updating the active Coupling Facility Resource Manager (CFRM). This must be done within the Parallel Sysplex. And last, but not least, as we extend this function, the concept of self discovery will be expanded so that less prompting will be necessary. An example is discovering the Coupling Facilities attached to a system, and querying their identification. We could then prime the Coupling Facility definitions in the CFRM with that information, thus eliminating required input and reducing the chance for an error when entering the Coupling Facility serial number. This means we must have direct access to the resources and having the driving and target systems in the same Parallel Sysplex facilitates this.

The customization process that msys for Setup provides is a multi-step process that allows customers to insert all the check and balances that they have today. At a very high level, the first step is the customization step. This generates the update requests but does not update the production environment. This is similar to the process where a separate system or sysplex is used today to make the updates to the target environment. The updates are actually in libraries that are duplicates of the production environment. When the updates are reviewed and tested, then a change window is used to copy the changes to production. The same method can be used with msys for Setup. The update process is divided into three subtasks which

allows the user to review the updates in as much detail as needed or required by installation's system management processes. Also, simulation of the update process can be used for further validation. Once comfortable with the update requests, a change window is used to execute the update requests and thus update the production environment.

z/OS introduces msys for Setup and sets the tone for the future of setup and administration in a z/OS environment. Parallel Sysplex exploits this technology to help with the numerous facets of setting up a sysplex. This includes incorporating best practices, being aware of product and inter-product dependencies, and consolidating the research required to configure a basic Parallel Sysplex. The thing not to lose sight of with this introduction is that Parallel Sysplex is just the first exploiter. The more important deliverable is the open infrastructure which provides these services to all the

other components and products that run on z/OS so that they can also become exploiters and offer the same easy to configure environment!

Additional information

For more information on z/OS msys for Setup, go to:

<http://www.ibm.com/eserver/zseries/msys>

For more information on Parallel Sysplex, go to:

<http://www.ibm.com/s390/ps0> •

What Happened to ServerPac for z/OS?

JOHN EELLS

What did we do to ServerPac?

Well, we've made more than a few changes. Now you can automate assigning data sets to volumes to get to IBM's recommended system layout in about 60 seconds. A new



Kim Mulé

View and change facility gives you 28 built-in groupings you can use for all sorts of things. You can rename all the data sets if you want to. You don't have to use "logical volumes" any more. And initial setup is done for Basic Sysplex, System Logger, WLM Goal mode, SMS, TCP/IP, and OCSF, and IBM's new License Manager. All coming your way in March, 2001.

A new look for Modify System Layout:

The first difference you'll notice about ServerPac is that the **Modify System Layout** dialog won't yell at you before you have a chance to change things!

Instead, you'll see a new **Modify System Layout Options** panel. If it's been a year since your last installation you won't have to remember what a "SUMD U" command means. You can either select from options written in English or use the commands, which remain available.

The old panel is still there, as Option P, but we think you will be more interested in using the two new options at the top.

The first one, *Create a Recommended System Layout*, will assign target and DLIB data sets to volumes automatically. The second, *View and change data sets by selected attributes*, is a powerful new facility for making changes to data sets.

View and change — Slicing and dicing the data set Lists

Let's take a look at the new *View and change* option first. Before, you could group data sets only by assigning them to logical volumes. Then, you could assign the logical volumes to physical volumes. Now, the you can create a display using any of 28 attributes so changes to things like volume assignment can be done much more quickly. You select values within each attribute to get exactly the list you want, and you can save the each list in a file. Here is the list of attributes:

- Whether the data set is to be allocated
- BLKSIZE
- Current data set size in cylinders
- DDDEF name
- DSORG
- DSNTYPE (HFS, LIBRARY, PDS)
- Data set element type (LMOD, PNLENU, EXEC, etc.)
- Currently assigned logical volume
- Must be in master catalog (Y or N)
- Device type
- Currently assigned physical volume
- RECFM
- LRECL
- Whether data set is new in this order - Whether rename is allowed (Yes, No, or Overridden)
- Data set type (Target, DLIB, or Operational)
- Volume sequence number (Tnn, Dnn, or Bnn)

- Subsystem Type (MVS, CICS, DB2, IMS, or NCP)
- Product, feature, or element name
- Special data set placement (FIRST or LAST)
- SMS-managed
- SMS-eligible
- SMS required
- Whether data set can be either a PDS or PDSE
- LPA required
- LPA eligible
- Link list eligible
- Volume Number (Tnn, Dnn, or Bnn)

There are new Change commands, too, and the Merge command now starts with the list you selected, which makes it much easier to merge like data sets. For instance, to merge all your US English panel data sets, select "Data Set Type," select "PNLENU" and "PNL," and issue the Merge line command...you're finished!

The first panel you will see when you select Option C from the **Modify System Layout Options** panel. Here, you pick one attribute to work with. For example, you might pick "Element Type" to display certain kinds of data sets. You will then see the **Select Values to Display** panel. Next, you pick the values the dialog found for that attribute.

You can pick one, some, or all the values shown. Once you have made your choices, enter the List primary command to see the data set list.

This list shows all the panel libraries in order. This is how you can quickly get to lists from which the Merge command is very useful. If issued from this panel, it would show you a "merge candidate list" of all the eligible panel libraries.

Charting the changes — The Change and Merge commands

People often ask us why, since IBM recommends no secondary space for link list libraries, ServerPac allocates them with secondary space by default. The real reason is, “Because that’s how it’s always been.” But the reason we never changed was because some people build maintenance systems for which, unlike production systems, secondary extents are desirable.

Now, you can quickly set the secondary space to zero for all the link list data sets. The CH S(pace) command has a new operand—P0—that you can use to set secondary space to zero. Using *View and change*, select “Link list eligible,” select “Yes,” and then issue the new change command (CH S P0) to set them all to zero...and you’re done! And, of course, secondary space is saved for each data set when you save a configuration for reuse with your next order.

There are other new change commands you can use to assign data sets directly to volumes (not “logical volumes,” real volumes), change the high-level qualifiers of data sets, and—lots of people asked for this one—allow all data sets to be renamed.

Here are the other new change commands:

- You can use CH PVOL to assign target, DLIB, or operational data sets directly to a volume.
- CH DSN *HLQ* changes just the high-level qualifiers of data sets.
- CH RENAME command overrides the “not renameable” attribute for data sets.

You can still issue Change commands from any data set list panel, not only from the new *View and change* panels. The Change command always starts with the list of data sets displayed, and shows you a list of candidate data sets from that list. Then you can exclude any you do not want to change and tell the dialog to apply the change to the rest of the data sets displayed.

The Merge command now starts with the data sets in the displayed list; it used to search the entire order for eligible data sets. When issued from the powerful *View and change* option, we think you will find it much more useful. (However, if you liked the way it used to work, you can just issue it from a list that includes all the data sets. For example, in *View and change* select “link list eligible” and then select both “yes” and “no” and then issue the merge command against the desired merge target data set.)

Recommended System Layout — The Automatic Data Set assignment story

Some time ago, a recommended system layout was published in *OS/390 Planning for Installation*. But it took a long time to install a ServerPac using that layout because the dialog did not lend itself well to that kind of structure. A spate of customer requirements (also called Requests) followed.

We are happy to report that with the ServerPac that becomes available at the same time as z/OS, a new function is introduced to help you get to the recommended system layout. Using this function is *much* faster than trying to create a layout manually. Most configurations take a minute or less to create.

The **Recommended System Layout** option will assign target and DLIB data sets to volumes for you according to a set of rules. But it won’t move existing data sets unless you ask it to.

You can have the dialog assign all target and DLIB data sets, to create a layout for the first time; or, have it assign only the new ones, to preserve an existing layout. You can even have it reassign some of the old ones and all the new ones. (This “Partial” option is akin to compressing a PDS; sooner or later, you’ve got to “clean up.”)

Once you pick All, New, or Partial, you see the current volume configuration.

When you’re ready, you issue the Create command to start automatic data set assignment. When it has finished, the current volume configuration is redisplayed to show you what happened.

With z/OS orders containing about 2,000 data sets, this is much faster than doing it by hand. This function leaves only the operational data sets for you to assign manually.

License Manager

There is a new component for z/OS, called IBM License Manager, or ILM. ILM will let you take advantage of workload pricing Charges. ILM requires a more robust environment, but ServerPac will help you with most of it, particularly if you use the Full-System Replacement installation path.

ServerPac setup will focus on your first system by setting up an ILM Application Agent and License Server on that image.

ServerPac will set up these functions for ILM:

- RACF - All the profiles you need for ILM and the services it uses are created by a new installation job. With some review and editing to match your installation's standards, this job should help you get a running start. Also, it defines all the new high-level qualifiers in the order for you.
- Basic Sysplex - ServerPac's default is now a Basic Sysplex configuration, also called a Monoplex.
- System Logger - System Logger is set up for License Manager...and for OPERLOG and LOGREC, too.
- SMS - SMS is set up with the SMS classes needed to support System Logger, and with minimally-intrusive ACS routines that should not interfere with the allocation of other data sets.
- WLM - The system now comes up in Goal Mode using the default WLM configuration. (You can tailor the WLM configuration with its ISPF dialog.)
- TCP/IP - Loopback communication is enabled, so the Application Agent and Server can talk to each other. There is also a sample FTP configuration.
- OCSF - This z/OS element is entirely set up and ready to use.

All these things are either done when your order is built (ILM Application Agent and Server, Basic Sysplex, WLM, OCSF, TCP/IP, and SMS), or done by new installation jobs (RACF and Logger).

If you use the Software Upgrade installation path, you might want to generate these jobs for use as samples by selecting the Full-System Replacement installation path, use GENSKEL to create the jobs, copy them, and then go back to the Software Upgrade path to perform the installation. The jobs most useful as samples include:

- ALLOCDS, which allocates SMS, Couple, and ILM data sets
- RESTORE, which restores the SMS data sets
- RACFDRV and RACFTGT, which create needed RACF profiles
- DEFNLOGS, which creates the log streams.

When you have finished with the ServerPac installation jobs, all that will be left is finishing the FTP setup (after all, we don't know what hardware you have), installing the ILM workstation code, and installing and modifying the certificates.

Giving you what you asked for

We think you'll find ServerPac is better than ever.

If you have been involved in user groups like SHARE, or follow the IBM-MAIN discussion group on the Usenet, you will probably recognize several requirements in this article. This release of ServerPac includes *a lot* of things our customers have asked for over the past few years. Keep those requests and requirements coming! •

Parallel Sysplex Update

MADLINE NICK AND DAVID RAFTEN

“May you live in exciting times.”
– *Ancient Chinese saying*

Did you ever wonder if you could select the best time that you would like to live in... would it be now, the past or the future? There are many of us who would select NOW!

These are exciting times for all, especially anyone working in the IT industry with an understanding of a web other than a place where spiders climb. The industry is moving so fast we have to stay awake 24 hours a day to not miss an opportunity, but this is not possible. So, we have to keep systems awake 24 hours and ready to change at a moments notice without warning. IBM has again stepped up to these challenges with our latest announcement of the zseries machines and z/OS operating system. There are numerous articles in this issue discussing this recent announcement. This article will focus on what has been announced as it relates to Parallel Sysplex and to our previous Hot Topics (issue 3) Newsletter .

Our last Hot Topics newsletter focused on Dynamic Resource Management and Parallel Sysplex dynamic workload balancing providing an explanation of base S/390 and OS/390 components such as Processor Resource Systems Management (PR/SM), workload Manager (WLM), I/O resource management and Parallel Sysplex functions. The articles in that newsletter gave you a strong understanding of the above base components which are the foundation on which a number of the following robust, recently announced, integration items are based.

In October, 2000, IBM announced a new processor family (zSeries) along with a new operating system, z/OS, each having some of it's core function tightly integrated with Parallel Sysplex, including the following:

- *Intelligent Resource Director:* Includes LPAR CPU Management, Dynamic Channel Path Management and Channel Subsystem Priority Queuing
- *New Stand-alone Coupling Facility:* The series 900 Model 100 with increased capacity, newer links, upgrade capability.
- *New Coupling Links:* Integrated Cluster Bus feature (ICB3), Inter-System Channel (ISC3), and Internal Coupling Channel (IC3) with “Peer Mode” capabilities.
- *GDPS / XRC:* Geographically Dispersed Parallel Sysplex (GDPS) support for the open IBM Extended Remote Copy (XRC) architecture.
- *Managed System Infrastructure for setup:* Consistent interfaces for installing and configuring products across the z/OS platform.

These items, as well as supporting the new z/Architecture, provide better system performance, provide better use of the hardware, enable simpler configuration, and ease the process of installing and managing your environment. All this combines to make z/OS the simplest environment to run your mixture of mission-critical workloads, with the lowest overall total cost of computing.

Intelligent Resource Director

IBM's exclusive Intelligent Resource Director (IRD) is made up of three functions which work

together to optimize your business' productivity and dynamically manage your CPU and I/O resources in today's ever-changing “dot.com” world:

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

When fully utilized, this exclusive function integrated the strengths of key S/390 and zseries platform leadership technologies, including z/OS Workload Manager, PR/SM (logical partitioning hardware technology), zseries Channel Subsystem and Parallel Sysplex Clustering technology. With this powerful combination, Intelligent Resource Director, *dynamically* moves resources to the application across multiple (LPAR Cluster) Images within a single server. This capability of a system to have application importance awareness across multiple images, to response to the needs of your most critical applications, is a quantum leap forward in systems management.

It enables the system to dynamically manage CPU and I/O resource and run at a higher utilization while still meeting the business goals of the most important work, which can help reduce the total cost of ownership of the solution while at the same time improve systems management. Put your business-critical application work under the management of IRD and improve your business's image and productivity as IRD moves the resources required to that application dynamically.

The z/OS Workload Manager is extended to work with PR/SM, Channel Subsystem and Parallel Sysplex on zseries servers to accomplish this distinctive single server solution. For optimum benefit, this solution requires WLM goal mode in order to dynamically manage computing and I/O resources across *multiple* logical partitions on the same series server based on the business goals of the workload running on these partitions. The scope of this management is the “LPAR Cluster” which is the set of z/OS images in the same parallel sysplex running as logical partitions on the same z900 server. 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 (in multi-system environment).

Intelligent Resource Director, or IRD, automates many of the manual planning and tuning efforts that go into maintaining these logical partitions. This is done with LPAR CPU Management to fine tune the LPAR weights and the number of logical CPs, Channel Subsystem Priority Queuing to prioritize channel access based on workload importance, and Dynamic Channel Path Management, or DCM, to automatically reconfigure the channel paths to the disk subsystem in response to changing workload demands. All this combines to:

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

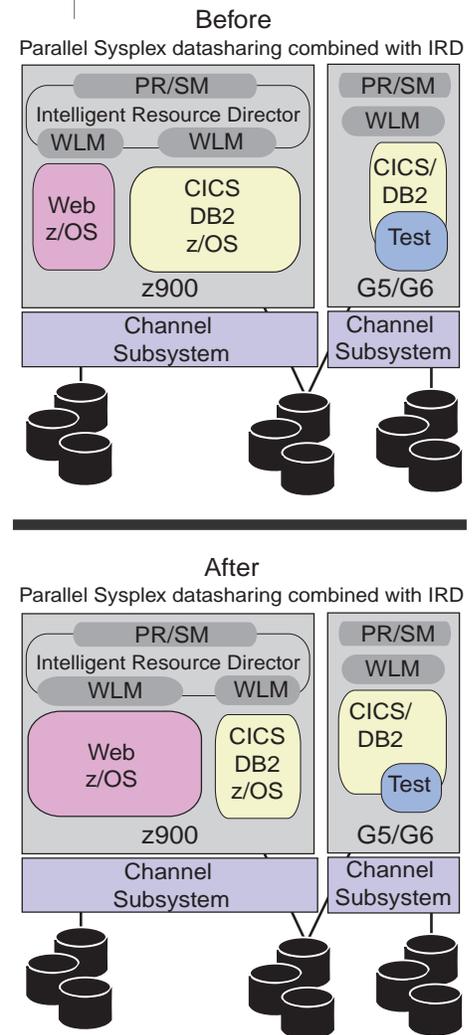
In a Parallel Sysplex with datasharing, Workload Manager directs work to the appropriate resources based on business policy. With IRD, resources are directed to the priority work. Together, Parallel Sysplex and IRD provide e-business workloads with a flexibility and responsiveness that is unrivaled in the industry!

Intelligent Resource Director together with Dynamic Workload Balancing

The purpose of IRD is to move CPU and I/O resources to where they are needed within a single server based on workload demand and business goals. The purpose of data sharing with dynamic workload balancing is to move transactions to the partition where they can obtain the best availability and responsiveness. While “responsiveness” depends on CPU and I/O resources, the two work hand in hand to provide availability and reduce system management activity. Keep in mind that IRD does not provide highest availability for software and many hardware outages. Data sharing with dynamic workload balancing does. IRD scope is limited to a single server while data sharing can span multiple servers. Finally, the combination of IRD and Parallel Sysplex datasharing optimizes the capability of running multiple heterogeneous workloads: take the following example; the first figure on the right represents the “before” view of a Parallel Sysplex with 2 servers, one server running IRD in the same Parallel Sysplex with CICS/DB2 datasharing.

Moving over to the second figure, you see a web campaign surge requiring more CPU and I/O

resources to meet its goals. With IRD the web application running on the z900 will be able to get almost all the resources on that server to handle the web campaign while not jeopardizing the CICS/DB2 workload. This is due to the fact that CPSM (CICSplex® Systems Manager) will just route the incoming CICS requests to the



G5/G6 server in the Sysplex, thinning down the LPAR on the z900 so that the web application meets its goals. As long as there is another server to handle the CICS work anywhere in the datasharing sysplex and the z900 is large enough to handle the campaign (with CUoD can grow to 16 way), everyone is happy.

For more on IRD, refer to Steve Hamilton and Madeline Nick's article in this issue.

IBM zSeries 900 Model 100 Coupling Facility

The centerpiece of the Parallel Sysplex is the Coupling Facility. To help make the times exciting, we have a new model.

The z900 Model 100 stand-alone Coupling Facility (CF) offers the industry's most powerful Coupling Facility with almost double the capacity of the 9672-R06. It supports up to nine processors, 15 logical partitions, with memory of up to 32 GB. But that is not all. The zSeries Model 100 also supports the new CF links to run in "Peer Mode" or the previous "Compatibility Mode". Up to 32 CF links can be configured on the Model 100 Coupling Facility.

The Model 100 CF can be upgraded to from a 9672-R06 and is upgradeable to a z900 server model. This provides protection of investment in IBM's Parallel Sysplex coupling technology and makes it easy to grow your 9672-R06.

New Coupling Facility links

The excitement builds with new Coupling Facility links, the means to connect to the new CFs. Instead of just "ISC" (InterSystem Communication HiPerLinks), ICB, or IC links, you can now configure "ISC3", "ICB3", or "IC3" links. While the ISC3 (when configured in compatibility mode) and the ICB2 links are used to communicate with 9672s in Compatibility Mode, the new ISC3 (when configured in peer mode) and the ICB3 links can be used in Peer Mode between two z900 servers. Peer mode provides optimal system performance and flexibility while potentially decreasing your overall link requirements.

Peer Mode Links support coupling between z900 servers and provide both sender and receiver capability on the same link. With these design improvements, the new Peer Mode ISC3 connections perform at 200 MBytes/sec, twice the former ISC capability of 100 Mbytes/sec. The new Peer Mode ICB3 connections perform at a theoretical peak of one GBytes/sec, effective rate of 800 MB/Sec. This is up from the former ICB capability measured at 250 MBytes/sec. With these faster links comes in increased coupling efficiency. Less time in the link means less delays in the processor and so, better system performance. Finally, when configuring links from an OS/390 or z/OS to an ICF on the same z900 server you have the option of configuring IC-3 Internal Coupling links. This provides approximately 1.25 Gbytes/second capacity.

But with these new peer links we didn't stop with just capacity and speed improvements, we also made protocol changes, increased the number of CF link message buffers, and the size of these buffers. These enhancements further improve your Coupling Facility responsiveness, especially at longer distances. AND we gave you the ability to use a Peer Mode link as a "CF Sender" and "CF Receiver" at the same time to save on hardware that needs to be configured.

Connections from a z900 to a 9672 must be in "Compatibility Mode" with the same characteristics as the standard 9672 ISC and ICB links.

The z900 supports a maximum of 32 ISC and/or ISC3 links and 16 ICB and/or ICB3 links with a total of these external links not exceeding 32. In addition, up to 16 IC and/or IC3 links can be configured.

GDPS / XRC

One of the most exciting (not good exciting) things that could happen to any computer installation is a disaster. For years, IBM has had the means to minimize the impact of this through a combination of hardware, software, and automation support in the Geographically Dispersed Parallel Sysplex (GDPS)[™].

Geographically Dispersed Parallel Sysplex (GDPS) complements a multisite Parallel Sysplex by providing a single, automated solution to dynamically manage storage subsystem mirroring, processors, and network resources to allow a business to attain continuous availability and near transparent business continuity/disaster recovery without data loss. GDPS provides the ability to perform a controlled site switch for both planned and unplanned site outages, while maintaining full data integrity across multiple storage subsystems.

Geographically Dispersed Parallel Sysplex (GDPS) has been enhanced to include support for the open IBM Extended Remote Copy (XRC) architecture. GDPS/XRC, in addition to the already available GDPS/PPRC support, based on the open IBM Peer-to-Peer Remote Copy (PPRC) architecture. GDPS/XRC is targeted for enterprises that:

- Need to copy data over extended distances (greater than 40 km) between data centers. This may be required either because there is a need to recover from regional disasters versus metropolitan disasters or because the customer has data centers already available at these extended distances.
- Do not have dark fiber available between data centers. GDPS, based on the synchronous

form of remote copy, PPRC, requires dark fiber between sites. GDPS/XRC, on the other hand, can be performed with communication links and channel extender technology.

- Need to maintain the highest levels of performance on their application host system.
- With XRC, the I/O operation is considered completed as soon as the data is hardened in the primary storage control unit. Subsequently, the data is copied asynchronously from the primary to the secondary storage control units, thereby not impacting the performance.
- Require a business continuance solution, instead of a continuous availability solution.
- Can tolerate a minimal amount of data loss (typically in the order of a few seconds).

Further information on GDPS can be found at <http://ibm.com/s390/pso>

Managed System Infrastructure for Setup makes z/OS easier to configure

We have new hardware, new links, and new functionality. We now help you with the installation process to get started on this using Managed System Infrastructure (msys).

With msys,

- Updates to the system are done automatically
- The system takes care of the syntax and interfaces
- The system does the calculations
- Only a small number of values is shown for which a decision is required
- Validity of a configuration is checked by the system
- On-line help is just a mouse-click away

msys will be exploited in z/OS R1 by Parallel Sysplex technology to allow setting up a Parallel Sysplex cluster much more quickly

and easily by defining and implementing the policies, parmlib specifications, and security settings necessary to configure a basic Parallel Sysplex environment.

This initial delivery focuses on Resource Sharing configurations, including support for Intelligent Resource Direct, XCF signaling and global resource serialization Star. To make the setup of msys itself easier, the z/OS LDAP server will provide executable scripts for an easy configuration as an z/OS Management Directory.

For more detail on msys for Setup, read Jose Castano's article in this issue.

While all of these changes do make the times new and exciting, it all combines to make your job a little bit LESS exciting! •

IBM's strategy for software service

STEPHEN SAROKA

S/390 Service Update Facility (SUF) is IBM's strategic Internet solution for obtaining S/390 software service. SUF is an Internet-based S/390 software service tool that makes ordering and receiving z/OS, OS/390, zVM, VM, and VSE software service quick and easy. With SUF, your system programmers can order both corrective and preventive service using the the Internet for delivery over the Internet. You can also obtain service on standard physical media if you still need it (of course, VSE preventive service, that is, a system refresh, can only be delivered on physical media). The SUF "Customer Application

Server" provides a "GUI interface" and can support multiple systems programmers supporting multiple S/390 servers. You can attach a mixture of OS/390, zOS, VM, zVM and VSE servers to a single SUF server simultaneously.



Kerry Colliia

The enhancements planned for SUF Version 3 Release 1, which will be available in March 2001, are:

- Simplified communications through your firewall using HTTP rather than FTP
- Streamlined configuration (the WebSphere Application Server is no longer required)
- A new and improved user interface
- SUF support on Linux/390 and Windows 2000

Remember these key points - SUF:

- Exploits the Internet for S/390 software service distribution
- Offers a graphical user interface based on Web browser technology
- Make use of the software service history of your systems
- Builds customized packages of just the service your system require
- Delivers preventive and corrective service
- Informs you about the status of your order
- Provides ++Enhanced Holddata with every OS/390 and z/OS service order

Why Is SUF valuable for you?

The S/390 Service Update Facility can be of value to you because:

- It provides an easier, faster way to order, receive, and install service.
- It reduces system programmer, research time and effort through a simplified software service maintenance process.
- It operates across multiple S/390 systems with a consistent Web-browser GUI that's easier to learn, use, and deploy within your enterprise.
- It provides better service, which results in more stable systems and less rediscovery of known defects.

- It lets your system programmers spend more time enhancing and exploiting your systems and less time maintaining them. (This is especially attractive to smaller businesses with limited system programming expertise, and resource.)
- It increases support staff productivity.
- It increases your satisfaction with IBM Service while reducing the size of your preventive maintenance.

For details regarding prerequisites, entitlement, and how to obtain SUF, visit the SUF home page at:

<http://www.ibm.com/s390/suf/> •

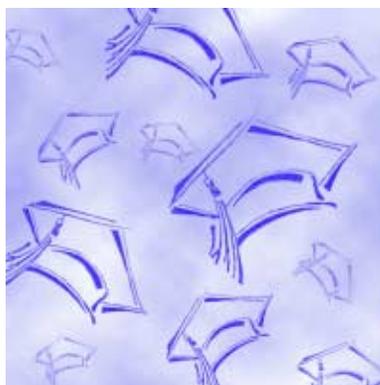
New courses!

REE HOWARD

We have 3 courses for the new zSeries:

zSeries 900 Architecture Planning and Implementation (OZ09)- Describes the new functions, features, planning and sysplex migration requirements for the z900 processor.

Sysplex migration to the zSeries 900 processors (OZ40) - Covers an Overview and Introduction of the new processor and architecture and includes new naming conventions, models, terminology, features and architecture, as well as discussing in detail connectivity concerns and requirements connected with a migration from an existing parallel sysplex to one with the new freeway processors without a sysplex wide outage.



Kim Mulé

License Manager Implementation for z/OS (OZ03) - Introduces the new License Manager base element within z/OS and how to implement these functions to simplify and centralize software usage billing for z/OS.

We have 1 course for Linux for S/390:

Linux for S/390 Implementation (ES04) - Learn how to install Linux on S/390 and dramatically increase the number of applications available to your S/390 hardware. From hand-held devices to the mainframe, Linux has grown up. The flexibility and openness of Linux when combined with S/390's quality of service provides a whole new range of capabilities. This course will cover Linux/390 distributions, installation, administration, access to OS/390 data, and prepare you to exploit the newest S/390 operating system.

For more information about our courses and/or to enroll, call 1-800-IBM-TEACH (426-8322) or visit: <http://ibm.com/services/learning/spotlight/s390.html> •

Our contributors

Franz-Peter Boley — Franz-Peter has been with IBM for eighteen years and has held a variety of software development positions within the IBM lab in Boeblingen, Germany. He is currently working on the marketing side for z/OS msys.

Susan Buley — Susan is a staff software engineer at IBM Poughkeepsie, has written information for several of the OS/390 based products including RACF, MVS, JES3, and OS/390 UNIX in her fourteen years at IBM. Susan is currently working on the z/OS UNIX System Services information, the XML Toolkit for OS/390 web site, as well as Linux for zSeries information.

Paul Carroll — Paul is a senior planner with over thirty-two years experience with IBM, the last ten years working with independent software vendors. Most of his career has been spent in development and management within mainframe computing, including an overseas assignment with CICS in Hursley. He currently manages the z/OS technical liaison (TLC) program.

Thomas Cierech — Tom is a senior software engineer at the IBM Poughkeepsie Lab. Tom's recent assignments include: designer of Measured Usage License Charges and S/390 Usage Pricing and software designer for Capacity Upgrade on Demand. Tom has been the S/390 design rep for the GUIDE/SHARE/Open Group's XSLM specification since October of 1996.

John Czukkermann — John is an advisory engineer working on the IBM License Manager team in Poughkeepsie, NY. His primary responsibility is to lead defining and managing the ILM Pre-Production Period for z/OS V1 R1.

Greg Dyck — Greg is a senior software engineer in the z/OS development organization. He is responsible for the design and implementation of system functions in many areas of z/OS.

John DiMartino — John is a student at the University of Nebraska, Omaha, studying Management Information Systems. John is currently working as a co-op at IBM Poughkeepsie. John's assignment is to write information for the Linux for zSeries product including these articles and items for the Linux for S/390 web site at <http://www.ibm.com/s390/linux/>.

John Eells — John is a former MVS and OS/390 system programmer, now a senior software engineer in IBM Software Delivery and Fulfillment. SDF delivers all IBM's software and service packages including ServerPac, CBPDQ, ESO, and SUF. His current assignments include ServerPac Design.

Jim Ground — Jim is a staff software engineer working in OS/390 and z/OS information development. He has six years of experience in technical writing and support. He holds an MS degree in Technical Communication from Rensselaer Polytechnic Institute. He has written extensively on MVS recoverable resource management services, as well as other elements and features of OS/390 and z/OS.

Steve Hamilton — Steve is an advisory information developer in the z/OS organization. He has supported Workload Management since 1996.

Madeline Nick — Madeline is a senior software engineer in the Parallel Sysplex Product Development team in Poughkeepsie, NY.

Vicky Pollien — Vicky is an advisory software engineer on the OS/390, z/OS PDT. She was the project leader for the reversion from OS/390 to z/OS, and has been working exclusively on WLC and ILM for the past year. Her recent past experience includes Problem and Project Management of a variety of key projects.

David Raften — David has been in the Parallel Sysplex Competency Center since 1994. He has a background in performance testing, DB2, IMS, and batch workloads.

William Rooney — William is a senior software engineer in the zSeries Software System Design organization where he is responsible for design of I/O functions. He is currently the lead designer for Dynamic Channel Path Management.

J.D. Ross — J.D. is a Software Engineer at IBM Poughkeepsie and has worked on documentation for OS/390 UNIX System Services and Linux for S/390. He is currently involved in documentation for z/OS UNIX System Services and is the webmaster of the Linux for S/390 website. Before joining IBM, J.D. served for several years as a system administrator, working with both the UNIX and Linux operating systems.

John Urbanic — John is an advisory information developer with eighteen years experience writing IBM software manuals. Although John can not take credit for creating the Internet, he did write one of IBM's first softcopy-only Web books in 1992.

Gail Whistance — Gail is an advisory programmer on the z/OS Workload Management team responsible for algorithm verification of new releases and for customer support.